

Book: Probability, Mathematical Statistics,
Stochastic Processes

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (<https://LibreTexts.org>) and like the hundreds of other texts available within this powerful platform, it is freely available for reading, printing and "consuming." Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects.

Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.



The LibreTexts mission is to unite students, faculty and scholars in a cooperative effort to develop an easy-to-use online platform for the construction, customization, and dissemination of OER content to reduce the burdens of unreasonable textbook costs to our students and society. The LibreTexts project is a multi-institutional collaborative venture to develop the next generation of open-access texts to improve postsecondary education at all levels of higher learning by developing an Open Access Resource environment. The project currently consists of 14 independently operating and interconnected libraries that are constantly being optimized by students, faculty, and outside experts to supplant conventional paper-based books. These free textbook alternatives are organized within a central environment that is both vertically (from advance to basic level) and horizontally (across different fields) integrated.

The LibreTexts libraries are Powered by [NICE CXOne](#) and are supported by the Department of Education Open Textbook Pilot Project, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptations contact info@LibreTexts.org. More information on our activities can be found via Facebook (<https://facebook.com/Libretexts>), Twitter (<https://twitter.com/libretexts>), or our blog (<http://Blog.Libretexts.org>).

This text was compiled on 03/19/2025

TABLE OF CONTENTS

Introduction

Licensing

Object Library

Credits

Sources and Resources

1: Foundations

- 1.1: Sets
- 1.2: Functions
- 1.3: Relations
- 1.4: Partial Orders
- 1.5: Equivalence Relations
- 1.6: Cardinality
- 1.7: Counting Measure
- 1.8: Combinatorial Structures
- 1.9: Topological Spaces
- 1.10: Metric Spaces
- 1.11: Measurable Spaces
- 1.12: Special Set Structures

2: Probability Spaces

- 2.1: Random Experiments
- 2.2: Events and Random Variables
- 2.3: Probability Measures
- 2.4: Conditional Probability
- 2.5: Independence
- 2.6: Convergence
- 2.7: Measure Spaces
- 2.8: Existence and Uniqueness
- 2.9: Probability Spaces Revisited
- 2.10: Stochastic Processes
- 2.11: Filtrations and Stopping Times

3: Distributions

- 3.1: Discrete Distributions
- 3.2: Continuous Distributions
- 3.3: Mixed Distributions
- 3.4: Joint Distributions
- 3.5: Conditional Distributions
- 3.6: Distribution and Quantile Functions
- 3.7: Transformations of Random Variables
- 3.8: Convergence in Distribution
- 3.9: General Distribution Functions

- 3.10: The Integral With Respect to a Measure
- 3.11: Properties of the Integral
- 3.12: General Measures
- 3.13: Absolute Continuity and Density Functions
- 3.14: Function Spaces

4: Expected Value

- 4.1: Definitions and Basic Properties
- 4.2: Additional Properties
- 4.3: Variance
- 4.4: Skewness and Kurtosis
- 4.5: Covariance and Correlation
- 4.6: Generating Functions
- 4.7: Conditional Expected Value
- 4.8: Expected Value and Covariance Matrices
- 4.9: Expected Value as an Integral
- 4.10: Conditional Expected Value Revisited
- 4.11: Vector Spaces of Random Variables
- 4.12: Uniformly Integrable Variables
- 4.13: Kernels and Operators

5: Special Distributions

- 5.1: Location-Scale Families
- 5.2: General Exponential Families
- 5.3: Stable Distributions
- 5.4: Infinitely Divisible Distributions
- 5.5: Power Series Distributions
- 5.6: The Normal Distribution
- 5.7: The Multivariate Normal Distribution
- 5.8: The Gamma Distribution
- 5.9: Chi-Square and Related Distribution
- 5.10: The Student t Distribution
- 5.11: The F Distribution
- 5.12: The Lognormal Distribution
- 5.13: The Folded Normal Distribution
- 5.14: The Rayleigh Distribution
- 5.15: The Maxwell Distribution
- 5.16: The Lévy Distribution
- 5.17: The Beta Distribution
- 5.18: The Beta Prime Distribution
- 5.19: The Arcsine Distribution
- 5.20: General Uniform Distributions
- 5.21: The Uniform Distribution on an Interval
- 5.22: Discrete Uniform Distributions
- 5.23: The Semicircle Distribution
- 5.24: The Triangle Distribution
- 5.25: The Irwin-Hall Distribution
- 5.26: The U-Power Distribution
- 5.27: The Sine Distribution
- 5.28: The Laplace Distribution
- 5.29: The Logistic Distribution

- 5.30: The Extreme Value Distribution
- 5.31: The Hyperbolic Secant Distribution
- 5.32: The Cauchy Distribution
- 5.33: The Exponential-Logarithmic Distribution
- 5.34: The Gompertz Distribution
- 5.35: The Log-Logistic Distribution
- 5.36: The Pareto Distribution
- 5.37: The Wald Distribution
- 5.38: The Weibull Distribution
- 5.39: Benford's Law
- 5.40: The Zeta Distribution
- 5.41: The Logarithmic Series Distribution

6: Random Samples

- 6.1: Introduction
- 6.2: The Sample Mean
- 6.3: The Law of Large Numbers
- 6.4: The Central Limit Theorem
- 6.5: The Sample Variance
- 6.6: Order Statistics
- 6.7: Sample Correlation and Regression
- 6.8: Special Properties of Normal Samples

7: Point Estimation

- 7.1: Estimators
- 7.2: The Method of Moments
- 7.3: Maximum Likelihood
- 7.4: Bayesian Estimation
- 7.5: Best Unbiased Estimators
- 7.6: Sufficient, Complete and Ancillary Statistics

8: Set Estimation

- 8.1: Introduction to Set Estimation
- 8.2: Estimation the Normal Model
- 8.3: Estimation in the Bernoulli Model
- 8.4: Estimation in the Two-Sample Normal Model
- 8.5: Bayesian Set Estimation

9: Hypothesis Testing

- 9.1: Introduction to Hypothesis Testing
- 9.2: Tests in the Normal Model
- 9.3: Tests in the Bernoulli Model
- 9.4: Tests in the Two-Sample Normal Model
- 9.5: Likelihood Ratio Tests
- 9.6: Chi-Square Tests

10: Geometric Models

- 10.1: Buffon's Problems
- 10.2: Bertrand's Paradox
- 10.3: Random Triangles

11: Bernoulli Trials

- 11.1: Introduction to Bernoulli Trials
- 11.2: The Binomial Distribution
- 11.3: The Geometric Distribution
- 11.4: The Negative Binomial Distribution
- 11.5: The Multinomial Distribution
- 11.6: The Simple Random Walk
- 11.7: The Beta-Bernoulli Process

12: Finite Sampling Models

- 12.1: Introduction to Finite Sampling Models
- 12.2: The Hypergeometric Distribution
- 12.3: The Multivariate Hypergeometric Distribution
- 12.4: Order Statistics
- 12.5: The Matching Problem
- 12.6: The Birthday Problem
- 12.7: The Coupon Collector Problem
- 12.8: Pólya's Urn Process
- 12.9: The Secretary Problem

13: Games of Chance

- 13.1: Introduction to Games of Chance
- 13.2: Poker
- 13.3: Simple Dice Games
- 13.4: Craps
- 13.5: Roulette
- 13.6: The Monty Hall Problem
- 13.7: Lotteries
- 13.8: The Red and Black Game
- 13.9: Timid Play
- 13.10: Bold Play
- 13.11: Optimal Strategies

14: The Poisson Process

- 14.1: Introduction to the Poisson Process
- 14.2: The Exponential Distribution
- 14.3: The Gamma Distribution
- 14.4: The Poisson Distribution
- 14.5: Thinning and Superposition
- 14.6: Non-homogeneous Poisson Processes
- 14.7: Compound Poisson Processes
- 14.8: Poisson Processes on General Spaces

15: Renewal Processes

- 15.1: Introduction
- 15.2: Renewal Equations
- 15.3: Renewal Limit Theorems
- 15.4: Delayed Renewal Processes
- 15.5: Alternating Renewal Processes
- 15.6: Renewal Reward Processes

16: Markov Processes

- 16.1: Introduction to Markov Processes
- 16.2: Potentials and Generators for General Markov Processes
- 16.3: Introduction to Discrete-Time Chains
- 16.4: Transience and Recurrence for Discrete-Time Chains
- 16.5: Periodicity of Discrete-Time Chains
- 16.6: Stationary and Limiting Distributions of Discrete-Time Chains
- 16.7: Time Reversal in Discrete-Time Chains
- 16.8: The Ehrenfest Chains
- 16.9: The Bernoulli-Laplace Chain
- 16.10: Discrete-Time Reliability Chains
- 16.11: Discrete-Time Branching Chain
- 16.12: Discrete-Time Queuing Chains
- 16.13: Discrete-Time Birth-Death Chains
- 16.14: Random Walks on Graphs
- 16.15: Introduction to Continuous-Time Markov Chains
- 16.16: Transition Matrices and Generators of Continuous-Time Chains
- 16.17: Potential Matrices
- 16.18: Stationary and Limiting Distributions of Continuous-Time Chains
- 16.19: Time Reversal in Continuous-Time Chains
- 16.20: Chains Subordinate to the Poisson Process
- 16.21: Continuous-Time Birth-Death Chains
- 16.22: Continuous-Time Queuing Chains
- 16.23: Continuous-Time Branching Chains

17: Martingales

- 17.1: Introduction to Martingales
- 17.2: Properties and Constructions
- 17.3: Stopping Times
- 17.4: Inequalities
- 17.5: Convergence
- 17.6: Backwards Martingales

18: Brownian Motion

- 18.1: Standard Brownian Motion
- 18.2: Brownian Motion with Drift and Scaling
- 18.3: The Brownian Bridge
- 18.4: Geometric Brownian Motion

[Index](#)

[Glossary](#)

[Detailed Licensing](#)

Licensing

A detailed breakdown of this resource's licensing can be found in [Back Matter/Detailed Licensing](#).

Table of Contents

<http://www.randomservices.org/random/index.html>

Random is a website devoted to probability, mathematical statistics, and stochastic processes, and is intended for teachers and students of these subjects. The site consists of an integrated set of components that includes expository text, interactive web apps, data sets, biographical sketches, and an object library. Please read the [Introduction](#) for more information about the content, structure, mathematical prerequisites, technologies, and organization of the project.

- **1: Foundations**

- [1.1: Sets](#)
- [1.2: Functions](#)
- [1.3: Relations](#)
- [1.4: Partial Orders](#)
- [1.5: Equivalence Relations](#)
- [1.6: Cardinality](#)
- [1.7: Counting Measure](#)
- [1.8: Combinatorial Structures](#)
- [1.9: Topological Spaces](#)
- [1.10: Metric Spaces](#)
- [1.11: Measurable Spaces](#)
- [1.12: Special Set Structures](#)

- **2: Probability Spaces**

- [2.1: Random Experiments](#)
- [2.2: Events and Random Variables](#)
- [2.3: Probability Measures](#)
- [2.4: Conditional Probability](#)
- [2.5: Independence](#)
- [2.6: Convergence](#)
- [2.7: Measure Spaces](#)
- [2.8: Existence and Uniqueness](#)
- [2.9: Probability Spaces Revisited](#)
- [2.10: Stochastic Processes](#)
- [2.11: Filtrations and Stopping Times](#)

- **3: Distributions**

- [3.1: Discrete Distributions](#)
- [3.2: Continuous Distributions](#)
- [3.3: Mixed Distributions](#)
- [3.4: Joint Distributions](#)
- [3.5: Conditional Distributions](#)
- [3.6: Distribution and Quantile Functions](#)
- [3.7: Transformations of Random Variables](#)
- [3.8: Convergence in Distribution](#)
- [3.9: General Distribution Functions](#)
- [3.10: The Integral With Respect to a Measure](#)
- [3.11: Properties of the Integral](#)
- [3.12: General Measures](#)
- [3.13: Absolute Continuity and Density Functions](#)
- [3.14: Function Spaces](#)

- [4: Expected Value](#)
 - [4.1: New Page](#)
 - [4.2: New Page](#)
 - [4.3: New Page](#)
 - [4.4: New Page](#)
 - [4.5: New Page](#)
 - [4.6: New Page](#)
 - [4.7: New Page](#)
 - [4.8: New Page](#)
 - [4.9: New Page](#)
 - [4.10: New Page](#)
- [5: Special Distributions](#)
 - [5.1: New Page](#)
 - [5.2: New Page](#)
 - [5.3: New Page](#)
 - [5.4: New Page](#)
 - [5.5: New Page](#)
 - [5.6: New Page](#)
 - [5.7: New Page](#)
 - [5.8: New Page](#)
 - [5.9: New Page](#)
 - [5.10: New Page](#)
- [6: Random Samples](#)
 - [6.1: New Page](#)
 - [6.2: New Page](#)
 - [6.3: New Page](#)
 - [6.4: New Page](#)
 - [6.5: New Page](#)
 - [6.6: New Page](#)
 - [6.7: New Page](#)
 - [6.8: New Page](#)
 - [6.9: New Page](#)
 - [6.10: New Page](#)
- [7: Point Estimation](#)
 - [7.1: New Page](#)
 - [7.2: New Page](#)
 - [7.3: New Page](#)
 - [7.4: New Page](#)
 - [7.5: New Page](#)
 - [7.6: New Page](#)
 - [7.7: New Page](#)
 - [7.8: New Page](#)
 - [7.9: New Page](#)
 - [7.10: New Page](#)
- [8: Set Estimation](#)
 - [8.1: New Page](#)
 - [8.2: New Page](#)
 - [8.3: New Page](#)

- [8.4: New Page](#)
 - [8.5: New Page](#)
 - [8.6: New Page](#)
 - [8.7: New Page](#)
 - [8.8: New Page](#)
 - [8.9: New Page](#)
 - [8.10: New Page](#)
- [9: Hypothesis Testing](#)
 - [9.1: New Page](#)
 - [9.2: New Page](#)
 - [9.3: New Page](#)
 - [9.4: New Page](#)
 - [9.5: New Page](#)
 - [9.6: New Page](#)
 - [9.7: New Page](#)
 - [9.8: New Page](#)
 - [9.9: New Page](#)
 - [9.10: New Page](#)
- [10: Geometric Models](#)
 - [10.1: New Page](#)
 - [10.2: New Page](#)
 - [10.3: New Page](#)
 - [10.4: New Page](#)
 - [10.5: New Page](#)
 - [10.6: New Page](#)
 - [10.7: New Page](#)
 - [10.8: New Page](#)
 - [10.9: New Page](#)
 - [10.10: New Page](#)
- [11: Bernoulli Trials](#)
 - [11.1: New Page](#)
 - [11.2: New Page](#)
 - [11.3: New Page](#)
 - [11.4: New Page](#)
 - [11.5: New Page](#)
 - [11.6: New Page](#)
 - [11.7: New Page](#)
 - [11.8: New Page](#)
 - [11.9: New Page](#)
 - [11.10: New Page](#)
- [12: Finite Sampling Models](#)
 - [12.1: New Page](#)
 - [12.2: New Page](#)
 - [12.3: New Page](#)
 - [12.4: New Page](#)
 - [12.5: New Page](#)
 - [12.6: New Page](#)
 - [12.7: New Page](#)
 - [12.8: New Page](#)

- [12.9: New Page](#)
- [12.10: New Page](#)
- [13: Games of Chance](#)
 - [13.1: New Page](#)
 - [13.2: New Page](#)
 - [13.3: New Page](#)
 - [13.4: New Page](#)
 - [13.5: New Page](#)
 - [13.6: New Page](#)
 - [13.7: New Page](#)
 - [13.8: New Page](#)
 - [13.9: New Page](#)
 - [13.10: New Page](#)
- [14: The Poisson Process](#)
 - [14.1: New Page](#)
 - [14.2: New Page](#)
 - [14.3: New Page](#)
 - [14.4: New Page](#)
 - [14.5: New Page](#)
 - [14.6: New Page](#)
 - [14.7: New Page](#)
 - [14.8: New Page](#)
 - [14.9: New Page](#)
 - [14.10: New Page](#)
- [15: Renewal Processes](#)
 - [15.1: New Page](#)
 - [15.2: New Page](#)
 - [15.3: New Page](#)
 - [15.4: New Page](#)
 - [15.5: New Page](#)
 - [15.6: New Page](#)
 - [15.7: New Page](#)
 - [15.8: New Page](#)
 - [15.9: New Page](#)
 - [15.10: New Page](#)
- [16: Markov Processes](#)
 - [16.1: New Page](#)
 - [16.2: New Page](#)
 - [16.3: New Page](#)
 - [16.4: New Page](#)
 - [16.5: New Page](#)
 - [16.6: New Page](#)
 - [16.7: New Page](#)
 - [16.8: New Page](#)
 - [16.9: New Page](#)
 - [16.10: New Page](#)
-

17: Martingales

- [17.1: New Page](#)
- [17.2: New Page](#)
- [17.3: New Page](#)
- [17.4: New Page](#)
- [17.5: New Page](#)
- [17.6: New Page](#)
- [17.7: New Page](#)
- [17.8: New Page](#)
- [17.9: New Page](#)
- [17.10: New Page](#)

- [18: Brownian Motion](#)

Credits

Sources and Resources

CHAPTER OVERVIEW

1: Foundations

In this chapter we review several mathematical topics that form the foundation of probability and mathematical statistics. These include the algebra of sets and functions, general relations with special emphasis on equivalence relations and partial orders, counting measure, and some basic combinatorial structures such as permutations and combinations. We also discuss some advanced topics from topology and measure theory. You may wish to review the topics in this chapter as the need arises.

[1.1: Sets](#)

[1.2: Functions](#)

[1.3: Relations](#)

[1.4: Partial Orders](#)

[1.5: Equivalence Relations](#)

[1.6: Cardinality](#)

[1.7: Counting Measure](#)

[1.8: Combinatorial Structures](#)

[1.9: Topological Spaces](#)

[1.10: Metric Spaces](#)

[1.11: Measurable Spaces](#)

[1.12: Special Set Structures](#)

This page titled [1: Foundations](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.1: Sets

Set theory is the foundation of probability and statistics, as it is for almost every branch of mathematics.

Sets and subsets

In this text, sets and their elements are primitive, self-evident concepts, an approach that is sometimes referred to as *naive set theory*.

A *set* is simply a collection of objects; the objects are referred to as *elements* of the set. The statement that x is an element of set S is written $x \in S$, and the negation that x is not an element of S is written as $x \notin S$. By definition, a set is completely determined by its elements; thus sets A and B are *equal* if they have the same elements:

$$A = B \text{ if and only if } x \in A \iff x \in B \quad (1.1.1)$$

Our next definition is the subset relation, another very basic concept.

If A and B are sets then A is a *subset* of B if every element of A is also an element of B :

$$A \subseteq B \text{ if and only if } x \in A \implies x \in B \quad (1.1.2)$$

Concepts in set theory are often illustrated with small, schematic sketches known as *Venn diagrams*, named for John Venn. The Venn diagram in the picture below illustrates the subset relation.




Figure 1.1.1: $A \subseteq B$

As noted earlier, membership is a primitive, undefined concept in naive set theory. However, the following construction, known as *Russell's paradox*, after the mathematician and philosopher Bertrand Russell, shows that we cannot be too cavalier in the construction of sets.

Let R be the set of all sets A such that $A \notin A$. Then $R \in R$ if and only if $R \notin R$.

Proof

The contradiction follows from the definition of R : If $R \in R$, then by definition, $R \notin R$. If $R \notin R$, then by definition, $R \in R$. The net result, of course, is that R is not a well-defined set.

Usually, the sets under discussion in a particular context are all subsets of a well-defined, specified set S , often called a *universal set*. The use of a universal set prevents the type of problem that arises in Russell's paradox. That is, if S is a given set and $p(x)$ is a *predicate* on S (that is, a valid mathematical statement that is either true or false for each $x \in S$), then $\{x \in S : p(x)\}$ is a valid subset of S . Defining a set in this way is known as *predicate form*. The other basic way to define a set is simply by listing its elements; this method is known as *list form*.

In contrast to a universal set, the *empty set*, denoted \emptyset , is the set with no elements.

$\emptyset \subseteq A$ for every set A .

Proof

$\emptyset \subseteq A$ means that $x \in \emptyset \implies x \in A$. Since the premise is false, the implication is true.

One step up from the empty set is a set with just one element. Such a set is called a *singleton set*. The subset relation is a partial order on the collection of subsets of S .

Suppose that A , B and C are subsets of a set S . Then

1. $A \subseteq A$ (the *reflexive property*).
2. If $A \subseteq B$ and $B \subseteq A$ then $A = B$ (the *anti-symmetric property*).
3. If $A \subseteq B$ and $B \subseteq C$ then $A \subseteq C$ (the *transitive property*).

Here are a couple of variations on the subset relation.

Suppose that A and B are sets.

1. If $A \subseteq B$ and $A \neq B$, then A is a *strict subset* of B and we sometimes write $A \subset B$.
2. If $\emptyset \subset A \subset B$, then A is called a *proper subset* of B .

The collection of all subsets of a given set frequently plays an important role, particularly when the given set is the universal set.

If S is a set, then the set of all subsets of S is known as the *power set* of S and is denoted $\mathcal{P}(S)$.

Special Sets

The following special sets are used throughout this text. Defining them will also give us practice using list and predicate form.

Special Sets

1. \mathbb{R} denotes the set of *real numbers* and is the universal set for the other subsets in this list.
2. $\mathbb{N} = \{0, 1, 2, \dots\}$ is the set of *natural numbers*
3. $\mathbb{N}_+ = \{1, 2, 3, \dots\}$ is the set of *positive integers*
4. $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ is the set of *integers*
5. $\mathbb{Q} = \{m/n : m \in \mathbb{Z} \text{ and } n \in \mathbb{N}_+\}$ is the set of *rational numbers*
6. $\mathbb{A} = \{x \in \mathbb{R} : p(x) = 0 \text{ for some polynomial } p \text{ with integer coefficients}\}$ is the set of *algebraic numbers*.

Note that $\mathbb{N}_+ \subset \mathbb{N} \subset \mathbb{Z} \subset \mathbb{Q} \subset \mathbb{A} \subset \mathbb{R}$. We will also occasionally need the set of *complex numbers* $\mathbb{C} = \{x + iy : x, y \in \mathbb{R}\}$ where i is the imaginary unit. The following special rational numbers turn out to be useful for various constructions.

For $n \in \mathbb{N}$, a rational number of the form $j/2^n$ where $j \in \mathbb{Z}$ is odd is a *dyadic rational* (or *binary rational*) of rank n .

1. For $n \in \mathbb{N}$, the set of dyadic rationals of rank n or less is $\mathbb{D}_n = \{j/2^n : j \in \mathbb{Z}\}$.
2. The set of all dyadic rationals is $\mathbb{D} = \{j/2^n : j \in \mathbb{Z} \text{ and } n \in \mathbb{N}\}$.

Note that $\mathbb{D}_0 = \mathbb{Z}$ and $\mathbb{D}_n \subset \mathbb{D}_{n+1}$ for $n \in \mathbb{N}$, and of course, $\mathbb{D} \subset \mathbb{Q}$. We use the usual notation for *intervals* of real numbers, but again the definitions provide practice with predicate notation.

Suppose that $a, b \in \mathbb{R}$ with $a < b$.

1. $[a, b] = \{x \in \mathbb{R} : a \leq x \leq b\}$. This interval is *closed*.
2. $(a, b) = \{x \in \mathbb{R} : a < x < b\}$. This interval is *open*.
3. $[a, b) = \{x \in \mathbb{R} : a \leq x < b\}$. This interval is *closed-open*.
4. $(a, b] = \{x \in \mathbb{R} : a < x \leq b\}$. This interval is *open-closed*.

The terms *open* and *closed* are actually topological concepts.

You may recall that $x \in \mathbb{R}$ is rational if and only if the decimal expansion of x either terminates or forms a repeating block. The binary rationals have simple binary expansions (that is, expansions in the base 2 number system).

A number $x \in \mathbb{R}$ is a binary rational of rank $n \in \mathbb{N}_+$ if and only if the binary expansion of x is finite, with 1 in position n (after the separator).

Proof

It suffices to consider $x \in (0, 1)$. The result is very simple so we just give the first few cases.

1. The number with rank 1 is $1/2$ with binary expansion 0.1
2. The numbers with rank 2 are $1/4$ with expansion 0.01 and $3/4$ with expansion 0.11
3. The numbers with rank 3 are $1/8$ with expansion 0.001, $3/8$ with expansion 0.011, $5/8$ with expansion 0.101, and $7/8$ with expansion 0.111.

Set Operations

We are now ready to review the basic operations of set theory. For the following definitions, suppose that A and B are subsets of a universal set, which we will denote by S .

The *union* of A and B is the set obtained by combining the elements of A and B .

$$A \cup B = \{x \in S : x \in A \text{ or } x \in B\} \quad (1.1.3)$$

The *intersection* of A and B is the set of elements common to both A and B :

$$A \cap B = \{x \in S : x \in A \text{ and } x \in B\} \quad (1.1.4)$$

If $A \cap B = \emptyset$ then A and B are *disjoint*.

So A and B are disjoint if the two sets have no elements in common.

The *set difference* of B and A is the set of elements that are in B but not in A :

$$B \setminus A = \{x \in S : x \in B \text{ and } x \notin A\} \quad (1.1.5)$$

Sometimes (particularly in older works and particularly when $A \subseteq B$), the notation $B - A$ is used instead of $B \setminus A$. When $A \subseteq B$, $B - A$ is known as *proper set difference*.

The *complement* of A is the set of elements that are *not* in A :

$$A^c = \{x \in S : x \notin A\} \quad (1.1.6)$$

Note that union, intersection, and difference are *binary* set operations, while complement is a *unary* set operation.

In the Venn diagram app, select each of the following and note the shaded area in the diagram.

1. A
2. B
3. A^c
4. B^c
5. $A \cup B$
6. $A \cap B$

Basic Rules

In the following theorems, A , B , and C are subsets of a universal set S . The proofs are straightforward, and just use the definitions and basic logic. Try the proofs yourself before reading the ones in the text.

$$A \cap B \subseteq A \subseteq A \cup B.$$

The *identity laws*:

1. $A \cup \emptyset = A$
2. $A \cap S = A$

So the empty set acts as an identity relative to the union operation, and the universal set acts as an identity relative to the intersection operation.

The *idempotent laws*:

1. $A \cup A = A$
2. $A \cap A = A$

The *complement laws*:

1. $A \cup A^c = S$
2. $A \cap A^c = \emptyset$

The *double complement law*: $(A^c)^c = A$

The *commutative laws*:

1. $A \cup B = B \cup A$
2. $A \cap B = B \cap A$

Proof

These results follow from the commutativity of the *or* and *and* logical operators.

The *associative laws*:

1. $A \cup (B \cup C) = (A \cup B) \cup C$
2. $A \cap (B \cap C) = (A \cap B) \cap C$

Proof

These results follow from the associativity of the *or* and *and* logical operators.

Thus, we can write $A \cup B \cup C$ without ambiguity. Note that x is an element of this set if and only if x is an element of at least one of the three given sets. Similarly, we can write $A \cap B \cap C$ without ambiguity. Note that x is an element of this set if and only if x is an element of all three of the given sets.

The *distributive laws*:

1. $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$
2. $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$

Proof

1. $x \in A \cap (B \cup C)$ if and only if $x \in A$ and $x \in B \cup C$ if and only if $x \in A$ and either $x \in B$ or $x \in C$ if and only if $x \in A$ and $x \in B$, or, $x \in A$ and $x \in C$ if and only if $x \in A \cap B$ or $x \in A \cap C$ if and only if $x \in (A \cap B) \cup (A \cap C)$.
2. The proof is exactly the same as (a), but with *or* and *and* interchanged.

So intersection distributes over union, and union distributes over intersection. It's interesting to compare the distributive properties of set theory with those of the real number system. If $x, y, z \in \mathbb{R}$, then $x(y+z) = (xy) + (xz)$, so multiplication distributes over addition, but it is *not true* that $x + (yz) = (x+y)(x+z)$, so addition does not distribute over multiplication. The following results are particularly important in probability theory.

DeMorgan's laws (named after Augustus DeMorgan):

1. $(A \cup B)^c = A^c \cap B^c$
2. $(A \cap B)^c = A^c \cup B^c$.

Proof

1. $x \in (A \cup B)^c$ if and only if $x \notin A \cup B$ if and only if $x \notin A$ and $x \notin B$ if and only if $x \in A^c$ and $x \in B^c$ if and only if $x \in A^c \cap B^c$
2. $x \in (A \cap B)^c$ if and only if $x \notin A \cap B$ if and only if $x \notin A$ or $x \notin B$ if and only if $x \in A^c$ or $x \in B^c$ if and only if $x \in A^c \cup B^c$

The following result explores the connections between the subset relation and the set operations.

The following statements are equivalent:

1. $A \subseteq B$
2. $B^c \subseteq A^c$
3. $A \cup B = B$
4. $A \cap B = A$
5. $A \setminus B = \emptyset$

Proof

1. Recall that $A \subseteq B$ means that $x \in A \implies x \in B$.
2. $B^c \subseteq A^c$ means that $x \notin B \implies x \notin A$. This is the contrapositive of (a) and hence is equivalent to (a).
3. If $A \subseteq B$ then clearly $A \cup B = B$. Conversely suppose $A \cup B = B$. If $x \in A$ then $x \in A \cup B$ so $x \in B$. Hence $A \subseteq B$.
4. If $A \subseteq B$ then clearly $A \cap B = A$. Conversely suppose $A \cap B = A$. If $x \in A$ then $x \in A \cap B$ and so $x \in B$. Hence $A \subseteq B$.
5. Suppose $A \subseteq B$. If $x \in A$ then $x \in B$ and so by definition, $x \notin A \setminus B$. If $x \notin A$ then again by definition, $x \notin A \setminus B$. Thus $A \setminus B = \emptyset$. Conversely suppose that $A \setminus B = \emptyset$. If $x \in A$ then $x \notin A \setminus B$ so $x \in B$. Thus $A \subseteq B$.

In addition to the [special sets](#) defined earlier, we also have the following:

More special sets

1. $\mathbb{R} \setminus \mathbb{Q}$ is the set of *irrational numbers*
2. $\mathbb{R} \setminus \mathbb{A}$ is the set of *transcendental numbers*

Since $\mathbb{Q} \subset \mathbb{A} \subset \mathbb{R}$ it follows that $\mathbb{R} \setminus \mathbb{A} \subset \mathbb{R} \setminus \mathbb{Q}$, that is, every transcendental number is also irrational.

Set difference can be expressed in terms of complement and intersection. *All* of the other set operations (complement, union, and intersection) can be expressed in terms of difference.

Results for set difference:

1. $B \setminus A = B \cap A^c$
2. $A^c = S \setminus A$
3. $A \cap B = A \setminus (A \setminus B)$
4. $A \cup B = S \setminus \{(S \setminus A) \setminus [(S \setminus A) \setminus (S \setminus B)]\}$

Proof

1. This is clear from the definition: $B \setminus A = B \cap A^c = \{x \in S : x \in B \text{ and } x \notin A\}$.
2. This follows from (a) with $B = S$.
3. Using (a), DeMorgan's law, and the distributive law, the right side is

$$A \cap (A \cap B^c)^c = A \cap (A^c \cup B) = (A \cap A^c) \cup (A \cap B) = \emptyset \cup (A \cap B) = A \cap B \quad (1.1.7)$$

4. Using (a), (b), DeMorgan's law, and the distributive law, the right side is

$$[A^c \cap (A^c \cap B)^c]^c = A \cup (A^c \cap B) = (A \cup A^c) \cap (A \cup B) = S \cap (A \cup B) = A \cup B \quad (1.1.8)$$

So in principle, we could do all of set theory using the one operation of set difference. But as (c) and (d) suggest, the results would be hideous.

$$(A \cup B) \setminus (A \cap B) = (A \setminus B) \cup (B \setminus A) \quad .$$

Proof

A direct proof is simple, but for practice let's give a proof using set algebra, in particular, DeMorgan's law, and the distributive law:

$$(A \cup B) \setminus (A \cap B) = (A \cup B) \cap (A \cap B)^c = (A \cup B) \cap (A^c \cup B^c) \quad (1.1.9)$$

$$= (A \cap A^c) \cup (B \cap A^c) \cup (A \cap B^c) \cup (B \cap B^c) \quad (1.1.10)$$

$$= \emptyset \cup (B \setminus A) \cup (A \setminus B) \cup \emptyset = (A \setminus B) \cup (B \setminus A) \quad (1.1.11)$$

The set in the previous result is called the *symmetric difference* of A and B , and is sometimes denoted $A \triangle B$. The elements of this set belong to one but not both of the given sets. Thus, the symmetric difference corresponds to *exclusive or* in the same way that union corresponds to *inclusive or*. That is, $x \in A \cup B$ if and only if $x \in A$ or $x \in B$ (or both); $x \in A \triangle B$ if and only if $x \in A$ or $x \in B$, but not both. On the other hand, the complement of the symmetric difference consists of the elements that belong to both or neither of the given sets:

$$(A \triangle B)^c = (A \cap B) \cup (A^c \cap B^c) = (A^c \cup B) \cap (B^c \cup A)$$

Proof

Again, a direct proof is simple, but let's give an algebraic proof for practice:

$$(A \triangle B)^c = [(A \cup B) \cap (A \cap B)^c]^c \quad (1.1.12)$$

$$= (A \cup B)^c \cup (A \cap B) = (A^c \cap B^c) \cup (A \cap B) \quad (1.1.13)$$

$$= (A^c \cup A) \cap (A^c \cup B) \cap (B^c \cup A) \cap (B^c \cup B) \quad (1.1.14)$$

$$= S \cap (A^c \cup B) \cap (B^c \cup A) \cap S = (A^c \cup B) \cap (B^c \cup A) \quad (1.1.15)$$

There are 16 different (in general) sets that can be constructed from two given events A and B .

Proof

S is the union of 4 pairwise disjoint sets: $A \cap B$, $A \cap B^c$, $A^c \cap B$, and $A^c \cap B^c$. If A and B are in “general position”, these 4 sets are distinct. Every set that can be constructed from A and B is a union of some (perhaps none, perhaps all) of these 4 sets. There are $2^4 = 16$ sub-collections of the 4 sets.

Open the Venn diagram app. This app lists the 16 sets that can be constructed from given sets A and B using the set operations.

1. Select each of the four subsets in the proof of the last exercise: $A \cap B$, $A \cap B^c$, $A^c \cap B$, and $A^c \cap B^c$. Note that these are disjoint and their union is S .
2. Select each of the other 12 sets and show how each is a union of some of the sets in (a).

General Operations

The operations of union and intersection can easily be extended to a finite or even an infinite collection of sets.

Definitions

Suppose that \mathcal{A} is a nonempty collection of subsets of a universal set S . In some cases, the subsets in \mathcal{A} may be naturally *indexed* by a nonempty index set I , so that $\mathcal{A} = \{A_i : i \in I\}$. (In a technical sense, any collection of subsets can be indexed.)

The *union* of the collection of sets \mathcal{A} is the set obtained by combining the elements of the sets in \mathcal{A} :

$$\bigcup \mathcal{A} = \{x \in S : x \in A \text{ for some } A \in \mathcal{A}\} \quad (1.1.16)$$

If $\mathcal{A} = \{A_i : i \in I\}$, so that the collection of sets is indexed, then we use the more natural notation:

$$\bigcup_{i \in I} A_i = \{x \in S : x \in A_i \text{ for some } i \in I\} \quad (1.1.17)$$

The *intersection* of the collection of sets \mathcal{A} is the set of elements common to all of the sets in \mathcal{A} :

$$\bigcap \mathcal{A} = \{x \in S : x \in A \text{ for all } A \in \mathcal{A}\} \quad (1.1.18)$$

If $\mathcal{A} = \{A_i : i \in I\}$, so that the collection of sets is indexed, then we use the more natural notation:

$$\bigcap_{i \in I} A_i = \{x \in S : x \in A_i \text{ for all } i \in I\} \quad (1.1.19)$$

Often the index set is an “integer interval” of \mathbb{N} . In such cases, an even more natural notation is to use the upper and lower limits of the index set. For example, if the collection is $\{A_i : i \in \mathbb{N}_+\}$ then we would write $\bigcup_{i=1}^{\infty} A_i$ for the union and $\bigcap_{i=1}^{\infty} A_i$ for the intersection. Similarly, if the collection is $\{A_i : i \in \{1, 2, \dots, n\}\}$ for some $n \in \mathbb{N}_+$, we would write $\bigcup_{i=1}^n A_i$ for the union and $\bigcap_{i=1}^n A_i$ for the intersection.

A collection of sets \mathcal{A} is *pairwise disjoint* if the intersection of any two sets in the collection is empty: $A \cap B = \emptyset$ for every $A, B \in \mathcal{A}$ with $A \neq B$.

A collection of sets \mathcal{A} is said to *partition* a set B if the collection \mathcal{A} is pairwise disjoint and $\bigcup \mathcal{A} = B$.

Partitions are intimately related to equivalence relations. As an example, for $n \in \mathbb{N}$, the set

$$\mathcal{D}_n = \left\{ \left[\frac{j}{2^n}, \frac{j+1}{2^n} \right) : j \in \mathbb{Z} \right\} \quad (1.1.20)$$

is a partition of \mathbb{R} into intervals of equal length $1/2^n$. Note that the endpoints are the [dyadic rationals](#) of rank n or less, and that \mathcal{D}_{n+1} can be obtained from \mathcal{D}_n by dividing each interval into two equal parts. This sequence of partitions is one of the reasons that the dyadic rationals are important.

Basic Rules

In the following problems, $\mathcal{A} = \{A_i : i \in I\}$ is a collection of subsets of a universal set S , indexed by a nonempty set I , and B is a subset of S .

The general *distributive laws*:

1. $(\bigcup_{i \in I} A_i) \cap B = \bigcup_{i \in I} (A_i \cap B)$
2. $(\bigcap_{i \in I} A_i) \cup B = \bigcap_{i \in I} (A_i \cup B)$

Restate the laws in the notation where the collection \mathcal{A} is not indexed.

Proof

1. x is an element of the set on the left or the right of the equation if and only if $x \in B$ and $x \in A_i$ for some $i \in I$.
2. x is an element of the set on the left or the right of the equation if and only if $x \in B$ or $x \in A_i$ for every $i \in I$.

$$(\bigcup \mathcal{A}) \cap B = \bigcup \{A \cap B : A \in \mathcal{A}\}, (\bigcap \mathcal{A}) \cup B = \bigcap \{A \cup B : A \in \mathcal{A}\}$$

The general *De Morgan's laws*:

1. $(\bigcup_{i \in I} A_i)^c = \bigcap_{i \in I} A_i^c$
2. $(\bigcap_{i \in I} A_i)^c = \bigcup_{i \in I} A_i^c$

Restate the laws in the notation where the collection \mathcal{A} is not indexed.

Proof

1. $x \in (\bigcup_{i \in I} A_i)^c$ if and only if $x \notin \bigcup_{i \in I} A_i$ if and only if $x \notin A_i$ for every $i \in I$ if and only if $x \in A_i^c$ for every $i \in I$ if and only if $x \in \bigcap_{i \in I} A_i^c$.
2. $x \in (\bigcap_{i \in I} A_i)^c$ if and only if $x \notin \bigcap_{i \in I} A_i$ if and only if $x \notin A_i$ for some $i \in I$ if and only if $x \in A_i^c$ for some $i \in I$ if and only if $x \in \bigcup_{i \in I} A_i^c$.

$$(\bigcup \mathcal{A})^c = \bigcap \{A^c : A \in \mathcal{A}\}, (\bigcap \mathcal{A})^c = \bigcup \{A^c : A \in \mathcal{A}\}$$

Suppose that the collection \mathcal{A} partitions S . For any subset B , the collection $\{A \cap B : A \in \mathcal{A}\}$ partitions B .

Proof

Suppose $\mathcal{A} = \{A_i : i \in I\}$ where I is an index set. If $i, j \in I$ with $i \neq j$ then $(A_i \cap B) \cap (A_j \cap B) = (A_i \cap A_j) \cap B = \emptyset \cap B = \emptyset$, so the collection $\{A_i \cap B : i \in I\}$ is disjoint. Moreover, by the distributive law,

$$\bigcup_{i \in I} (A_i \cap B) = \left(\bigcup_{i \in I} A_i \right) \cap B = S \cap B = B \quad (1.1.21)$$




Figure 1.1.2: A partition of S induces a partition of B

Suppose that $\{A_i : i \in \mathbb{N}_+\}$ is a collection of subsets of a universal set S

1. $\bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} A_k = \{x \in S : x \in A_k \text{ for infinitely many } k \in \mathbb{N}_+\}$
2. $\bigcup_{n=1}^{\infty} \bigcap_{k=n}^{\infty} A_k = \{x \in S : x \in A_k \text{ for all but finitely many } k \in \mathbb{N}_+\}$

Proof

1. Note that $x \in \bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} A_k$ if and only if for every $n \in \mathbb{N}_+$ there exists $k \geq n$ such that $x \in A_k$. In turn, this occurs if and only if $x \in A_k$ for infinitely many $k \in \mathbb{N}_+$.
2. Note that $x \in \bigcup_{n=1}^{\infty} \bigcap_{k=n}^{\infty} A_k$ if and only if there exists $n \in \mathbb{N}_+$ such that $x \in A_k$ for every $k \geq n$. In turn, this occurs if and only if $x \in A_k$ for all but finitely many $k \in \mathbb{N}_+$.

The sets in the previous result turn out to be important in the study of probability.

Product sets

Definitions

Product sets are sets of *sequences*. The defining property of a sequence, of course, is that *order* as well as *membership* is important.

Let us start with *ordered pairs*. In this case, the defining property is that $(a, b) = (c, d)$ if and only if $a = c$ and $b = d$. Interestingly, the structure of an ordered pair can be defined just using set theory. The construction in the result below is due to Kazimierz Kuratowski

Define $(a, b) = \{\{a\}, \{a, b\}\}$. This definition captures the defining property of an ordered pair.

Proof

Suppose that $(a, b) = (c, d)$ so that $\{\{a\}, \{a, b\}\} = \{\{c\}, \{c, d\}\}$. In the case that $a = b$ note that $(a, b) = \{\{a\}\}$. Thus we must have $\{c\} = \{c, d\} = \{a\}$ and hence $c = d = a$, and in particular, $a = c$ and $b = d$. In the case that $a \neq b$, we must have $\{c\} = \{a\}$ and hence $c = a$. But we cannot have $\{c, d\} = \{a\}$ because then $(c, d) = \{\{a\}\}$ and hence $\{a, b\} = \{a\}$, which would force $a = b$, a contradiction. Thus we must have $\{c, d\} = \{a, b\}$. Since $c = a$ and $a \neq b$ we must have $d = b$. The converse is trivial: if $a = c$ and $b = d$ then $\{a\} = \{c\}$ and $\{a, b\} = \{c, d\}$ so $(a, b) = (c, d)$.

Of course, it's important not to confuse the ordered pair (a, b) with the open interval (a, b) , since the same notation is used for both. Usually it's clear from context which type of object is referred to. For *ordered triples*, the defining property is $(a, b, c) = (d, e, f)$ if and only if $a = d$, $b = e$, and $c = f$. Ordered triples can be defined in terms of ordered pairs, which via the last result, uses only set theory.

Define $(a, b, c) = (a, (b, c))$. This definition captures the defining property of an ordered triple.

Proof

Suppose $(a, b, c) = (d, e, f)$. Then $(a, (b, c)) = (d, (e, f))$. Hence by the definition of an ordered pair, we must have $a = d$ and $(b, c) = (e, f)$. Using the definition again we have $b = e$ and $c = f$. Conversely, if $a = d$, $b = e$, and $c = f$, then $(b, c) = (e, f)$ and hence $(a, (b, c)) = (d, (e, f))$. Thus $(a, b, c) = (d, e, f)$.

All of this is just to show how complicated structures can be built from simpler ones, and ultimately from set theory. But enough of that! More generally, two ordered sequences of the same size (finite or infinite) are the same if and only if their corresponding coordinates agree. Thus for $n \in \mathbb{N}_+$, the definition for n -tuples is $(x_1, x_2, \dots, x_n) = (y_1, y_2, \dots, y_n)$ if and only if $x_i = y_i$ for all $i \in \{1, 2, \dots, n\}$. For infinite sequences, $(x_1, x_2, \dots) = (y_1, y_2, \dots)$ if and only if $x_i = y_i$ for all $i \in \mathbb{N}_+$.

Suppose now that we have a sequence of n sets, (S_1, S_2, \dots, S_n) , where $n \in \mathbb{N}_+$. The *Cartesian product* of the sets is defined as follows:

$$S_1 \times S_2 \times \cdots \times S_n = \{(x_1, x_2, \dots, x_n) : x_i \in S_i \text{ for } i \in \{1, 2, \dots, n\}\} \quad (1.1.22)$$

Cartesian products are named for René Descartes. If $S_i = S$ for each i , then the Cartesian product set can be written compactly as S^n , a *Cartesian power*. In particular, recall that \mathbb{R} denotes the set of real numbers so that \mathbb{R}^n is n -dimensional *Euclidean space*, named after Euclid, of course. The elements of $\{0, 1\}^n$ are called *bit strings* of length n . As the name suggests, we sometimes represent elements of this product set as strings rather than sequences (that is, we omit the parentheses and commas). Since the coordinates just take two values, there is no risk of confusion.

Suppose that we have an infinite sequence of sets (S_1, S_2, \dots) . The *Cartesian product* of the sets is defined by

$$S_1 \times S_2 \times \cdots = \{(x_1, x_2, \dots) : x_i \in S_i \text{ for each } i \in \{1, 2, \dots\}\} \quad (1.1.23)$$

When $S_i = S$ for $i \in \mathbb{N}_+$, the Cartesian product set is sometimes written as a Cartesian power as S^∞ or as $S^{\mathbb{N}_+}$. An explanation for the last notation, as well as a much more general construction for products of sets, is given in the next section on functions. Also, notation similar to that of general union and intersection is often used for Cartesian product, with \prod as the operator. So

$$\prod_{i=1}^n S_i = S_1 \times S_2 \times \cdots \times S_n, \quad \prod_{i=1}^{\infty} S_i = S_1 \times S_2 \times \cdots \quad (1.1.24)$$

Rules for Product Sets

We will now see how the set operations relate to the Cartesian product operation. Suppose that S and T are sets and that $A \subseteq S$, $B \subseteq S$ and $C \subseteq T$, $D \subseteq T$. The sets in the theorems below are subsets of $S \times T$.

The most important rules that relate Cartesian product with union, intersection, and difference are the *distributive rules*:

Distributive rules for product sets

1. $A \times (C \cup D) = (A \times C) \cup (A \times D)$
2. $(A \cup B) \times C = (A \times C) \cup (B \times C)$
3. $A \times (C \cap D) = (A \times C) \cap (A \times D)$
4. $(A \cap B) \times C = (A \times C) \cap (B \times C)$
5. $A \times (C \setminus D) = (A \times C) \setminus (A \times D)$
6. $(A \setminus B) \times C = (A \times C) \setminus (B \times C)$

Proof

1. $(x, y) \in A \times (C \cup D)$ if and only if $x \in A$ and $y \in C \cup D$ if and only if $x \in A$ and either $y \in C$ or $y \in D$ if and only if $x \in A$ and $y \in C$, or $x \in A$ and $y \in D$ if and only if $(x, y) \in A \times C$ or $(x, y) \in A \times D$ if and only if $(x, y) \in (A \times C) \cup (A \times D)$.
2. Similar to (a), but with the roles of the coordinates reversed.
3. $(x, y) \in A \times (C \cap D)$ if and only if $x \in A$ and $y \in C \cap D$ if and only if $x \in A$ and $y \in C$ and $y \in D$ if and only if $(x, y) \in A \times C$ and $(x, y) \in A \times D$ if and only if $(x, y) \in (A \times C) \cap (A \times D)$.
4. Similar to (c) but with the roles of the coordinates reversed.
5. $(x, y) \in A \times (C \setminus D)$ if and only if $x \in A$ and $y \in C \setminus D$ if and only if $x \in A$ and $y \in C$ and $y \notin D$ if and only if $(x, y) \in A \times C$ and $(x, y) \notin A \times D$ if and only if $(x, y) \in (A \times C) \setminus (A \times D)$.
6. Similar to (e) but with the roles of the coordinates reversed.

In general, the product of unions is larger than the corresponding union of products.

$$(A \cup B) \times (C \cup D) = (A \times C) \cup (A \times D) \cup (B \times C) \cup (B \times D)$$

Proof

$(x, y) \in (A \cup B) \times (C \cup D)$ if and only if $x \in A \cup B$ and $y \in C \cup D$ if and only if at least one of the following is true: $x \in A$ and $y \in C$, $x \in A$ and $y \in D$, $x \in B$ and $y \in C$, $x \in B$ and $y \in D$ if and only if $(x, y) \in (A \times C) \cup (A \times D) \cup (B \times C) \cup (B \times D)$

So in particular it follows that $(A \times C) \cup (B \times D) \subseteq (A \cup B) \times (C \cup D)$. On the other hand, the product of intersections is the same as the corresponding intersection of products.

$$(A \times C) \cap (B \times D) = (A \cap B) \times (C \cap D)$$

Proof

$(x, y) \in (A \times C) \cap (B \times D)$ if and only if $(x, y) \in A \times C$ and $(x, y) \in B \times D$ if and only if $x \in A$ and $y \in C$ and $x \in B$ and $y \in D$ if and only if $x \in A \cap B$ and $y \in C \cap D$ if and only if $(x, y) \in (A \cap B) \times (C \cap D)$.

In general, the product of differences is smaller than the corresponding difference of products.

$$(A \setminus B) \times (C \setminus D) = [(A \times C) \setminus (A \times D)] \setminus [(B \times C) \setminus (B \times D)]$$

Proof

$(x, y) \in (A \setminus B) \times (C \setminus D)$ if and only if $x \in A \setminus B$ and $y \in C \setminus D$ if and only if $x \in A$ and $x \notin B$ and $y \in C$ and $y \notin D$. On the other hand, $(x, y) \in [(A \times C) \setminus (A \times D)] \setminus [(B \times C) \setminus (B \times D)]$ if and only if $(x, y) \in (A \times C) \setminus (A \times D)$ and $(x, y) \notin (B \times C) \setminus (B \times D)$. The first statement means that $x \in A$ and $y \in C$ and $y \notin D$. The second statement is the negation of $x \in B$ and $y \in C$ and $y \notin D$. The two statements both hold if and only if $x \in A$ and $x \notin B$ and $y \in C$ and $y \notin D$.

So in particular it follows that $(A \setminus B) \times (C \setminus D) \subseteq (A \times C) \setminus (B \times D)$.

Projections and Cross Sections

In this discussion, suppose again that S and T are nonempty sets, and that $C \subseteq S \times T$.

Cross Sections

1. The cross section of C in the first coordinate at $x \in S$ is $C_x = \{y \in T : (x, y) \in C\}$
2. The cross section of C at in the second coordinate at $y \in T$ is

$$C^y = \{x \in S : (x, y) \in C\} \quad (1.1.25)$$

Note that $C_x \subseteq T$ for $x \in S$ and $C^y \subseteq S$ for $y \in T$.

Projections

1. The projection of C onto T is $C_T = \{y \in T : (x, y) \in C \text{ for some } x \in S\}$.
2. The projection of C onto S is $C^S = \{x \in S : (x, y) \in C \text{ for some } y \in T\}$.

The projections are the unions of the appropriate cross sections.

Unions

1. $C_T = \bigcup_{x \in S} C_x$
2. $C^S = \bigcup_{y \in T} C^y$

Cross sections are preserved under the set operations. We state the result for cross sections at $x \in S$. By symmetry, of course, analogous results hold for cross sections at $y \in T$.

Suppose that $C, D \subseteq S \times T$. Then for $x \in S$,

1. $(C \cup D)_x = C_x \cup D_x$
2. $(C \cap D)_x = C_x \cap D_x$
3. $(C \setminus D)_x = C_x \setminus D_x$

Proof

1. $y \in (C \cup D)_x$ if and only if $(x, y) \in C \cup D$ if and only if $(x, y) \in C$ or $(x, y) \in D$ if and only if $y \in C_x$ or $y \in D_x$.
2. The proof is just like (a), with *and* replacing *or*.
3. The proof is just like (a), with *and not* replacing *or*.

For projections, the results are a bit more complicated. We give the results for projections onto T ; naturally the results for projections onto S are analogous.

Suppose again that $C, D \subseteq S \times T$. Then

1. $(C \cup D)_T = C_T \cup D_T$
2. $(C \cap D)_T \subseteq C_T \cap D_T$
3. $(C_T)^c \subseteq (C^c)_T$

Proof

1. Suppose that $y \in (C \cup D)_T$. Then there exists $x \in S$ such that $(x, y) \in C \cup D$. Hence $(x, y) \in C$ so $y \in C_T$, or $(x, y) \in D$ so $y \in D_T$. In either case, $y \in C_T \cup D_T$. Conversely, suppose that $y \in C_T \cup D_T$. Then $y \in C_T$ or $y \in D_T$. If $y \in C_T$ then there exists $x \in S$ such that $(x, y) \in C$. But then $(x, y) \in C \cup D$ so $y \in (C \cup D)_T$. Similarly if $y \in D_T$ then $y \in (C \cup D)_T$.
2. Suppose that $y \in (C \cap D)_T$. Then there exists $x \in S$ such that $(x, y) \in C \cap D$. Hence $(x, y) \in C$ so $y \in C_T$ and $(x, y) \in D$ so $y \in D_T$. Therefore $y \in C_T \cap D_T$.
3. Suppose that $y \in (C_T)^c$. Then $y \notin C_T$, so for every $x \in S$, $(x, y) \notin C$. Fix $x_0 \in S$. Then $(x_0, y) \notin C$ so $(x_0, y) \in C^c$ and therefore $y \in (C^c)_T$.

It's easy to see that equality does not hold in general in parts (b) and (c). In part (b) for example, suppose that $A_1, A_2 \subseteq S$ are nonempty and disjoint and $B \subseteq T$ is nonempty. Let $C = A_1 \times B$ and $D = A_2 \times B$. Then $C \cap D = \emptyset$ so $(C \cap D)_T = \emptyset$. But $C_T = D_T = B$. In part (c) for example, suppose that A is a nonempty proper subset of S and B is a nonempty proper subset of T . Let $C = A \times B$. Then $C_T = B$ so $(C_T)^c = B^c$. On the other hand, $C^c = (A^c \times B) \cup (A \times B^c) \cup (A^c \times B^c)$, so $(C^c)_T = T$.

Cross sections and projections will be extended to very general product sets in the next section on Functions.

Computational Exercises

Subsets of \mathbb{R}

The universal set is $[0, \infty)$. Let $A = [0, 5]$ and $B = (3, 7)$. Express each of the following in terms of intervals:

1. $A \cap B$
2. $A \cup B$
3. $A \setminus B$
4. $B \setminus A$
5. A^c

Answer

1. $(3, 5]$
2. $[0, 7)$
3. $[0, 3]$
4. $(5, 7)$
5. $(5, \infty)$

The universal set is \mathbb{N} . Let $A = \{n \in \mathbb{N} : n \text{ is even}\}$ and let $B = \{n \in \mathbb{N} : n \leq 9\}$. Give each of the following:

1. $A \cap B$ in list form

2. $A \cup B$ in predicate form
3. $A \setminus B$ in list form
4. $B \setminus A$ in list form
5. A^c in predicate form
6. B^c in list form

Answer

1. $\{0, 2, 4, 6, 8\}$
2. $\{n \in \mathbb{N} : n \text{ is even or } n \leq 9\}$
3. $\{10, 12, 14, \dots\}$
4. $\{1, 3, 5, 7, 9\}$
5. $\{n \in \mathbb{N} : n \text{ is odd}\}$
6. $\{10, 11, 12, \dots\}$

Coins and Dice

Let $S = \{1, 2, 3, 4\} \times \{1, 2, 3, 4, 5, 6\}$. This is the set of outcomes when a 4-sided die and a 6-sided die are tossed. Further let $A = \{(x, y) \in S : x = 2\}$ and $B = \{(x, y) \in S : x + y = 7\}$. Give each of the following sets in list form:

1. A
2. B
3. $A \cap B$
4. $A \cup B$
5. $A \setminus B$
6. $B \setminus A$

Answer

1. $\{(2, 1), (2, 2), (2, 3), (2, 4), (2, 5), (2, 6)\}$
2. $\{(1, 6), (2, 5), (3, 4), (4, 3)\}$
3. $\{(2, 5)\}$
4. $\{(2, 1), (2, 2), (2, 3), (2, 4), (2, 5), (2, 6), (1, 6), (3, 4), (4, 3)\}$
5. $\{(2, 1), (2, 2), (2, 3), (2, 4), (2, 6)\}$
6. $\{(1, 6), (3, 4), (4, 3)\}$

Let $S = \{0, 1\}^3$. This is the set of outcomes when a coin is tossed 3 times (0 denotes tails and 1 denotes heads). Further let $A = \{(x_1, x_2, x_3) \in S : x_2 = 1\}$ and $B = \{(x_1, x_2, x_3) \in S : x_1 + x_2 + x_3 = 2\}$. Give each of the following sets in list form, using bit-string notation:

1. S
2. A
3. B
4. A^c
5. B^c
6. $A \cap B$
7. $A \cup B$
8. $A \setminus B$
9. $B \setminus A$

Answer

1. $\{000, 100, 010, 001, 110, 101, 011, 111\}$
2. $\{010, 110, 011, 111\}$
3. $\{110, 011, 101\}$
4. $\{000, 100, 001, 101\}$
5. $\{000, 100, 010, 001, 111\}$
6. $\{110, 011\}$
7. $\{010, 110, 011, 111, 101\}$
8. $\{010, 111\}$
9. $\{101\}$

Let $S = \{0, 1\}^2$. This is the set of outcomes when a coin is tossed twice (0 denotes tails and 1 denotes heads). Give $\mathcal{P}(S)$ in list form.

Answer

$\{\emptyset, \{00\}, \{01\}, \{10\}, \{11\}, \{00, 01\}, \{00, 10\}, \{00, 11\}, \{01, 10\}, \{01, 11\}, \{10, 11\}, \{00, 01, 10\}, \{00, 01, 11\}, \{00, 10, 11\}, \{01, 10, 11\}, \{00, 01, 10, 11\}\}$

Cards

A standard *card deck* can be modeled by the Cartesian product set

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (1.1.26)$$

where the first coordinate encodes the *denomination* or *kind* (ace, 2–10, jack, queen, king) and where the second coordinate encodes the *suit* (clubs, diamonds, hearts, spades). Sometimes we represent a card as a *string* rather than an ordered pair (for example $q\heartsuit$ for the queen of hearts). For the problems in this subsection, the card deck D is the universal set.

Let H denote the set of hearts and F the set of face cards. Find each of the following:

1. $H \cap F$
2. $H \setminus F$
3. $F \setminus H$
4. $H \triangle F$

Answer

1. $\{j\heartsuit, q\heartsuit, k\heartsuit\}$
2. $\{1\heartsuit, 2\heartsuit, 3\heartsuit, 4\heartsuit, 5\heartsuit, 6\heartsuit, 7\heartsuit, 8\heartsuit, 9\heartsuit, 10\heartsuit\}$
3. $\{j\clubsuit, q\clubsuit, k\clubsuit, j\diamondsuit, q\diamondsuit, k\diamondsuit, j\spadesuit, q\spadesuit, k\spadesuit\}$
4. $\{1\heartsuit, 2\heartsuit, 3\heartsuit, 4\heartsuit, 5\heartsuit, 6\heartsuit, 7\heartsuit, 8\heartsuit, 9\heartsuit, 10\heartsuit, j\clubsuit, q\clubsuit, k\clubsuit, j\diamondsuit, q\diamondsuit, k\diamondsuit, j\spadesuit, q\spadesuit, k\spadesuit\}$

A *bridge hand* is a subset of D with 13 cards. Often bridge hands are described by giving the cross sections by suit.

Suppose that N is a bridge hand, held by a player named North, defined by

$$N^\clubsuit = \{2, 5, q\}, N^\diamondsuit = \{1, 5, 8, q, k\}, N^\heartsuit = \{8, 10, j, q\}, N^\spadesuit = \{1\} \quad (1.1.27)$$

Find each of the following:

1. The nonempty cross sections of N by denomination.
2. The projection of N onto the set of suits.
3. The projection of N onto the set of denominations

Answer

1. $N_1 = \{\diamondsuit, \spadesuit\}, N_2 = \{\clubsuit\}, N_5 = \{\clubsuit, \diamondsuit\}, N_8 = \{\diamondsuit, \heartsuit\}, N_{10} = \{\heartsuit\}, N_j = \{\heartsuit\}, N_q = \{\clubsuit, \diamondsuit, \heartsuit\}, N_k = \{\diamondsuit\}$
2. $\{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\}$
3. $\{1, 2, 5, 8, 10, j, q, k\}$

By contrast, it is usually more useful to describe a *poker hand* by giving the cross sections by denomination. In the usual version of *draw poker*, a hand is a subset of D with 5 cards.

Suppose that B is a poker hand, held by a player named Bill, with

$$B_1 = \{\clubsuit, \spadesuit\}, B_8 = \{\clubsuit, \spadesuit\}, B_q = \{\heartsuit\} \quad (1.1.28)$$

Find each of the following:

1. The nonempty cross sections of B by suit.
2. The projection of B onto the set of suits.
3. The projection of B onto the set of denominations

Answer

1. $B^\clubsuit = \{1, 8\}, B^\heartsuit = \{q\}, B^\spadesuit = \{1, 8\}$
2. $\{\clubsuit, \heartsuit, \spadesuit\}$
3. $\{1, 8, q\}$

The poker hand in the last exercise is known as a *dead man's hand*. Legend has it that Wild Bill Hickock held this hand at the time of his murder in 1876.

General unions and intersections

For the problems in this subsection, the universal set is \mathbb{R} .

Let $A_n = [0, 1 - \frac{1}{n}]$ for $n \in \mathbb{N}_+$. Find

1. $\bigcap_{n=1}^{\infty} A_n$
2. $\bigcup_{n=1}^{\infty} A_n$
3. $\bigcap_{n=1}^{\infty} A_n^c$

4. $\bigcup_{n=1}^{\infty} A_n^c$

Answer

1. $\{0\}$
2. $[0, 1)$
3. $(-\infty, 0) \cup [1, \infty)$
4. $\mathbb{R} - \{0\}$

Let $A_n = (2 - \frac{1}{n}, 5 + \frac{1}{n})$ for $n \in \mathbb{N}_+$. Find

1. $\bigcap_{n=1}^{\infty} A_n$
2. $\bigcup_{n=1}^{\infty} A_n$
3. $\bigcap_{n=1}^{\infty} A_n^c$
4. $\bigcup_{n=1}^{\infty} A_n^c$

Answer

1. $[2, 5]$
2. $(1, 6)$
3. $(-\infty, 1] \cup [6, \infty)$
4. $(-\infty, 2) \cup (5, \infty)$

Subsets of \mathbb{R}^2

Let T be the closed triangular region in \mathbb{R}^2 with vertices $(0, 0)$, $(1, 0)$, and $(1, 1)$. Find each of the following:

1. The cross section T_x for $x \in \mathbb{R}$
2. The cross section T^y for $y \in \mathbb{R}$
3. The projection of T onto the horizontal axis
4. The projection of T onto the vertical axis

Answer

1. $T_x = [0, x]$ for $x \in [0, 1]$, $T_x = \emptyset$ otherwise
2. $T^y = [y, 1]$ for $y \in [0, 1]$, $T^y = \emptyset$ otherwise
3. $[0, 1]$
4. $[0, 1]$

This page titled [1.1: Sets](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.2: Functions

Functions play a central role in probability and statistics, as they do in every other branch of mathematics. For the most part, the proofs in this section are straightforward, so be sure to try them yourself before reading the ones in the text.

Definitions and Properties

Basic Definitions

We start with the formal, technical definition of a function. It's not very intuitive, but has the advantage that it only requires set theory.

A function f from a set S into a set T is a subset of the product set $S \times T$ with the property that for each element $x \in S$, there exists a unique element $y \in T$ such that $(x, y) \in f$. If f is a function from S to T we write $f : S \rightarrow T$. If $(x, y) \in f$ we write $y = f(x)$.

Less formally, a function f from S into T is a “rule” (or “procedure” or “algorithm”) that assigns to each $x \in S$ a unique element $f(x) \in T$. The definition of a function as a set of ordered pairs, is due to Kazimierz Kuratowski. The term *map* or *mapping* is also used in place of function, so we could say that f maps S into T .




Figure 1.2.1: A function f from S into T

The sets S and T in the definition are clearly important.

Suppose that $f : S \rightarrow T$.

1. The set S is the *domain* of f .
2. The set T is the *range space* or *co-domain* of f .
3. The *range* of f is the set of function values. That is, $\text{range}(f) = \{y \in T : y = f(x) \text{ for some } x \in S\}$.

The domain and range are completely specified by a function. That's not true of the co-domain: if f is a function from S into T , and U is another set with $T \subseteq U$, then we can also think of f as a function from S into U . The following definitions are natural and important.

Suppose again that $f : S \rightarrow T$.

1. f maps S *onto* T if $\text{range}(f) = T$. That is, for each $y \in T$ there exists $x \in S$ such that $f(x) = y$.
2. f is *one-to-one* if distinct elements in the domain are mapped to distinct elements in the range. That is, if $u, v \in S$ and $u \neq v$ then $f(u) \neq f(v)$.

Clearly a function always maps its domain onto its range. Note also that f is one-to-one if $f(u) = f(v)$ implies $u = v$ for $u, v \in S$.

Inverse functions

A function that is one-to-one and onto can be “reversed” in a sense.

If f maps S one-to-one onto T , the *inverse* of f is the function f^{-1} from T onto S given by

$$f^{-1}(y) = x \iff f(x) = y; \quad x \in S, y \in T \quad (1.2.1)$$

If you like to think of a function as a set of ordered pairs, then $f^{-1} = \{(y, x) \in T \times S : (x, y) \in f\}$. The fact that f is one-to-one and onto ensures that f^{-1} is a valid function from T onto S . Sets S and T are in *one-to-one correspondence* if there exists a one-to-one function from S onto T . One-to-one correspondence plays an essential role in the study of cardinality.

Restrictions

The domain of a function can be *restricted* to create a new function.

Suppose that $f : S \rightarrow T$ and that $A \subseteq S$. The function $f_A : A \rightarrow T$ defined by $f_A(x) = f(x)$ for $x \in A$ is the *restriction of f to A* .

As a set of ordered pairs, note that $f_A = \{(x, y) \in f : x \in A\}$.

Composition

Composition is perhaps the most important way to combine two functions to create another function.

Suppose that $g : R \rightarrow S$ and $f : S \rightarrow T$. The *composition* of f with g is the function $f \circ g : R \rightarrow T$ defined by

$$(f \circ g)(x) = f(g(x)), \quad x \in R \quad (1.2.2)$$

Composition is *associative*:

Suppose that $h : R \rightarrow S$, $g : S \rightarrow T$, and $f : T \rightarrow U$. Then

$$f \circ (g \circ h) = (f \circ g) \circ h \quad (1.2.3)$$

Proof

Note that both functions map R into U . Using the definition of composition, the value of both functions at $x \in R$ is $f(g(h(x)))$.

Thus we can write $f \circ g \circ h$ without ambiguity. On the other hand, composition is not *commutative*. Indeed depending on the domains and co-domains, $f \circ g$ might be defined when $g \circ f$ is not. Even when both are defined, they may have different domains and co-domains, and so of course cannot be the same function. Even when both are defined and have the same domains and co-domains, the two compositions will not be the same in general. Examples of all of these cases are given in the computational exercises below.

Suppose that $g : R \rightarrow S$ and $f : S \rightarrow T$.

1. If f and g are one-to-one then $f \circ g$ is one-to-one.
2. If f and g are onto then $f \circ g$ is onto.

Proof

1. Suppose that $u, v \in R$ and $(f \circ g)(u) = (f \circ g)(v)$. Then $f(g(u)) = f(g(v))$. Since f is one-to-one, $g(u) = g(v)$. Since g is one-to-one, $u = v$.
2. Suppose that $z \in T$. Since f is onto, there exist $y \in S$ with $f(y) = z$. Since g is onto, there exists $x \in R$ with $g(x) = y$. Then $(f \circ g)(x) = f(g(x)) = f(y) = z$.

The *identity function* on a set S is the function I_S from S onto S defined by $I_S(x) = x$ for $x \in S$

The identity function acts like an identity with respect to the operation of composition.

If $f : S \rightarrow T$ then

1. $f \circ I_S = f$
2. $I_T \circ f = f$

Proof

1. Note that $f \circ I_S : S \rightarrow T$. For $x \in S$, $(f \circ I_S)(x) = f(I_S(x)) = f(x)$.
2. Note that $I_T \circ f : S \rightarrow T$. For $x \in S$, $(I_T \circ f)(x) = I_T(f(x)) = f(x)$.

The inverse of a function is really the inverse with respect to composition.

Suppose that f is a one-to-one function from S onto T . Then

1. $f^{-1} \circ f = I_S$
2. $f \circ f^{-1} = I_T$

Proof

1. Note that $f^{-1} \circ f : S \rightarrow S$. For $x \in S$, $(f^{-1} \circ f)(x) = f^{-1}(f(x)) = x$.
2. Note that $f \circ f^{-1} : T \rightarrow T$. For $y \in T$, $(f \circ f^{-1})(y) = f(f^{-1}(y)) = y$.

An element $x \in S^n$ can be thought of as a function from $\{1, 2, \dots, n\}$ into S . Similarly, an element $x \in S^\infty$ can be thought of as a function from \mathbb{N}_+ into S . For such a sequence x , of course, we usually write x_i instead of $x(i)$. More generally, if S and T are sets, then the set of all functions from S into T is denoted by T^S . In particular, as we noted in the last section, S^∞ is also (and more accurately) written as $S^{\mathbb{N}_+}$.

Suppose that g is a one-to-one function from R onto S and that f is a one-to-one function from S onto T . Then $(f \circ g)^{-1} = g^{-1} \circ f^{-1}$.

Proof

Note that $(f \circ g)^{-1} : T \rightarrow R$ and $g^{-1} \circ f^{-1} : T \rightarrow R$. For $y \in T$, let $x = (f \circ g)^{-1}(y)$. Then $(f \circ g)(x) = y$ so that $f(g(x)) = y$ and hence $g(x) = f^{-1}(y)$ and finally $x = g^{-1}(f^{-1}(y))$.

Inverse Images

Inverse images of a function play a fundamental role in probability, particularly in the context of random variables.

Suppose that $f : S \rightarrow T$. If $A \subseteq T$, the *inverse image* of A under f is the subset of S given by

$$f^{-1}(A) = \{x \in S : f(x) \in A\} \quad (1.2.4)$$

So $f^{-1}(A)$ is the subset of S consisting of those elements that map into A .




Figure 1.2.2: The inverse image of A under f

Technically, the inverse images define a new function from $\mathcal{P}(T)$ into $\mathcal{P}(S)$. We use the same notation as for the *inverse function*, which is defined when f is one-to-one and onto. These are very different functions, but usually no confusion results. The following important theorem shows that inverse images preserve all set operations.

Suppose that $f : S \rightarrow T$, and that $A, B \subseteq T$. Then

1. $f^{-1}(A \cup B) = f^{-1}(A) \cup f^{-1}(B)$
2. $f^{-1}(A \cap B) = f^{-1}(A) \cap f^{-1}(B)$
3. $f^{-1}(A \setminus B) = f^{-1}(A) \setminus f^{-1}(B)$
4. If $A \subseteq B$ then $f^{-1}(A) \subseteq f^{-1}(B)$
5. If A and B are disjoint, so are $f^{-1}(A)$ and $f^{-1}(B)$

Proof

1. $x \in f^{-1}(A \cup B)$ if and only if $f(x) \in A \cup B$ if and only if $f(x) \in A$ or $f(x) \in B$ if and only if $x \in f^{-1}(A)$ or $x \in f^{-1}(B)$ if and only if $x \in f^{-1}(A) \cup f^{-1}(B)$
2. The proof is the same as (a), with intersection replacing union and with *and* replacing *or* throughout.
3. The proof is the same as (a), with set difference replacing union and with *and not* replacing *or* throughout.

4. Suppose $A \subseteq B$. If $x \in f^{-1}(A)$ then $f(x) \in A$ and hence $f(x) \in B$, so $x \in f^{-1}(B)$.
5. If A and B are disjoint, then from (b), $f^{-1}(A) \cap f^{-1}(B) = f^{-1}(A \cap B) = f^{-1}(\emptyset) = \emptyset$.

The result in part (a) holds for arbitrary unions, and the result in part (b) holds for arbitrary intersections. No new ideas are involved; only the notation is more complicated.

Suppose that $\{A_i : i \in I\}$ is a collection of subsets of T , where I is a nonempty index set. Then

1. $f^{-1}(\bigcup_{i \in I} A_i) = \bigcup_{i \in I} f^{-1}(A_i)$
2. $f^{-1}(\bigcap_{i \in I} A_i) = \bigcap_{i \in I} f^{-1}(A_i)$

Proof

1. $x \in f^{-1}(\bigcup_{i \in I} A_i)$ if and only if $f(x) \in \bigcup_{i \in I} A_i$ if and only if $f(x) \in A_i$ for some $i \in I$ if and only if $x \in f^{-1}(A_i)$ for some $i \in I$ if and only if $x \in \bigcup_{i \in I} f^{-1}(A_i)$
2. The proof is the same as (a), with intersection replacing union and with *for every* replacing *for some*.

Forward Images

Forward images of a function are a naturally complement to inverse images.

Suppose again that $f : S \rightarrow T$. If $A \subseteq S$, the *forward image* of A under f is the subset of T given by

$$f(A) = \{f(x) : x \in A\} \quad (1.2.5)$$

So $f(A)$ is the range of f restricted to A .




Figure 1.2.3: The forward image of A under f

Technically, the forward images define a new function from $\mathcal{P}(S)$ into $\mathcal{P}(T)$, but we use the same symbol f for this new function as for the underlying function from S into T that we started with. Again, the two functions are very different, but usually no confusion results.

It might seem that forward images are more natural than inverse images, but in fact, the inverse images are much more important than the forward ones (at least in probability and measure theory). Fortunately, the inverse images are nicer as well—unlike the inverse images, the forward images do not preserve all of the set operations.

Suppose that $f : S \rightarrow T$, and that $A, B \subseteq S$. Then

1. $f(A \cup B) = f(A) \cup f(B)$.
2. $f(A \cap B) \subseteq f(A) \cap f(B)$. Equality holds if f is one-to-one.
3. $f(A) \setminus f(B) \subseteq f(A \setminus B)$. Equality holds if f is one-to-one.
4. If $A \subseteq B$ then $f(A) \subseteq f(B)$.

Proof

1. Suppose $y \in f(A \cup B)$. Then $y = f(x)$ for some $x \in A \cup B$. If $x \in A$ then $y \in f(A)$ and if $x \in B$ then $y \in f(B)$. In both cases $y \in f(A) \cup f(B)$. Conversely suppose $y \in f(A) \cup f(B)$. If $y \in f(A)$ then $y = f(x)$ for some $x \in A$. But then $x \in A \cup B$ so $y \in f(A \cup B)$. Similarly, if $y \in f(B)$ then $y = f(x)$ for some $x \in B$. But then $x \in A \cup B$ so $y \in f(A \cup B)$.
2. If $y \in f(A \cap B)$ then $y = f(x)$ for some $x \in A \cap B$. But then $x \in A$ so $y \in f(A)$ and $x \in B$ so $y \in f(B)$ and hence $y \in f(A) \cap f(B)$. Conversely, suppose that $y \in f(A) \cap f(B)$. Then $y \in f(A)$ and $y \in f(B)$, so there exists $x \in A$ with $f(x) = y$ and there exists $u \in B$ with $f(u) = y$. At this point, we can go no further. But if f is one-to-one, then $u = x$ and hence $x \in A$ and $x \in B$. Thus $x \in A \cap B$ so $y \in f(A \cap B)$.

3. Suppose $y \in f(A) \setminus f(B)$. Then $y \in f(A)$ and $y \notin f(B)$. Hence $y = f(x)$ for some $x \in A$ and $y \neq f(u)$ for every $u \in B$. Thus, $x \notin B$ so $x \in A \setminus B$ and hence $y \in f(A \setminus B)$. Conversely, suppose $y \in f(A \setminus B)$. Then $y = f(x)$ for some $x \in A \setminus B$. Hence $x \in A$ so $y \in f(A)$. Again, the proof breaks down at this point. However, if f is one-to-one and $f(u) = y$ for some $u \in B$, then $u = x$ so $x \in B$, a contradiction. Hence $f(u) \neq y$ for every $u \in B$ so $y \notin f(B)$. Thus $y \in f(A \setminus B)$.
4. Suppose $A \subseteq B$. If $y \in f(A)$ then $y = f(x)$ for some $x \in A$. But then $x \in B$ so $y \in f(B)$.

The result in part (a) hold for arbitrary unions, and the results in part (b) hold for arbitrary intersections. No new ideas are involved; only the notation is more complicated.

Suppose that $\{A_i : i \in I\}$ is a collection of subsets of S , where I is a nonempty index set. Then

1. $f\left(\bigcup_{i \in I} A_i\right) = \bigcup_{i \in I} f(A_i)$.
2. $f\left(\bigcap_{i \in I} A_i\right) \subseteq \bigcap_{i \in I} f(A_i)$. Equality holds if f is one-to-one.

Proof

1. $y \in f\left(\bigcup_{i \in I} A_i\right)$ if and only if $y = f(x)$ for some $x \in \bigcup_{i \in I} A_i$ if and only if $y = f(x)$ for some $x \in A_i$ and some $i \in I$ if and only if $y \in f(A_i)$ for some $i \in I$ if and only if $y \in \bigcup_{i \in I} f(A_i)$.
2. If $y \in f\left(\bigcap_{i \in I} A_i\right)$ then $y = f(x)$ for some $x \in \bigcap_{i \in I} A_i$. Hence $x \in A_i$ for every $i \in I$ so $y \in f(A_i)$ for every $i \in I$ and thus $y \in \bigcap_{i \in I} f(A_i)$. Conversely, suppose that $y \in \bigcap_{i \in I} f(A_i)$. Then $y \in f(A_i)$ for every $i \in I$. Hence for every $i \in I$ there exists $x_i \in A_i$ with $y = f(x_i)$. If f is one-to-one, $x_i = x_j$ for all $i, j \in I$. Call the common value x . Then $x \in A_i$ for every $i \in I$ so $x \in \bigcap_{i \in I} A_i$ and therefore $y \in f\left(\bigcap_{i \in I} A_i\right)$.

Suppose again that $f : S \rightarrow T$. As noted earlier, the forward images of f define a function from $\mathcal{P}(S)$ into $\mathcal{P}(T)$ and the inverse images define a function from $\mathcal{P}(T)$ into $\mathcal{P}(S)$. One might hope that these functions are inverses of one another, but alas no.

Suppose that $f : S \rightarrow T$.

1. $A \subseteq f^{-1}[f(A)]$ for $A \subseteq S$. Equality holds if f is one-to-one.
2. $f[f^{-1}(B)] \subseteq B$ for $B \subseteq T$. Equality holds if f is onto.

Proof

1. If $x \in A$ then $f(x) \in f(A)$ and hence $x \in f^{-1}[f(A)]$. Conversely suppose that $x \in f^{-1}[f(A)]$. Then $f(x) \in f(A)$ so $f(x) = f(u)$ for some $u \in A$. At this point we can go no further. But if f is one-to-one, then $u = x$ and hence $x \in A$.
2. Suppose $y \in f[f^{-1}(B)]$. Then $y = f(x)$ for some $x \in f^{-1}(B)$. But then $y = f(x) \in B$. Conversely suppose that f is onto and $y \in B$. There exist $x \in S$ with $f(x) = y$. Hence $x \in f^{-1}(B)$ and so $y \in f[f^{-1}(B)]$.

Spaces of Real Functions

Real-valued function on a given set S are of particular importance. The usual arithmetic operations on such functions are defined *pointwise*.

Suppose that $f, g : S \rightarrow \mathbb{R}$ and $c \in \mathbb{R}$, then $f + g, f - g, fg, cf, f/g : S \rightarrow \mathbb{R}$ are defined as follows for all $x \in S$.

1. $(f + g)(x) = f(x) + g(x)$
2. $(f - g)(x) = f(x) - g(x)$
3. $(fg)(x) = f(x)g(x)$
4. $(cf)(x) = cf(x)$
5. $(f/g)(x) = f(x)/g(x)$ assuming that $g(x) \neq 0$ for $x \in S$.

Now let \mathcal{V} denote the collection of all functions from the given set S into \mathbb{R} . A fact that is very important in probability as well as other branches of analysis is that \mathcal{V} , with addition and scalar multiplication as defined above, is a *vector space*. The *zero function* $\mathbf{0}$ is defined, of course, by $\mathbf{0}(x) = 0$ for all $x \in S$.

$(\mathcal{V}, +, \cdot)$ is a vector space over \mathbb{R} . That is, for all $f, g, h \in \mathcal{V}$ and $a, b \in \mathbb{R}$

1. $f + g = g + f$, the *commutative property* of vector addition.
2. $f + (g + h) = (f + g) + h$, the *associative property* of vector addition.

3. $a(f + g) = af + ag$, scalar multiplication *distributes* over vector addition.
4. $(a + b)f = af + bf$, scalar multiplication *distributive* over scalar addition.
5. $f + \mathbf{0} = f$, the existence of an *zero vector*.
6. $f + (-f) = \mathbf{0}$, the existence of *additive inverses*.
7. $1 \cdot f = f$, the *unity property*.

Proof

Each of these properties follows from the corresponding property in \mathbb{R} .

Various subspaces of \mathcal{V} are important in probability as well. We will return to the discussion of vector spaces of functions in the sections on partial orders and in the advanced sections on metric spaces and measure theory.

Indicator Functions

For our next discussion, suppose that S is the universal set, so that all other sets mentioned are subsets of S .

Suppose that $A \subseteq S$. The *indicator function* of A is the function $\mathbf{1}_A : S \rightarrow \{0, 1\}$ defined as follows:

$$\mathbf{1}_A(x) = \begin{cases} 1, & x \in A \\ 0, & x \notin A \end{cases} \quad (1.2.6)$$

Thus, the indicator function of A simply *indicates* whether or not $x \in A$ for each $x \in S$. Conversely, any function on S that just takes the values 0 and 1 is an indicator function.

If $f : S \rightarrow \{0, 1\}$ then f is the indicator function of the set $A = f^{-1}\{1\} = \{x \in S : f(x) = 1\}$.

Thus, there is a one-to-one correspondence between $\mathcal{P}(S)$, the power set of S , and the collection of indicator functions $\{0, 1\}^S$. The next result shows how the *set algebra* of subsets corresponds to the *arithmetic algebra* of the indicator functions.

Suppose that $A, B \subseteq S$. Then

1. $\mathbf{1}_{A \cap B} = \mathbf{1}_A \mathbf{1}_B = \min\{\mathbf{1}_A, \mathbf{1}_B\}$
2. $\mathbf{1}_{A \cup B} = 1 - (1 - \mathbf{1}_A)(1 - \mathbf{1}_B) = \max\{\mathbf{1}_A, \mathbf{1}_B\}$
3. $\mathbf{1}_{A^c} = 1 - \mathbf{1}_A$
4. $\mathbf{1}_{A \setminus B} = \mathbf{1}_A(1 - \mathbf{1}_B)$
5. $A \subseteq B$ if and only if $\mathbf{1}_A \leq \mathbf{1}_B$

Proof

1. Note that both functions on the right just take the values 0 and 1. Moreover, $\mathbf{1}_A(x)\mathbf{1}_B(x) = \min\{\mathbf{1}_A(x), \mathbf{1}_B(x)\} = 1$ if and only if $x \in A$ and $x \in B$.
2. Note that both function on the right just take the values 0 and 1. Moreover, $1 - (1 - \mathbf{1}_A(x))(1 - \mathbf{1}_B(x)) = \max\{\mathbf{1}_A(x), \mathbf{1}_B(x)\} = 1$ if and only if $x \in A$ or $x \in B$.
3. Note that $1 - \mathbf{1}_A$ just takes the values 0 and 1. Moreover, $1 - \mathbf{1}_A(x) = 1$ if and only if $x \notin A$.
4. Note that $\mathbf{1}_{A \setminus B} = \mathbf{1}_{A \cap B^c} = \mathbf{1}_A \mathbf{1}_{B^c} = \mathbf{1}_A(1 - \mathbf{1}_B)$ by parts (a) and (c).
5. Since both functions just take the values 0 and 1, note that $\mathbf{1}_A \leq \mathbf{1}_B$ if and only if $\mathbf{1}_A(x) = 1$ implies $\mathbf{1}_B(x) = 1$. But in turn, this is equivalent to $A \subseteq B$.

The results in part (a) extends to arbitrary intersections and the results in part (b) extends to arbitrary unions.

Suppose that $\{A_i : i \in I\}$ is a collection of subsets of S , where I is a nonempty index set. Then

1. $\mathbf{1}_{\bigcap_{i \in I} A_i} = \prod_{i \in I} \mathbf{1}_{A_i} = \min\{\mathbf{1}_{A_i} : i \in I\}$
2. $\mathbf{1}_{\bigcup_{i \in I} A_i} = 1 - \prod_{i \in I} (1 - \mathbf{1}_{A_i}) = \max\{\mathbf{1}_{A_i} : i \in I\}$

Proof

In general, a product over an infinite index set may not make sense. But if all of the factors are either 0 or 1, as they are here, then we can simply define the product to be 1 if all of the factors are 1, and 0 otherwise.

1. The functions in the middle and on the right just take the values 0 and 1. Moreover, both take the value 1 at $x \in S$ if and only if $x \in A_i$ for every $i \in I$.
2. The functions in the middle and on the right just take the values 0 and 1. Moreover, both take the value 1 at $x \in S$ if and only if $x \in A_i$ for some $i \in I$.

Multisets

A *multiset* is like an ordinary set, except that elements may be repeated. A multiset A (with elements from a universal set S) can be uniquely associated with its *multiplicity function* $m_A : S \rightarrow \mathbb{N}$, where $m_A(x)$ is the number of times that element x is in A for each $x \in S$. So the multiplicity function of a multiset plays the same role that an indicator function does for an ordinary set. Multisets arise naturally when objects are sampled with replacement (but without regard to order) from a population. Various sampling models are explored in the section on Combinatorial Structures. We will not go into detail about the operations on multisets, but the definitions are straightforward generalizations of those for ordinary sets.

Suppose that A and B are multisets with elements from the universal set S . Then

1. $A \subseteq B$ if and only if $m_A \leq m_B$
2. $m_{A \cup B} = \max\{m_A, m_B\}$
3. $m_{A \cap B} = \min\{m_A, m_B\}$
4. $m_{A+B} = m_A + m_B$

Product Spaces

Using functions, we can generalize the Cartesian products studied earlier. In this discussion, we suppose that S_i is a set for each i in a nonempty index set I .

Define the *product set*

$$\prod_{i \in I} S_i = \left\{ x : x \text{ is a function from } I \text{ into } \bigcup_{i \in I} S_i \text{ such that } x(i) \in S_i \text{ for each } i \in I \right\} \quad (1.2.7)$$

Note that except for being nonempty, there are no assumptions on the cardinality of the index set I . Of course, if $I = \{1, 2, \dots, n\}$ for some $n \in \mathbb{N}_+$, or if $I = \mathbb{N}_+$ then this construction reduces to $S_1 \times S_2 \times \dots \times S_n$ and to $S_1 \times S_2 \times \dots$, respectively. Since we want to make the notation more closely resemble that of simple Cartesian products, we will write x_i instead of $x(i)$ for the value of the function x at $i \in I$, and we sometimes refer to this value as the *i th coordinate* of x . Finally, note that if $S_i = S$ for each $i \in I$, then $\prod_{i \in I} S_i$ is simply the set of all functions from I into S , which we denoted by S^I above.

For $j \in I$ define the *projection* $p_j : \prod_{i \in I} S_i \rightarrow S_j$ by $p_j(x) = x_j$ for $x \in \prod_{i \in I} S_i$.

So $p_j(x)$ is just the j th coordinate of x . The projections are of basic importance for product spaces. In particular, we have a better way of looking at projections of a subset of a product set.

For $A \subseteq \prod_{i \in I} S_i$ and $j \in I$, the *forward image* $p_j(A)$ is the projection of A onto S_j .

Proof

Note that $p_j(A) = \{p_j(x) : x \in A\} = \{x_j : x \in A\}$, the set of all j th coordinates of the points in A .

So the properties of projection that we studied in the last section are just special cases of the properties of forward images. Projections also allow us to get *coordinate functions* in a simple way.

Suppose that R is a set, and that $f : R \rightarrow \prod_{i \in I} S_i$. If $j \in I$ then $p_j \circ f : R \rightarrow S_j$ is the *j th coordinate function* of f .

Proof

Note that for $x \in R$, $(p_j \circ f)(x) = p_j[f(x)] = f_j(x)$, the j th coordinate of $f(x) \in \prod_{i \in I} S_i$.

This will look more familiar for a simple cartesian product. If $f : R \rightarrow S_1 \times S_2 \times \dots \times S_n$, then $f = (f_1, f_2, \dots, f_n)$ where $f_j : R \rightarrow S_j$ is the j th coordinate function for $j \in \{1, 2, \dots, n\}$.

Cross sections of a subset of a product set can be expressed in terms of **inverse images** of a function. First we need some additional notation. Suppose that our index set I has at least two elements. For $j \in I$ and $u \in S_j$, define $j_u : \prod_{i \in I - \{j\}} S_i \rightarrow \prod_{i \in I} S_i$ by $j_u(x) = y$ where $y_i = x_i$ for $i \in I - \{j\}$ and $y_j = u$. In words, j_u takes a point $x \in \prod_{i \in I - \{j\}} S_i$ and assigns u to coordinate j to produce the point $y \in \prod_{i \in I} S_i$.

In the setting above, if $j \in I$, $u \in S_j$ and $A \subseteq \prod_{i \in I} S_i$ then $j_u^{-1}(A)$ is the cross section of A in the j th coordinate at u .

Proof

This follows from the definition of cross section: $j_u^{-1}(A)$ is the set of all $x \in \prod_{i \in I - \{j\}} S_i$ such that y defined above is in A and has j th coordinate u .

Let's look at this for the product of two sets S and T . For $x \in S$, the function $1_x : T \rightarrow S \times T$ is given by $1_x(y) = (x, y)$. Similarly, for $y \in T$, the function $2_y : S \rightarrow S \times T$ is given by $2_y(x) = (x, y)$. Suppose now that $A \subseteq S \times T$. If $x \in S$, then $1_x^{-1}(A) = \{y \in T : (x, y) \in A\}$, the very definition of the cross section of A in the first coordinate at x . Similarly, if $y \in T$, then $2_y^{-1}(A) = \{x \in S : (x, y) \in A\}$, the very definition of the cross section of A in the second coordinate at y . This construction is not particularly important except to show that cross sections are inverse images. Thus the fact that cross sections preserve all of the set operations is a simple consequence of the fact that inverse images generally preserve set operations.

Operators

Sometimes functions have special interpretations in certain settings.

Suppose that S is a set.

1. A function $f : S \rightarrow S$ is sometimes called a *unary operator* on S .
2. A function $g : S \times S \rightarrow S$ is sometimes called a *binary operator* on S .

As the names suggests, a unary operator f *operates* on an element $x \in S$ to produce a new element $f(x) \in S$. Similarly, a binary operator g *operates* on a pair of elements $(x, y) \in S \times S$ to produce a new element $g(x, y) \in S$. The *arithmetic operators* are quintessential examples:

The following are operators on \mathbb{R} :

1. $\text{minus}(x) = -x$ is a unary operator.
2. $\text{sum}(x, y) = x + y$ is a binary operator.
3. $\text{product}(x, y) = x y$ is a binary operator.
4. $\text{difference}(x, y) = x - y$ is a binary operator.

For a fixed universal set S , the *set operators* studied in the section on Sets provide other examples.

For a given set S , the following are operators on $\mathcal{P}(S)$:

1. $\text{complement}(A) = A^c$ is a unary operator.
2. $\text{union}(A, B) = A \cup B$ is a binary operator.
3. $\text{intersect}(A, B) = A \cap B$ is a binary operator.
4. $\text{difference}(A, B) = A \setminus B$ is a binary operator.

As these examples illustrate, a binary operator is often written as $x f y$ rather than $f(x, y)$. Still, it is useful to know that operators are simply functions of a special type.

Suppose that f is a unary operator on a set S , g is a binary operator on S , and that $A \subseteq S$.

1. A is *closed* under f if $x \in A$ implies $f(x) \in A$.
2. A is *closed* under g if $(x, y) \in A \times A$ implies $g(x, y) \in A$.

Thus if A is closed under the unary operator f , then f restricted to A is unary operator on A . Similary if A is closed under the binary operator g , then g restricted to $A \times A$ is a binary operator on A . Let's return to our most basic example.

For the arithmetic operators on \mathbb{R} ,

1. \mathbb{N} is closed under plus and times, but not under minus and difference.
2. \mathbb{Z} is closed under plus, times, minus, and difference.
3. \mathbb{Q} is closed under plus, times, minus, and difference.

Many properties that you are familiar with for special operators (such as the arithmetic and set operators) can now be formulated generally.

Suppose that f and g are binary operators on a set S . In the following definitions, x , y , and z are arbitrary elements of S .

1. f is *commutative* if $f(x, y) = f(y, x)$, that is, $x f y = y f x$
2. f is *associative* if $f(x, f(y, z)) = f(f(x, y), z)$, that is, $x f (y f z) = (x f y) f z$
3. g *distributes* over f (on the left) if $g(x, f(y, z)) = f(g(x, y), g(x, z))$, that is, $x g (y f z) = (x g y) f (x g z)$

The Axiom of Choice

Suppose that \mathcal{S} is a collection of nonempty subsets of a set S . The *axiom of choice* states that there exists a function $f : \mathcal{S} \rightarrow S$ with the property that $f(A) \in A$ for each $A \in \mathcal{S}$. The function f is known as a *choice function*.

Stripped of most of the mathematical jargon, the idea is very simple. Since each set $A \in \mathcal{S}$ is nonempty, we can select an element of A ; we will call the element we select $f(A)$ and thus our selections define a function. In fact, you may wonder why we need an axiom at all. The problem is that we have not given a *rule* (or procedure or algorithm) for selecting the elements of the sets in the collection. Indeed, we may not know enough about the sets in the collection to define a specific rule, so in such a case, the axiom of choice simply guarantees the existence of a choice function. Some mathematicians, known as *constructionists* do not accept the axiom of choice, and insist on well defined rules for constructing functions.

A nice consequence of the axiom of choice is a type of duality between one-to-one functions and onto functions.

Suppose that f is a function from a set S onto a set T . There exists a one-to-one function g from T into S .

Proof.

For each $y \in T$, the set $f^{-1}\{y\}$ is non-empty, since f is onto. By the axiom of choice, we can select an element $g(y)$ from $f^{-1}\{y\}$ for each $y \in T$. The resulting function g is one-to-one.

Suppose that f is a one-to-one function from a set S into a set T . There exists a function g from T onto S .

Proof.

Fix a special element $x_0 \in S$. If $y \in \text{range}(f)$, there exists a unique $x \in S$ with $f(x) = y$. Define $g(y) = x$. If $y \notin \text{range}(f)$, define $g(y) = x_0$. The function g is onto.

Computational Exercises

Some Elementary Functions

Each of the following rules defines a function from \mathbb{R} into \mathbb{R} .

- $f(x) = x^2$
- $g(x) = \sin(x)$
- $h(x) = \lfloor x \rfloor$
- $u(x) = \frac{e^x}{1+e^x}$

Find the range of the function and determine if the function is one-to-one in each of the following cases:

1. f
2. g
3. h
4. u

Answer

1. Range $[0, \infty)$. Not one-to-one.
2. Range $[-1, 1]$. Not one-to-one.
3. Range \mathbb{Z} . Not one-to-one.
4. Range $(0, 1)$. One-to-one.

Find the following inverse images:

1. $f^{-1}[4, 9]$
2. $g^{-1}\{0\}$
3. $h^{-1}\{2, 3, 4\}$

Answer

1. $[-3, -2] \cup [2, 3]$
2. $\{n\pi : n \in \mathbb{Z}\}$
3. $[2, 5)$

The function u is one-to-one. Find (that is, give the domain and rule for) the inverse function u^{-1} .

Answer

$$u^{-1}(p) = \ln\left(\frac{p}{1-p}\right) \text{ for } p \in (0, 1)$$

Give the rule and find the range for each of the following functions:

1. $f \circ g$
2. $g \circ f$
3. $h \circ g \circ f$

Answer

1. $(f \circ g)(x) = \sin^2(x)$. Range $[0, 1]$
2. $(g \circ f)(x) = \sin(x^2)$. Range $[-1, 1]$
3. $(h \circ g \circ f)(x) = \lfloor \sin(x^2) \rfloor$. Range $\{-1, 0, 1\}$

Note that $f \circ g$ and $g \circ f$ are well-defined functions from \mathbb{R} into \mathbb{R} , but $f \circ g \neq g \circ f$.

Dice

Let $S = \{1, 2, 3, 4, 5, 6\}^2$. This is the set of possible outcomes when a pair of standard dice are thrown. Let f, g, u , and v be the functions from S into \mathbb{Z} defined by the following rules:

- $f(x, y) = x + y$
- $g(x, y) = y - x$
- $u(x, y) = \min\{x, y\}$
- $v(x, y) = \max\{x, y\}$

In addition, let F and U be the functions defined by $F = (f, g)$ and $U = (u, v)$.

Find the range of each of the following functions:

1. f
2. g
3. u
4. v
5. U

Answer

1. $\{2, 3, 4, \dots, 12\}$
2. $\{-5, -4, \dots, 4, 5\}$

3. $\{1, 2, 3, 4, 5, 6\}$
4. $\{1, 2, 3, 4, 5, 6\}$
5. $\{(i, j) \in \{1, 2, 3, 4, 5, 6\}^2 : i \leq j\}$

Give each of the following inverse images in list form:

1. $f^{-1}\{6\}$
2. $u^{-1}\{3\}$
3. $v^{-1}\{4\}$
4. $U^{-1}\{(3, 4)\}$

Answer

1. $\{(1, 5), (2, 4), (3, 3), (4, 2), (5, 1)\}$
2. $\{(3, 3), (3, 4), (4, 3), (3, 5), (5, 3), (3, 6), (6, 3)\}$
3. $\{(1, 4), (4, 1), (2, 4), (4, 2), (3, 4), (4, 3), (4, 4)\}$
4. $\{(3, 4), (4, 3)\}$

Find each of the following compositions:

1. $f \circ U$
2. $g \circ U$
3. $u \circ F$
4. $v \circ F$
5. $F \circ U$
6. $U \circ F$

Answer

1. $f \circ U = f$
2. $g \circ U = |g|$
3. $u \circ F = g$
4. $v \circ F = f$
5. $F \circ U = (f, |g|)$
6. $U \circ F = (g, f)$

Note that while $f \circ U$ is well-defined, $U \circ f$ is not. Note also that $f \circ U = f$ even though U is not the identity function on S .

Bit Strings

Let $n \in \mathbb{N}_+$ and let $S = \{0, 1\}^n$ and $T = \{0, 1, \dots, n\}$. Recall that the elements of S are bit strings of length n , and could represent the possible outcomes of n tosses of a coin (where 1 means heads and 0 means tails). Let $f : S \rightarrow T$ be the function defined by $f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n x_i$. Note that $f(\mathbf{x})$ is just the number of 1s in the bit string \mathbf{x} . Let $g : T \rightarrow S$ be the function defined by $g(k) = \mathbf{x}_k$ where \mathbf{x}_k denotes the bit string with k 1s followed by $n - k$ 0s.

Find each of the following

1. $f \circ g$
2. $g \circ f$

Answer

1. $f \circ g : T \rightarrow T$ and $(f \circ g)(k) = k$.
2. $g \circ f : S \rightarrow S$ and $(g \circ f)(\mathbf{x}) = \mathbf{x}_k$ where $k = f(\mathbf{x}) = \sum_{i=1}^n x_i$. In words, $(g \circ f)(\mathbf{x})$ is the bit string with the same number of 1s as \mathbf{x} , but rearranged so that all the 1s come first.

In the previous exercise, note that $f \circ g$ and $g \circ f$ are both well-defined, but have different domains, and so of course are not the same. Note also that $f \circ g$ is the identity function on T , but f is not the inverse of g . Indeed f is not one-to-one, and so does not have an inverse. However, f restricted to $\{\mathbf{x}_k : k \in T\}$ (the range of g) is one-to-one and is the inverse of g .

Let $n = 4$. Give $f^{-1}(\{k\})$ in list form for each $k \in T$.

Answer

1. $f^{-1}(\{0\}) = \{0000\}$
2. $f^{-1}(\{1\}) = \{1000, 0100, 0010, 0001\}$
3. $f^{-1}(\{2\}) = \{1100, 1010, 1001, 0110, 0101, 0011\}$
4. $f^{-1}(\{3\}) = \{1110, 1101, 1011, 0111\}$
5. $f^{-1}(\{4\}) = \{1111\}$

Again let $n = 4$. Let $A = \{1000, 1010\}$ and $B = \{1000, 1100\}$. Give each of the following in list form:

1. $f(A)$
2. $f(B)$
3. $f(A \cap B)$
4. $f(A) \cap f(B)$
5. $f^{-1}(f(A))$

Answer

1. $\{1, 2\}$
2. $\{1, 2\}$
3. $\{1\}$
4. $\{1, 2\}$
5. $\{1000, 0100, 0010, 0001, 1100, 1010, 1001, 0110, 0101, 0011\}$

In the previous exercise, note that $f(A \cap B) \subset f(A) \cap f(B)$ and $A \subset f^{-1}(f(A))$.

Indicator Functions

Suppose that A and B are subsets of a universal set S . Express, in terms of $\mathbf{1}_A$ and $\mathbf{1}_B$, the indicator function of each of the 14 non-trivial sets that can be constructed from A and B . Use the Venn diagram app to help.

Answer

1. $\mathbf{1}_A$
2. $\mathbf{1}_B$
3. $\mathbf{1}_{A^c} = 1 - \mathbf{1}_A$
4. $\mathbf{1}_{B^c} = 1 - \mathbf{1}_B$
5. $\mathbf{1}_{A \cap B} = \mathbf{1}_A \mathbf{1}_B$
6. $\mathbf{1}_{A \cup B} = \mathbf{1}_A + \mathbf{1}_B - \mathbf{1}_A \mathbf{1}_B$
7. $\mathbf{1}_{A \cap B^c} = \mathbf{1}_A - \mathbf{1}_A \mathbf{1}_B$
8. $\mathbf{1}_{B \cap A^c} = \mathbf{1}_B - \mathbf{1}_A \mathbf{1}_B$
9. $\mathbf{1}_{A \cup B^c} = 1 - \mathbf{1}_B + \mathbf{1}_A \mathbf{1}_B$
10. $\mathbf{1}_{B \cup A^c} = 1 - \mathbf{1}_A + \mathbf{1}_A \mathbf{1}_B$
11. $\mathbf{1}_{A^c \cap B^c} = 1 - \mathbf{1}_A - \mathbf{1}_B + \mathbf{1}_A \mathbf{1}_B$
12. $\mathbf{1}_{A^c \cup B^c} = 1 - \mathbf{1}_A \mathbf{1}_B$
13. $\mathbf{1}_{(A \cap B^c) \cup (B \cap A^c)} = \mathbf{1}_A + \mathbf{1}_B - 2\mathbf{1}_A \mathbf{1}_B$
14. $\mathbf{1}_{(A \cap B) \cup (A^c \cap B^c)} = 1 - \mathbf{1}_A - \mathbf{1}_B + 2\mathbf{1}_A \mathbf{1}_B$

Suppose that A , B , and C are subsets of a universal set S . Give the indicator function of each of the following, in terms of $\mathbf{1}_A$, $\mathbf{1}_B$, and $\mathbf{1}_C$ in sum-product form:

1. $D = \{x \in S : x \text{ is an element of exactly one of the given sets}\}$
2. $E = \{x \in S : x \text{ is an element of exactly two of the given sets}\}$

Answer

1. $\mathbf{1}_D = \mathbf{1}_A + \mathbf{1}_B + \mathbf{1}_C - 2(\mathbf{1}_A \mathbf{1}_B + \mathbf{1}_A \mathbf{1}_C + \mathbf{1}_B \mathbf{1}_C) + 3\mathbf{1}_A \mathbf{1}_B \mathbf{1}_C$
2. $\mathbf{1}_E = \mathbf{1}_A \mathbf{1}_B + \mathbf{1}_A \mathbf{1}_C + \mathbf{1}_B \mathbf{1}_C - 3\mathbf{1}_A \mathbf{1}_B \mathbf{1}_C$

Operators

Recall the standard arithmetic operators on \mathbb{R} discussed above.

We all know that sum is commutative and associative, and that product distributes over sum.

1. Is difference commutative?
2. Is difference associative?
3. Does product distribute over difference?
4. Does sum distributed over product?

Answer

1. No. $x - y \neq y - x$
2. No. $x - (y - z) \neq (x - y) - z$
3. Yes. $x(y - z) = (xy) - (xz)$
4. No. $x + (yz) \neq (x + y)(x + z)$

Multisets

Express the multiset A in list form that has the multiplicity function $m : \{a, b, c, d, e\} \rightarrow \mathbb{N}$ given by $m(a) = 2$, $m(b) = 3$, $m(c) = 1$, $m(d) = 0$, $m(e) = 4$.

Answer

$$A = \{a, a, b, b, b, c, e, e, e, e\}$$

Express the prime factors of 360 as a multiset in list form.

Answer

$$\{2, 2, 2, 3, 3, 5\}$$

This page titled [1.2: Functions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.3: Relations

Relations play a fundamental role in probability theory, as in most other areas of mathematics.

Definitions and Constructions

Suppose that S and T are sets. A *relation* from S to T is a subset of the product set $S \times T$.

1. The *domain* of R is the set of first coordinates: $\text{domain}(R) = \{x \in S : (x, y) \in R \text{ for some } y \in T\}$.
2. The *range* of R is the set of second coordinates: $\text{range}(R) = \{y \in T : (x, y) \in R \text{ for some } x \in S\}$.

A relation from a set S to itself is a relation *on* S .

As the name suggests, a relation R from S into T is supposed to define a *relationship* between the elements of S and the elements of T , and so we usually use the more suggestive notation $x R y$ when $(x, y) \in R$. Note that the domain of R is the projections of R onto S and the range of R is the projection of R onto T .

Basic Examples

Suppose that S is a set and recall that $\mathcal{P}(S)$ denotes the power set of S , the collection of all subsets of S . The *membership relation* \in from S to $\mathcal{P}(S)$ is perhaps the most important and basic relationship in mathematics. Indeed, for us, it's a primitive (undefined) relationship—given x and A , we assume that we understand the meaning of the statement $x \in A$.

Another basic primitive relation is the *equality relation* $=$ on a given set of objects S . That is, given two objects x and y , we assume that we understand the meaning of the statement $x = y$.

Other basic relations that you have seen are

1. The subset relation \subseteq on $\mathcal{P}(S)$.
2. The order relation \leq on \mathbb{R}

These two belong to a special class of relations known as partial orders that we will study in the next section. Note that a function f from S into T is a special type of relation. To compare the two types of notation (relation and function), note that $x f y$ means that $y = f(x)$.

Constructions

Since a relation is just a set of ordered pairs, the set operations can be used to build new relations from existing ones.

if Q and R are relations from S to T , then so are $Q \cup R$, $Q \cap R$, $Q \setminus R$.

1. $x(Q \cup R)y$ if and only if $x Q y$ or $x R y$.
2. $x(Q \cap R)y$ if and only if $x Q y$ and $x R y$.
3. $x(Q \setminus R)y$ if and only if $x Q y$ but not $x R y$.
4. If $Q \subseteq R$ then $x Q y$ implies $x R y$.

If R is a relation from S to T and $Q \subseteq R$, then Q is a relation from S to T .

The restriction of a relation defines a new relation.

If R is a relation on S and $A \subseteq S$ then $R_A = R \cap (A \times A)$ is a relation on A , called the *restriction* of R to A .

The inverse of a relation also defines a new relation.

If R is a relation from S to T , the *inverse* of R is the relation from T to S defined by

$$y R^{-1} x \text{ if and only if } x R y \quad (1.3.1)$$

Equivalently, $R^{-1} = \{(y, x) : (x, y) \in R\}$. Note that any function f from S into T has an inverse *relation*, but only when the f is one-to-one is the inverse relation also a function (the inverse function). Composition is another natural way to create new relations

from existing ones.

Suppose that Q is a relation from S to T and that R is a relation from T to U . The *composition* $Q \circ R$ is the relation from S to U defined as follows: for $x \in S$ and $z \in U$, $x(Q \circ R)z$ if and only if there exists $y \in T$ such that $x Q y$ and $y R z$.

Note that the notation is inconsistent with the notation used for composition of functions, essentially because relations are read from left to right, while functions are read from right to left. Hopefully, the inconsistency will not cause confusion, since we will always use function notation for functions.

Basic Properties

The important classes of relations that we will study in the next couple of sections are characterized by certain basic properties. Here are the definitions:

Suppose that R is a relation on S .

1. R is *reflexive* if $x R x$ for all $x \in S$.
2. R is *irreflexive* if no $x \in S$ satisfies $x R x$.
3. R is *symmetric* if $x R y$ implies $y R x$ for all $x, y \in S$.
4. R is *anti-symmetric* if $x R y$ and $y R x$ implies $x = y$ for all $x, y \in S$.
5. R is *transitive* if $x R y$ and $y R z$ implies $x R z$ for all $x, y, z \in S$.

The proofs of the following results are straightforward, so be sure to try them yourself before reading the ones in the text.

A relation R on S is reflexive if and only if the equality relation $=$ on S is a subset of R .

Proof

This follows from the definitions. R is reflexive if and only if $(x, x) \in R$ for all $x \in S$.

A relation R on S is symmetric if and only if $R^{-1} = R$.

Proof

Suppose that R is symmetric. If $(x, y) \in R$ then $(y, x) \in R$ and hence $(x, y) \in R^{-1}$. If $(x, y) \in R^{-1}$ then $(y, x) \in R$ and hence $(x, y) \in R$. Thus $R = R^{-1}$. Conversely, suppose $R = R^{-1}$. If $(x, y) \in R$ then $(x, y) \in R^{-1}$ and hence $(y, x) \in R$.

A relation R on S is transitive if and only if $R \circ R \subseteq R$.

Proof

Suppose that R is transitive. If $(x, z) \in R \circ R$ then there exists $y \in S$ such that $(x, y) \in R$ and $(y, z) \in R$. But then $(x, z) \in R$ by transitivity. Hence $R \circ R \subseteq R$. Conversely, suppose that $R \circ R \subseteq R$. If $(x, y) \in R$ and $(y, z) \in R$ then $(x, z) \in R \circ R$ and hence $(x, z) \in R$. Hence R is transitive.

A relation R on S is antisymmetric if and only if $R \cap R^{-1}$ is a subset of the equality relation $=$ on S .

Proof

Restated, this result is that R is antisymmetric if and only if $(x, y) \in R \cap R^{-1}$ implies $x = y$. Thus suppose that R is antisymmetric. If $(x, y) \in R \cap R^{-1}$ then $(x, y) \in R$ and $(x, y) \in R^{-1}$. But then $(y, x) \in R$ so by antisymmetry, $x = y$. Conversely suppose that $(x, y) \in R \cap R^{-1}$ implies $x = y$. If $(x, y) \in R$ and $(y, x) \in R$ then $(x, y) \in R^{-1}$ and hence $(x, y) \in R \cap R^{-1}$. Thus $x = y$ so R is antisymmetric.

Suppose that Q and R are relations on S . For each property below, if both Q and R have the property, then so does $Q \cap R$.

1. reflexive
2. symmetric
3. transitive

Proof

1. Suppose that Q and R are reflexive. Then $(x, x) \in Q$ and $(x, x) \in R$ for each $x \in S$ and hence $(x, x) \in Q \cap R$ for each $x \in S$. Thus $Q \cap R$ is reflexive.
2. Suppose that Q and R are symmetric. If $(x, y) \in Q \cap R$ then $(x, y) \in Q$ and $(x, y) \in R$. Hence $(y, x) \in Q$ and $(y, x) \in R$ so $(y, x) \in Q \cap R$. Hence $Q \cap R$ is symmetric.
3. Suppose that Q and R are transitive. If $(x, y) \in Q \cap R$ and $(y, z) \in Q \cap R$ then $(x, y) \in Q$, $(x, y) \in R$, $(y, z) \in Q$, and $(y, z) \in R$. Hence $(x, z) \in Q$ and $(x, z) \in R$ so $(x, z) \in Q \cap R$. Hence $Q \cap R$ is transitive.

Suppose that R is a relation on a set S .

1. Give an explicit definition for the property R is not reflexive.
2. Give an explicit definition for the property R is not irreflexive.
3. Are any of the properties R is reflexive, R is not reflexive, R is irreflexive, R is not irreflexive equivalent?

Answer

1. R is not reflexive if and only if there exists $x \in S$ such that $(x, x) \notin R$.
2. R is not irreflexive if and only if there exists $x \in S$ such that $(x, x) \in R$.
3. Nope.

Suppose that R is a relation on a set S .

1. Give an explicit definition for the property R is not symmetric.
2. Give an explicit definition for the property R is not antisymmetric.
3. Are any of the properties R is symmetric, R is not symmetric, R is antisymmetric, R is not antisymmetric equivalent?

Answer

1. R is not symmetric if and only if there exist $x, y \in S$ such that $(x, y) \in R$ and $(y, x) \notin R$.
2. R is not antisymmetric if and only if there exist distinct $x, y \in S$ such that $(x, y) \in R$ and $(y, x) \in R$.
3. Nope.

Computational Exercises

Let R be the relation defined on \mathbb{R} by $x R y$ if and only if $\sin(x) = \sin(y)$. Determine if R has each of the following properties:

1. reflexive
2. symmetric
3. transitive
4. irreflexive
5. antisymmetric

Answer

1. yes
2. yes
3. yes
4. no
5. no

The relation R in the previous exercise is a member of an important class of equivalence relations.

Let R be the relation defined on \mathbb{R} by $x R y$ if and only if $x^2 + y^2 \leq 1$. Determine if R has each of the following properties:

1. reflexive
2. symmetric
3. transitive
4. irreflexive
5. antisymmetric

Answer

1. no
2. yes
3. no
4. no
5. no

This page titled [1.3: Relations](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.4: Partial Orders

Partial orders are a special class of relations that play an important role in probability theory.

Basic Theory

Definitions

A *partial order* on a set S is a relation \preceq on S that is reflexive, anti-symmetric, and transitive. The pair (S, \preceq) is called a *partially ordered set*. So for all $x, y, z \in S$:

1. $x \preceq x$, the *reflexive property*
2. If $x \preceq y$ and $y \preceq x$ then $x = y$, the *antisymmetric property*
3. If $x \preceq y$ and $y \preceq z$ then $x \preceq z$, the *transitive property*

As the name and notation suggest, a partial order is a type of ordering of the elements of S . Partial orders occur naturally in many areas of mathematics, including probability. A partial order on a set naturally gives rise to several other relations on the set.

Suppose that \preceq is a partial order on a set S . The relations $\succeq, \prec, \succ, \perp$, and \parallel are defined as follows:

1. $x \succeq y$ if and only if $y \preceq x$.
2. $x \prec y$ if and only if $x \preceq y$ and $x \neq y$.
3. $x \succ y$ if and only if $y \prec x$.
4. $x \perp y$ if and only if $x \preceq y$ or $y \preceq x$.
5. $x \parallel y$ if and only if neither $x \preceq y$ nor $y \preceq x$.

Note that \succeq is the inverse of \preceq , and \succ is the inverse of \prec . Note also that $x \preceq y$ if and only if either $x \prec y$ or $x = y$, so the relation \prec completely determines the relation \preceq . The relation \prec is sometimes called a *strict* or *strong* partial order to distinguish it from the ordinary (weak) partial order \preceq . Finally, note that $x \perp y$ means that x and y are *related* in the partial order, while $x \parallel y$ means that x and y are *unrelated* in the partial order. Thus, the relations \perp and \parallel are complements of each other, as sets of ordered pairs. A total or linear order is a partial order in which there are no unrelated elements.

A partial order \preceq on S is a *total order* or *linear order* if for every $x, y \in S$, either $x \preceq y$ or $y \preceq x$.

Suppose that \preceq_1 and \preceq_2 are partial orders on a set S . Then \preceq_1 is a *sub-order* of \preceq_2 , or equivalently \preceq_2 is an *extension* of \preceq_1 if and only if $x \preceq_1 y$ implies $x \preceq_2 y$ for $x, y \in S$.

Thus if \preceq_1 is a suborder of \preceq_2 , then as sets of ordered pairs, \preceq_1 is a subset of \preceq_2 . We need one more relation that arises naturally from a partial order.

Suppose that \preceq is a partial order on a set S . For $x, y \in S$, y is said to *cover* x if $x \prec y$ but no element $z \in S$ satisfies $x \prec z \prec y$.

If S is finite, the *covering relation* completely determines the partial order, by virtue of the transitive property.

Suppose that \preceq is a partial order on a finite set S . The *covering graph* or *Hasse graph* of (S, \preceq) is the directed graph with vertex set S and directed edge set E , where $(x, y) \in E$ if and only if y covers x .

Thus, $x \prec y$ if and only if there is a directed path in the graph from x to y . Hasse graphs are named for the German mathematician Helmut Hasse. The graphs are often drawn with the edges directed upward. In this way, the directions can be inferred without having to actually draw arrows.

Basic Examples

Of course, the ordinary order \leq is a total order on the set of real numbers \mathbb{R} . The subset partial order is one of the most important in probability theory:

Suppose that S is a set. The *subset relation* \subseteq is a partial order on $\mathcal{P}(S)$, the power set of S .

Proof

We proved this result in the section on sets. To review, recall that for $A, B \in \mathcal{P}(S)$, $A \subseteq B$ means that $x \in A$ implies $x \in B$. Also $A = B$ means that $x \in A$ if and only if $x \in B$. Thus

1. $A \subseteq A$
2. $A \subseteq B$ and $B \subseteq A$ if and only if $A = B$
3. $A \subseteq B$ and $B \subseteq C$ imply $A \subseteq C$

Here is a partial order that arises naturally from arithmetic.

Let $|$ denote the *division relation* on the set of positive integers \mathbb{N}_+ . That is, $m | n$ if and only if there exists $k \in \mathbb{N}_+$ such that $n = km$. Then

1. $|$ is a partial order on \mathbb{N}_+ .
2. $|$ is a sub-order of the ordinary order \leq .

Proof

1. Clearly $n | n$ for $n \in \mathbb{N}_+$, since $n = 1 \cdot n$, so $|$ is reflexive. Suppose $m | n$ and $n | m$, where $m, n \in \mathbb{N}_+$. Then there exist $j, k \in \mathbb{N}_+$ such that $n = km$ and $m = jn$. Substituting gives $n = jkn$, and hence $j = k = 1$. Thus $m = n$ so $|$ is antisymmetric. Finally, suppose $m | n$ and $n | p$, where $m, n, p \in \mathbb{N}_+$. Then there exists $j, k \in \mathbb{N}_+$ such that $n = jm$ and $p = kn$. Substituting gives $p = jkm$, so $m | p$. Thus $|$ is transitive.
2. If $m, n \in \mathbb{N}_+$ and $m | n$, then there exists $k \in \mathbb{N}_+$ such that $n = km$. Since $k \geq 1$, $m \leq n$.

The set of functions from a set into a partial ordered set can itself be partially ordered in a natural way.

Suppose that S is a set and that (T, \preceq_T) is a partially ordered set, and let \mathcal{S} denote the set of functions $f : S \rightarrow T$. The relation \preceq on \mathcal{S} defined by $f \preceq g$ if and only if $f(x) \preceq_T g(x)$ for all $x \in S$ is a partial order on \mathcal{S} .

Proof

Suppose that $f, g, h \in \mathcal{S}$.

1. $f(x) \preceq_T f(x)$ for all $x \in S$, so $f \preceq f$.
2. If $f \preceq g$ and $g \preceq f$ then $f(x) \preceq_T g(x)$ and $g(x) \preceq_T f(x)$ for all $x \in S$. Hence $f(x) = g(x)$ for all $x \in S$ so $f = g$.
3. If $f \preceq g$ and $g \preceq h$ then $f(x) \preceq_T g(x)$ and $g(x) \preceq_T h(x)$ for all $x \in S$. Hence $f(x) \preceq_T h(x)$ for all $x \in S$ so $f \preceq h$.

Note that we don't need a partial order on the domain S .

Basic Properties

The proofs of the following basic properties are straightforward. Be sure to try them yourself before reading the ones in the text.

The inverse of a partial order is also a partial order.

Proof

Clearly the reflexive, antisymmetric and transitive properties hold for \succeq .

If \preceq is a partial order on S and A is a subset of S , then the restriction of \preceq to A is a partial order on A .

Proof

The reflexive, antisymmetric, and transitive properties given above hold for all $x, y, z \in S$ and hence hold for all $x, y, z \in A$.

The following theorem characterizes relations that correspond to strict order.

Let S be a set. A relation \preceq is a partial order on S if and only if \prec is transitive and irreflexive.

Proof

Suppose that \preceq is a partial order on S . Recall that \prec is defined by $x \prec y$ if and only if $x \preceq y$ and $x \neq y$. If $x \prec y$ and $y \prec z$ then $x \preceq y$ and $y \preceq z$, and so $x \preceq z$. On the other hand, if $x = z$ then $x \preceq y$ and $y \preceq x$ so $x = y$, a contradiction. Hence $x \neq z$ and so $x \prec z$. Therefore \prec is transitive. If $x \prec y$ then $x \neq y$ by definition, so \prec is irreflexive.

Conversely, suppose that \prec is a transitive and irreflexive relation on S . Recall that \preceq is defined by $x \preceq y$ if and only if $x \prec y$ or $x = y$. By definition then, \preceq is reflexive: $x \preceq x$ for every $x \in S$. Next, suppose that $x \preceq y$ and $y \preceq x$. If $x \prec y$ and $y \prec x$ then

$x \prec x$ by the transitive property of \prec . But this is a contradiction by the irreflexive property, so we must have $x = y$. Thus \preceq is antisymmetric. Suppose $x \preceq y$ and $y \preceq z$. There are four cases:

1. If $x \prec y$ and $y \prec z$ then $x \prec z$ by the transitive property of \prec .
2. If $x = y$ and $y \prec z$ then $x \prec z$ by substitution.
3. If $x \prec y$ and $y = z$ then $x \prec z$ by substitution.
4. If $x = y$ and $y = z$ then $x = z$ by the transitive property of $=$.

In all cases we have $x \preceq z$ so \preceq is transitive. Hence \preceq is a partial order on S .

Monotone Sets and Functions

Partial orders form a natural setting for increasing and decreasing sets and functions. Here are the definitions:

Suppose that \preceq is a partial order on a set S and that $A \subseteq S$. In the following definitions, x, y are arbitrary elements of S .

1. A is *increasing* if $x \in A$ and $x \preceq y$ imply $y \in A$.
2. A is *decreasing* if $y \in A$ and $x \preceq y$ imply $x \in A$.

Suppose that S is a set with partial order \preceq_S , T is a set with partial order \preceq_T , and that $f : S \rightarrow T$. In the following definitions, x, y are arbitrary elements of S .

1. f is *increasing* if and only if $x \preceq_S y$ implies $f(x) \preceq_T f(y)$.
2. f is *decreasing* if and only if $x \preceq_S y$ implies $f(x) \succeq_T f(y)$.
3. f is *strictly increasing* if and only if $x \prec_S y$ implies $f(x) \prec_T f(y)$.
4. f is *strictly decreasing* if and only if $x \prec_S y$ implies $f(x) \succ_T f(y)$.

Recall the definition of the **indicator function** $\mathbf{1}_A$ associated with a subset A of a universal set S : For $x \in S$, $\mathbf{1}_A(x) = 1$ if $x \in A$ and $\mathbf{1}_A(x) = 0$ if $x \notin A$.

Suppose that \preceq is a partial order on a set S and that $A \subseteq S$. Then

1. A is increasing if and only if $\mathbf{1}_A$ is increasing.
2. A is decreasing if and only if $\mathbf{1}_A$ is decreasing.

Proof

1. A is increasing if and only if $x \in A$ and $x \preceq y$ implies $y \in A$ if and only if $\mathbf{1}_A(x) = 1$ and $x \preceq y$ implies $\mathbf{1}_A(y) = 1$ if and only if $\mathbf{1}_A$ is increasing.
2. A is decreasing if and only if $y \in A$ and $x \preceq y$ implies $x \in A$ if and only if $\mathbf{1}_A(y) = 1$ and $x \preceq y$ implies $\mathbf{1}_A(x) = 1$ if and only if $\mathbf{1}_A$ is decreasing.

Isomorphism

Two partially ordered sets (S, \preceq_S) and (T, \preceq_T) are said to be *isomorphic* if there exists a one-to-one function f from S onto T such that $x \preceq_S y$ if and only if $f(x) \preceq_T f(y)$, for all $x, y \in S$. The function f is an *isomorphism*.

Generally, a mathematical *space* often consists of a set and various structures defined in terms of the set, such as relations, operators, or a collection of subsets. Loosely speaking, two mathematical spaces of the same type are *isomorphic* if there exists a one-to-one function from one of the sets onto the other that *preserves* the structures, and again, the function is called an *isomorphism*. The basic idea is that isomorphic spaces are mathematically identical, except for superficial matters of appearance. The word *isomorphism* is from the Greek and means *equal shape*.

Suppose that the partially ordered sets (S, \preceq_S) and (T, \preceq_T) are isomorphic, and that $f : S \rightarrow T$ is an isomorphism. Then f and f^{-1} are strictly increasing.

Proof

We need to show that for $x, y \in S$, $x \prec_S y$ if and only if $f(x) \prec_T f(y)$. If $x \prec_S y$ then by definition, $f(x) \preceq_T f(y)$. But if $f(x) = f(y)$ then $x = y$ since f is one-to-one. This is a contradiction, so $f(x) \prec_T f(y)$. Similarly, if $f(x) \prec_T f(y)$ then by definition, $x \preceq_S y$. But if $x = y$ then $f(x) = f(y)$, a contradiction. Hence $x \prec_S y$.

In a sense, the subset partial order is universal—every partially ordered set is isomorphic to (\mathcal{S}, \subseteq) for some collection of sets \mathcal{S} .

Suppose that \preceq is a partial order on a set S . Then there exists $\mathcal{S} \subseteq \mathcal{P}(S)$ such that (S, \preceq) is isomorphic to (\mathcal{S}, \subseteq) .

Proof

For each $x \in S$, let $A_x = \{u \in S : u \preceq x\}$, and then let $\mathcal{S} = \{A_x : x \in S\}$, so that $\mathcal{S} \subseteq \mathcal{P}(S)$. We will show that the function $x \mapsto A_x$ from S onto \mathcal{S} is one-to-one, and satisfies

$$x \preceq y \iff A_x \subseteq A_y \quad (1.4.1)$$

First, suppose that $x, y \in S$ and $A_x = A_y$. Then $x \in A_x$ so $x \in A_y$ and hence $x \preceq y$. Similarly, $y \in A_y$ so $y \in A_x$ and hence $y \preceq x$. Thus $x = y$, so the mapping is one-to-one. Next, suppose that $x \preceq y$. If $u \in A_x$ then $u \preceq x$ so $u \preceq y$ by the transitive property, and hence $u \in A_y$. Thus $A_x \subseteq A_y$. Conversely, suppose $A_x \subseteq A_y$. As before, $x \in A_x$, so $x \in A_y$ and hence $x \preceq y$.

Extremal Elements

Various types of *extremal elements* play important roles in partially ordered sets. Here are the definitions:

Suppose that \preceq is a partial order on a set S and that $A \subseteq S$.

1. An element $a \in A$ is the *minimum* element of A if and only if $a \preceq x$ for every $x \in A$.
2. An element $a \in A$ is a *minimal* element of A if and only if no $x \in A$ satisfies $x \prec a$.
3. An element $b \in A$ is the *maximum* element of A if and only if $b \succeq x$ for every $x \in A$.
4. An element $b \in A$ is a *maximal* element of A if and only if no $x \in A$ satisfies $x \succ b$.

In general, a set can have several maximal and minimal elements (or none). On the other hand,

The minimum and maximum elements of A , if they exist, are unique. They are denoted $\min(A)$ and $\max(A)$, respectively.

Proof

Suppose that a, b are minimum elements of A . Since $a, b \in A$ we have $a \preceq b$ and $b \preceq a$, so $a = b$ by the antisymmetric property. The proof for the maximum element is analogous.

Minimal, maximal, minimum, and maximum elements of a set must belong to that set. The following definitions relate to upper and lower bounds of a set, which do not have to belong to the set.

Suppose again that \preceq is a partial order on a set S and that $A \subseteq S$. Then

1. An element $u \in S$ is a *lower bound* for A if and only if $u \preceq x$ for every $x \in A$.
2. An element $v \in S$ is an *upper bound* for A if and only if $v \succeq x$ for every $x \in A$.
3. The *greatest lower bound* or *infimum* of A , if it exists, is the maximum of the set of lower bounds of A .
4. The *least upper bound* or *supremum* of A , if it exists, is the minimum of the set of upper bounds of A .

By (20), the greatest lower bound of A is unique, if it exists. It is denoted $\text{glb}(A)$ or $\inf(A)$. Similarly, the least upper bound of A is unique, if it exists, and is denoted $\text{lub}(A)$ or $\sup(A)$. Note that every element of S is a lower bound and an upper bound for \emptyset , since the conditions in the definition hold vacuously.

The symbols \wedge and \vee are also used for infimum and supremum, respectively, so $\bigwedge A = \inf(A)$ and $\bigvee A = \sup(A)$ if they exist. In particular, for $x, y \in S$, operator notation is more commonly used, so $x \wedge y = \inf\{x, y\}$ and $x \vee y = \sup\{x, y\}$. Partially ordered sets for which these elements always exist are important, and have a special name.

Suppose that \preceq is a partial order on a set S . Then (S, \preceq) is a *lattice* if $x \wedge y$ and $x \vee y$ exist for every $x, y \in S$.

For the subset partial order, the \inf and \sup operators correspond to intersection and union, respectively:

Let S be a set and consider the subset partial order \subseteq on $\mathcal{P}(S)$, the power set of S . Let \mathcal{A} be a nonempty subset of $\mathcal{P}(S)$, that is, a nonempty collection of subsets of S . Then

1. $\inf(\mathcal{A}) = \bigcap \mathcal{A}$
2. $\sup(\mathcal{A}) = \bigcup \mathcal{A}$

Proof

1. First, $\bigcap \mathcal{A} \subseteq A$ for every $A \in \mathcal{A}$ and hence $\bigcap \mathcal{A}$ is a lower bound of \mathcal{A} . If B is a lower bound of \mathcal{A} then $B \subseteq A$ for every $A \in \mathcal{A}$ and hence $B \subseteq \bigcap \mathcal{A}$. Therefore $\bigcap \mathcal{A}$ is the greatest lower bound.
2. First, $A \subseteq \bigcup \mathcal{A}$ for every $A \in \mathcal{A}$ and hence $\bigcup \mathcal{A}$ is an upper bound of \mathcal{A} . If B is an upper bound of \mathcal{A} then $A \subseteq B$ for every $A \in \mathcal{A}$ and hence $\bigcup \mathcal{A} \subseteq B$. Therefore $\bigcup \mathcal{A}$ is the least upper bound.

In particular, $A \wedge B = A \cap B$ and $A \vee B = A \cup B$, so $(\mathcal{P}(S), \subseteq)$ is a lattice.

Consider the division partial order $|$ on the set of positive integers \mathbb{N}_+ and let A be a nonempty subset of \mathbb{N}_+ .

1. $\inf(A)$ is the *greatest common divisor* of A , usually denoted $\gcd(A)$ in this context.
2. If A is infinite then $\sup(A)$ does not exist. If A is finite then $\sup(A)$ is the *least common multiple* of A , usually denoted $\text{lcm}(A)$ in this context.

Suppose that S is a set and that $f : S \rightarrow S$. An element $z \in S$ is said to be a *fixed point* of f if $f(z) = z$.

The following result explores a basic fixed point theorem for a partially ordered set. The theorem is important in the study of [cardinality](#).

Suppose that \preceq is a partial order on a set S with the property that $\sup(A)$ exists for every $A \subseteq S$. If $f : S \rightarrow S$ is increasing, then f has a fixed point.

Proof.

Let $A = \{x \in S : x \preceq f(x)\}$ and let $z = \sup(A)$. If $x \in A$ then $x \preceq z$ so $x \preceq f(x) \preceq f(z)$. Hence $f(z)$ is an upper bound of A so $z \preceq f(z)$. But then $f(z) \preceq f(f(z))$ so $f(z) \in A$. Hence $f(z) \preceq z$. Therefore $f(z) = z$.

Note that the hypotheses of the theorem require that $\sup(\emptyset) = \min(S)$ exists. The set $A = \{x \in S : x \preceq f(x)\}$ is nonempty since $\min(S) \in A$.

If \preceq is a total order on a set S with the property that every nonempty subset of S has a minimum element, then S is said to be *well ordered* by \preceq . One of the most important examples is \mathbb{N}_+ , which is well ordered by the ordinary order \leq . On the other hand, the *well ordering principle*, which is equivalent to the [axiom of choice](#), states that every nonempty set can be well ordered.

Orders on Product Spaces

Suppose that S and T are sets with partial orders \preceq_S and \preceq_T respectively. Define the relation \preceq on $S \times T$ by $(x, y) \preceq (z, w)$ if and only if $x \preceq_S z$ and $y \preceq_T w$.

1. The relation \preceq is a partial order on $S \times T$, called, appropriately enough, the *product order*.
2. Suppose that $(S, \preceq_S) = (T, \preceq_T)$. If S has at least 2 elements, then \preceq is not a total order on S^2 .

Proof




Figure 1.4.1: The product order on \mathbb{R}^2 . The region shaded red is the set of points $\succeq (x, y)$. The region shaded blue is the set of points $\preceq (x, y)$. The region shaded white is the set of points that are not comparable with (x, y) .

Product order extends in a straightforward way to the Cartesian product of a finite or an infinite sequence of partially ordered spaces. For example, suppose that S_i is a set with partial order \preceq_i for each $i \in \{1, 2, \dots, n\}$, where $n \in \mathbb{N}_+$. The product order \preceq on the product set $S_1 \times S_2 \times \dots \times S_n$ is defined as follows: for $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and $\mathbf{y} = (y_1, y_2, \dots, y_n)$ in the product set, $\mathbf{x} \preceq \mathbf{y}$ if and only if $x_i \preceq_i y_i$ for each $i \in \{1, 2, \dots, n\}$. We can generalize this further to arbitrary [product sets](#). Suppose that S_i is a set for each i in a nonempty (both otherwise arbitrary) index set I . Recall that

$$\prod_{i \in I} S_i = \left\{ x : x \text{ is a function from } I \text{ into } \bigcup_{i \in I} S_i \text{ such that } x(i) \in S_i \text{ for each } i \in I \right\} \quad (1.4.2)$$

To make the notation look more like a simple Cartesian product, we will write x_i instead of $x(i)$ for the value of a function x in the product set at $i \in I$.

Suppose that S_i is a set with partial order \preceq_i for each i in a nonempty index set I . Define the relation \preceq on $\prod_{i \in I} S_i$ by $x \preceq y$ if and only if $x_i \preceq_i y_i$ for each $i \in I$. Then \preceq is a partial order on the product set, known again as the *product order*.

Proof

In spite of the abstraction, the proof is perfectly straightforward. Suppose that $x, y, z \in \prod_{i \in I} S_i$.

1. $x_i \preceq_i x_i$ for every $i \in I$, and hence $x \preceq x$. Thus \preceq is reflexive.
2. Suppose that $x \preceq y$ and $y \preceq x$. Then $x_i \preceq_i y_i$ and $y_i \preceq_i x_i$ for each $i \in I$. Hence $x_i = y_i$ for each $i \in I$ and so $x = y$. Thus \preceq is antisymmetric.
3. Suppose that $x \preceq y$ and $y \preceq z$. Then $x_i \preceq_i y_i$ and $y_i \preceq_i z_i$ for each $i \in I$. Hence $x_i \preceq_i z_i$ for each $i \in I$, so $x \preceq z$. Thus \preceq is transitive.

Note again that no assumptions are made on the index set I , other than it be nonempty. In particular, no order is necessary on I . The next result gives a very different type of order on a product space.

Suppose again that S and T are sets with partial orders \preceq_S and \preceq_T respectively. Define the relation \preceq on $S \times T$ by $(x, y) \preceq (z, w)$ if and only if either $x \prec_S z$, or $x = z$ and $y \preceq_T w$.

1. The relation \preceq is a partial order on $S \times T$, called the *lexicographic order* or *dictionary order*.
2. If \preceq_S and \preceq_T are total orders on S and T , respectively, then \preceq is a total order on $S \times T$.

Proof




Figure 1.4.2: The lexicographic order on \mathbb{R}^2 . The region shaded red is the set of points $\succeq (x, y)$. The region shaded blue is the set of points $\preceq (x, y)$.

As with the product order, the lexicographic order can be generalized to a collection of partially ordered spaces. However, we need the index set to be totally ordered.

Suppose that S_i is a set with partial order \preceq_i for each i in a nonempty index set I . Suppose also that \leq is a total order on I . Define the relation \preceq on the product set $\prod_{i \in I} S_i$ as follows: $x \prec y$ if and only if there exists $j \in I$ such that $x_i = y_i$ if $i < j$ and $x_j \prec_j y_j$. Then

1. \preceq is a partial order on S , known again as the *lexicographic order*.
2. If \preceq_i is a total order for each $i \in I$, and I is well ordered by \leq , then \preceq is a total order on S .

Proof

1. By the result on [strong orders](#), we need to show that \prec is irreflexive and transitive. First, no $x \in \prod_{i \in I} S_i$ satisfies $x \prec x$ since $x_i = x_i$ for all $i \in I$. Hence \prec is irreflexive. Next, suppose that $x, y, z \in \prod_{i \in I} S_i$ and that $x \prec y$ and $y \prec z$. Then there exists $j \in I$ such that $x_i = y_i$ if $i < j$ and $x_j \prec_j y_j$. Similarly, there exists $k \in I$ such that $y_i = z_i$ if $i < k$ and $y_k \prec_k z_k$. Again, since I is totally ordered, either $j < k$ or $k < j$ or $j = k$. If $j < k$, then $x_i = y_i = z_i$ if $i < j$ and $x_j \prec_j y_j = z_j$. If $k < j$, then $x_i = y_i = z_i$ if $i < k$ and $x_k = y_k \prec_k z_k$. If $j = k$, then $x_i = y_i = z_i$ if $i < j$ and $x_j \prec_j y_j \prec_j z_j$. In all cases, $x \prec z$ so \prec is transitive.
2. Suppose now that \preceq_i is a total order on S_i for each $i \in I$ and that I is well ordered by \leq . Let $x, y \in \prod_{i \in I} S_i$ with $x \neq y$. Let $J = \{i \in I : x_i \neq y_i\}$. Then $J \neq \emptyset$ by assumption, and hence has a minimum element j . If $i < j$ then $i \notin J$ and hence $x_i = y_i$. On the other hand, $x_j \neq y_j$ since $j \in J$ and therefore, since \preceq_j is totally ordered, we must have either $x_j \prec_j y_j$ or $y_j \prec_j x_j$. In the first case, $x \prec y$ and in the second case $y \prec x$. Hence \preceq is totally ordered.

The term *lexicographic* comes from the way that we order words alphabetically: We look at the first letter; if these are different, we know how to order the words. If the first letters are the same, we look at the second letter; if these are different, we know how to order the words. We continue in this way until we find letters that are different, and we can order the words. In fact, the lexicographic order is sometimes referred to as the *first difference* order. Note also that if S_i is a set and \preceq_i a total order on S_i for $i \in I$, then by the well ordering principle, there exists a well ordering \leq of I , and hence there exists a lexicographic total order on the product space $\prod_{i \in I} S_i$. As a mathematical structure, the lexicographic order is not as obscure as you might think.

(\mathbb{R}, \leq) is isomorphic to the lexicographic product of (\mathbb{Z}, \leq) with $([0, 1), \leq)$, where \leq is the ordinary order for real numbers.

Proof

Every $x \in \mathbb{R}$ can be uniquely expressed in the form $x = n + t$ where $n = \lfloor x \rfloor \in \mathbb{Z}$ is the *integer part* and $t = x - n \in [0, 1)$ is the *remainder*. Thus $x \mapsto (n, t)$ is a one-to-one function from \mathbb{R} onto $\mathbb{Z} \times [0, 1)$. For example, 5.3 maps to (5, 0.3), while -6.7 maps to (-7, 0.3). Suppose that $x = m + s$, $y = n + t \in \mathbb{R}$, where of course $m, n \in \mathbb{Z}$ are the integer parts of x and y , respectively, and $s, t \in [0, 1)$ are the corresponding remainders. Then $x < y$ if and only if $m < n$ or $m = n$ and $s < t$. Again, to illustrate with *real* real numbers, we can tell that $5.3 < 7.8$ just by comparing the integer parts: $5 < 7$. We can ignore the remainders. On the other hand, to see that $6.4 < 6.7$ we need to compare the remainders: $0.4 < 0.7$ since the integer parts are the same.

Limits of Sequences of Real Numbers

Suppose that (a_1, a_2, \dots) is a sequence of real numbers.

The sequence $\inf\{a_n, a_{n+1}, \dots\}$ is increasing in $n \in \mathbb{N}_+$.

Since the sequence of infimums in the last result is increasing, the limit exists in $\mathbb{R} \cup \{\infty\}$, and is called the *limit inferior* of the original sequence:

$$\liminf_{n \rightarrow \infty} a_n = \lim_{n \rightarrow \infty} \inf\{a_n, a_{n+1}, \dots\} \quad (1.4.3)$$

The sequence $\sup\{a_n, a_{n+1}, \dots\}$ is decreasing in $n \in \mathbb{N}_+$.

Since the the sequence of supremums in the last result is decreasing, the limit exists in $\mathbb{R} \cup \{-\infty\}$, and is called the *limit superior* of the original sequence:

$$\limsup_{n \rightarrow \infty} a_n = \lim_{n \rightarrow \infty} \sup\{a_n, a_{n+1}, \dots\} \quad (1.4.4)$$

Note that $\liminf_{n \rightarrow \infty} a_n \leq \limsup_{n \rightarrow \infty} a_n$ and equality holds if and only if $\lim_{n \rightarrow \infty} a_n$ exists (and is the common value).

Vector Spaces of Functions

Suppose that S is a nonempty set, and recall that the set \mathcal{V} of functions $f : S \rightarrow \mathbb{R}$ is a vector space, under the usual pointwise definition of addition and scalar multiplication. As noted in (9), \mathcal{V} is also a partial ordered set, under the pointwise partial order: $f \preceq g$ if and only if $f(x) \leq g(x)$ for all $x \in S$. Consistent with the definitions (19), $f \in \mathcal{V}$ is *bounded* if there exists $C \in (0, \infty)$ such that $|f(x)| \leq C$ for all $x \in S$. Now let \mathcal{U} denote the set of bounded functions $f : S \rightarrow \mathbb{R}$, and for $f \in \mathcal{U}$ define

$$\|f\| = \sup\{|f(x)| : x \in S\} \quad (1.4.5)$$

\mathcal{U} is a vector subspace of \mathcal{V} and $\|\cdot\|$ is a norm on \mathcal{U} .

Proof

To show that \mathcal{U} is a subspace, we just have to note that it is closed under addition and scalar multiplication. That is, if $f, g : S \rightarrow \mathbb{R}$ are bounded, and if $c \in \mathbb{R}$, then $f + g$ and cf are bounded. Next we show that $\|\cdot\|$ satisfies the axioms of a norm. Again, let $f, g \in \mathcal{U}$ and $c \in \mathbb{R}$

1. Clearly $\|f\| \geq 0$ and $\|f\| = 0$ if and only if $f(x) = 0$ for all $x \in S$ if and only if $f = \mathbf{0}$, the zero function on S .
2. $\|cf\| = \sup\{|cf(x)| : x \in S\} = |c| \sup\{|f(x)| : x \in S\} = |c| \|f\|$
3. By the usual triangle inequality on \mathbb{R} , $|f(x) + g(x)| \leq |f(x)| + |g(x)|$ for $x \in S$. Hence

$$\sup\{|f(x) + g(x)| : x \in S\} \leq \sup\{|f(x)| + |g(x)| : x \in S\} \leq \sup\{|f(x)| : x \in S\} + \sup\{|g(x)| : x \in S\} \quad (1.4.6)$$

That is, $\|f + g\| \leq \|f\| + \|g\|$.

Recall that part (a) is the *positive property*, part (b) is the *scaling property*, and part (c) is the *triangle inequality*.

Appropriately enough, $\|\cdot\|$ is called the *supremum norm* on \mathcal{U} . Vector spaces of bounded, real-valued functions, with the supremum norm are especially important in probability and random processes. We will return to this discussion again in the advanced sections on metric spaces and measure theory.

Computational Exercises


Let $S = \{2, 3, 4, 6, 12\}$.

1. Sketch the Hasse graph corresponding to the ordinary order \leq on S .
2. Sketch the Hasse graph corresponding to the division partial order $|$ on S .

Answer


1.

The Hasse graph of (S, \leq)

 Hasse graph

2.

The Hasse graph of $(S, |)$

 Hasse graph

Consider the ordinary order \leq on the set of real numbers \mathbb{R} , and let $A = [a, b)$ where $a < b$. Find each of the following that exist:

1. The set of minimal elements of A
2. The set of maximal elements of A
3. $\min(A)$
4. $\max(A)$
5. The set of lower bounds of A
6. The set of upper bounds of A
7. $\inf(A)$
8. $\sup(A)$

Answer

1. $\{a\}$
2. \emptyset
3. a
4. Does not exist
5. $(-\infty, a]$
6. $[b, \infty)$
7. a
8. b

Again consider the division partial order $|$ on the set of positive integers \mathbb{N}_+ and let $A = \{2, 3, 4, 6, 12\}$. Find each of the following that exist:

1. The set of minimal elements of A
2. The set of maximal elements of A
3. $\min(A)$
4. $\max(A)$
5. The set of lower bounds of A
6. The set of upper bounds of A
7. $\inf(A)$
8. $\sup(A)$.

Answer

1. $\{2, 3\}$
2. $\{12\}$
3. Does not exist
4. 12
5. $\{1\}$
6. $\{12, 24, 36, \dots\}$

7. 1
8. 12

Let $S = \{a, b, c\}$.

1. Give $\mathcal{P}(S)$ in list form.
2. Describe the Hasse graph of $(\mathcal{P}(S), \subseteq)$

Answer

1. $\mathcal{P}(S) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, S\}$
2. For $A \in \mathcal{P}(S)$ and $x \in S \setminus A$, there is a directed edge from A to $A \cup \{x\}$

Note that the Hasse graph of \supseteq looks the same as the graph of \subseteq , except for the labels on the vertices. This symmetry is because of the complement relationship.

Let $S = \{a, b, c, d\}$.

1. Give $\mathcal{P}(S)$ in list form.
2. Describe the Hasse graph of $(\mathcal{P}(S), \subseteq)$

Answer

1. $\mathcal{P}(S) = \{\emptyset, \{a\}, \{b\}, \{c\}, \{d\}, \{a, b\}, \{a, c\}, \{a, d\}, \{b, c\}, \{b, d\}, \{c, d\}, \{a, b, c\}, \{a, b, d\}, \{a, c, d\}, \{b, c, d\}, S\}$
2. For $A \in \mathcal{P}(S)$ and $x \in S \setminus A$, there is a directed edge from A to $A \cup \{x\}$

Note again that the Hasse graph of \supseteq looks the same as the graph of \subseteq , except for the labels on the vertices. This symmetry is because the complement relationship.

Suppose that A and B are subsets of a universal set S . Let \mathcal{A} denote the collection of the 16 subsets of S that can be constructed from A and B using the set operations. Show that (\mathcal{A}, \subseteq) is isomorphic to the partially ordered set in the previous exercise. Use the Venn diagram app to help.

Proof

Let $a = A \cap B$, $b = A \cap B^c$, $c = A^c \cap B$, $d = A^c \cap B^c$. Our basic assumption is that A and B are in “general position”, so that a, b, c, d are distinct and nonempty. Note also that $\{a, b, c, d\}$ partitions S . Now, map each subset \mathcal{S} of $\{a, b, c, d\}$ to $\bigcup \mathcal{S}$. This function is an isomorphism from \mathcal{S} to \mathcal{A} . That is, for \mathcal{S} and \mathcal{T} subsets of $\{a, b, c, d\}$, $\mathcal{S} \subseteq \mathcal{T}$ if and only if $\bigcup \mathcal{S} \subseteq \bigcup \mathcal{T}$.

This page titled [1.4: Partial Orders](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.5: Equivalence Relations

Basic Theory

Definitions

A relation \approx on a nonempty set S that is reflexive, symmetric, and transitive is an *equivalence relation* on S . Thus, for all $x, y, z \in S$,

1. $x \approx x$, the *reflexive property*.
2. If $x \approx y$ then $y \approx x$, the *symmetric property*.
3. If $x \approx y$ and $y \approx z$ then $x \approx z$, the *transitive property*.

As the name and notation suggest, an equivalence relation is intended to define a type of equivalence among the elements of S . Like [partial orders](#), equivalence relations occur naturally in most areas of mathematics, including probability.

Suppose that \approx is an equivalence relation on S . The *equivalence class* of an element $x \in S$ is the set of all elements that are equivalent to x , and is denoted

$$[x] = \{y \in S : y \approx x\} \quad (1.5.1)$$

Results

The most important result is that an equivalence relation on a set S defines a partition of S , by means of the equivalence classes.

Suppose that \approx is an equivalence relation on a set S .

1. If $x \approx y$ then $[x] = [y]$.
2. If $x \not\approx y$ then $[x] \cap [y] = \emptyset$.
3. The collection of (distinct) equivalence classes is a partition of S into nonempty sets.

Proof

1. Suppose that $x \approx y$. If $u \in [x]$ then $u \approx x$ and hence $u \approx y$ by the transitive property. Hence $u \in [y]$. Similarly, if $u \in [y]$ then $u \approx y$. But $y \approx x$ by the symmetric property, and hence $u \approx x$ by the transitive property. Hence $u \in [x]$.
2. Suppose that $x \not\approx y$. If $u \in [x] \cap [y]$, then $u \in [x]$ and $u \in [y]$, so $u \approx x$ and $u \approx y$. But then $x \approx u$ by the symmetric property, and then $x \approx y$ by the transitive property. This is a contradiction, so $[x] \cap [y] = \emptyset$.
3. From (a) and (b), the (distinct) equivalence classes are disjoint. If $x \in S$, then $x \approx x$ by the reflexive property, and hence $x \in [x]$. Therefore $\bigcup_{x \in S} [x] = S$.

Sometimes the set \mathcal{S} of equivalence classes is denoted S/\approx . The idea is that the equivalence classes are new “objects” obtained by “identifying” elements in S that are equivalent. Conversely, every partition of a set defines an equivalence relation on the set.

Suppose that \mathcal{S} is a collection of nonempty sets that partition a given set S . Define the relation \approx on S by $x \approx y$ if and only if $x \in A$ and $y \in A$ for some $A \in \mathcal{S}$.

1. \approx is an equivalence relation.
2. \mathcal{S} is the set of equivalence classes.

Proof

1. If $x \in S$, then $x \in A$ for some $A \in \mathcal{S}$, since \mathcal{S} partitions S . Hence $x \approx x$, and so the reflexive property holds. Next, \approx is trivially symmetric by definition. Finally, suppose that $x \approx y$ and $y \approx z$. Then $x, y \in A$ for some $A \in \mathcal{S}$ and $y, z \in B$ for some $B \in \mathcal{S}$. But then $y \in A \cap B$. The sets in \mathcal{S} are disjoint, so $A = B$. Hence $x, z \in A$, so $x \approx z$. Thus \approx is transitive.
2. If $x \in S$, then $x \in A$ for a unique $A \in \mathcal{S}$, and then by definition, $[x] = A$.




Figure 1.5.1: A partition of S . Any two points in the same partition set are equivalent.

Sometimes the equivalence relation \approx associated with a given partition \mathcal{S} is denoted S/\mathcal{S} . The idea, of course, is that elements in the same set of the partition are equivalent.

The process of forming a partition from an equivalence relation, and the process of forming an equivalence relation from a partition are inverses of each other.

1. If we start with an equivalence relation \approx on S , form the associated partition, and then construct the equivalence relation associated with the partition, then we end up with the original equivalence relation. In modular notation, $S/(S/\approx)$ is the same as \approx .
2. If we start with a partition \mathcal{S} of S , form the associated equivalence relation, and then form the partition associated with the equivalence relation, then we end up with the original partition. In modular notation, $S/(S/\mathcal{S})$ is the same as \mathcal{S} .

Suppose that S is a nonempty set. The most basic equivalence relation on S is the *equality relation* $=$. In this case $[x] = \{x\}$ for each $x \in S$. At the other extreme is the *trivial relation* \approx defined by $x \approx y$ for all $x, y \in S$. In this case S is the only equivalence class.

Every function f defines an equivalence relation on its domain, known as the *equivalence relation associated with f* . Moreover, the equivalence classes have a simple description in terms of the inverse images of f .

Suppose that $f : S \rightarrow T$. Define the relation \approx on S by $x \approx y$ if and only if $f(x) = f(y)$.

1. The relation \approx is an equivalence relation on S .
2. The set of equivalence classes is $\mathcal{S} = \{f^{-1}\{t\} : t \in \text{range}(f)\}$.
3. The function $F : \mathcal{S} \rightarrow T$ defined by $F([x]) = f(x)$ is well defined and is one-to-one.

Proof

1. If $x \in S$ then trivially $f(x) = f(x)$, so $x \approx x$. Hence \approx is reflexive. If $x \approx y$ then $f(x) = f(y)$ so trivially $f(y) = f(x)$ and hence $y \approx x$. Thus \approx is symmetric. If $x \approx y$ and $y \approx z$ then $f(x) = f(y)$ and $f(y) = f(z)$, so trivially $f(x) = f(z)$ and so $x \approx z$. Hence \approx is transitive.
2. Recall that $t \in \text{range}(f)$ if and only if $f(x) = t$ for some $x \in S$. Then by definition, $[x] = f^{-1}\{t\} = \{y \in S : f(y) = t\} = \{y \in S : f(y) = f(x)\}$
3. From (2), $[x] = [y]$ if and only if $x \approx y$ if and only if $f(x) = f(y)$. This shows both that F is well defined, and that F is one-to-one.




Figure 1.5.2: The equivalence relation on S associated with $f : S \rightarrow T$

Suppose again that $f : S \rightarrow T$.

1. If f is one-to-one then the equivalence relation associated with f is the equality relation, and hence $[x] = \{x\}$ for each $x \in S$.
2. If f is a constant function then the equivalence relation associated with f is the trivial relation, and hence S is the only equivalence class.

Proof

1. If f is one-to-one, then $x \approx y$ if and only if $f(x) = f(y)$ if and only if $x = y$.
2. If f is constant on S then $f(x) = f(y)$ and hence $x \approx y$ for all $x, y \in S$.

Equivalence relations associated with functions are universal: every equivalence relation is of this form:

Suppose that \approx is an equivalence relation on a set S . Define $f : S \rightarrow \mathcal{P}(S)$ by $f(x) = [x]$. Then \approx is the equivalence relation associated with f .

Proof

From (6), $x \approx y$ if and only if $[x] = [y]$ if and only if $f(x) = f(y)$.

The intersection of two equivalence relations is another equivalence relation.

Suppose that \approx and \cong are equivalence relations on a set S . Let \equiv denote the intersection of \approx and \cong (thought of as subsets of $S \times S$). Equivalently, $x \equiv y$ if and only if $x \approx y$ and $x \cong y$.

1. \equiv is an equivalence relation on S .

2. $[x]_{\equiv} = [x]_{\approx} \cap [x]_{\cong}$.

Suppose that we have a relation that is reflexive and transitive, but fails to be a partial order because it's not anti-symmetric. The relation and its inverse naturally lead to an equivalence relation, and then in turn, the original relation defines a true partial order on the equivalence classes. This is a common construction, and the details are given in the next theorem.

Suppose that \preceq is a relation on a set S that is reflexive and transitive. Define the relation \approx on S by $x \approx y$ if and only if $x \preceq y$ and $y \preceq x$.

1. \approx is an equivalence relation on S .

2. If A and B are equivalence classes and $x \preceq y$ for some $x \in A$ and $y \in B$, then $u \preceq v$ for all $u \in A$ and $v \in B$.

3. Define the relation \preceq on the collection of equivalence classes \mathcal{S} by $A \preceq B$ if and only if $x \preceq y$ for some (and hence all) $x \in A$ and $y \in B$. Then \preceq is a partial order on \mathcal{S} .

Proof

1. If $x \in S$ then $x \preceq x$ since \preceq is reflexive. Hence $x \approx x$, so \approx is reflexive. Clearly \approx is symmetric by the symmetry of the definition. Suppose that $x \approx y$ and $y \approx z$. Then $x \preceq y$, $y \preceq z$, $z \preceq y$ and $y \preceq x$. Hence $x \preceq z$ and $z \preceq x$ since \preceq is transitive. Therefore $x \approx z$ so \approx is transitive.

2. Suppose that A and B are equivalence classes of \approx and that $x \preceq y$ for some $x \in A$ and $y \in B$. If $u \in A$ and $v \in B$, then $x \approx u$ and $y \approx v$. Therefore $u \preceq x$ and $y \preceq v$. By transitivity, $u \preceq v$.

3. Suppose that $A \in \mathcal{S}$. If $x, y \in A$ then $x \approx y$ and hence $x \preceq y$. Therefore $A \preceq A$ and so \preceq is reflexive. Next suppose that $A, B \in \mathcal{S}$ and that $A \preceq B$ and $B \preceq A$. If $x \in A$ and $y \in B$ then $x \preceq y$ and $y \preceq x$. Hence $x \approx y$ so $A = B$. Therefore \preceq is antisymmetric. Finally, suppose that $A, B, C \in \mathcal{S}$ and that $A \preceq B$ and $B \preceq C$. Note that $B \neq \emptyset$ so let $y \in B$. If $x \in A$, $z \in C$ then $x \preceq y$ and $y \preceq z$. Hence $x \preceq z$ and therefore $A \preceq C$. So \preceq is transitive.

A prime example of the construction in the previous theorem occurs when we have a function whose range space is partially ordered. We can construct a partial order on the equivalence classes in the domain that are associated with the function.

Suppose that S and T are sets and that \preceq_T is a partial order on T . Suppose also that $f : S \rightarrow T$. Define the relation \preceq_S on S by $x \preceq_S y$ if and only if $f(x) \preceq_T f(y)$.

1. \preceq_S is reflexive and transitive.

2. The equivalence relation on S constructed in (10) is the equivalence relation associated with f , as in (6).

3. \preceq_S can be extended to a partial order on the equivalence classes corresponding to f .

Proof

1. If $x \in S$ then $f(x) \preceq_T f(x)$ since \preceq_T is reflexive, and hence $x \preceq_S x$. Thus \preceq_S is reflexive. Suppose that $x, y, z \in S$ and that $x \preceq_S y$ and $y \preceq_S z$. Then $f(x) \preceq_T f(y)$ and $f(y) \preceq_T f(z)$. Hence $f(x) \preceq_T f(z)$ since \preceq_T is transitive. Thus \preceq_S is transitive.

2. For the equivalence relation \approx on S constructed in (10), $x \approx y$ if and only if $x \preceq_S y$ and $y \preceq_S x$ if and only if $f(x) \preceq_T f(y)$ and $f(y) \preceq_T f(x)$ if and only if $f(x) = f(y)$, since \preceq_T is antisymmetric. Thus \approx is the equivalence relation associated with f .

3. This follows immediately from (10) and parts (a) and (b). If $u, v \in \text{range}(f)$, then $f^{-1}(\{u\}) \preceq_S f^{-1}(\{v\})$ if and only if $u \preceq_T v$.

Examples and Applications

Simple functions

Give the equivalence classes explicitly for the functions from \mathbb{R} into \mathbb{R} defined below:

1. $f(x) = x^2$.
2. $g(x) = \lfloor x \rfloor$.
3. $h(x) = \sin(x)$.

Answer

1. $[x] = \{x, -x\}$
2. $[x] = [\lfloor x \rfloor, \lfloor x \rfloor + 1)$
3. $[x] = \{x + 2n\pi : n \in \mathbb{Z}\} \cup \{(2n+1)\pi - x : n \in \mathbb{Z}\}$

Calculus

Suppose that I is a fixed interval of \mathbb{R} , and that S is the set of differentiable functions from I into \mathbb{R} . Consider the equivalence relation associated with the derivative operator D on S , so that $D(f) = f'$. For $f \in S$, give a simple description of $[f]$.

Answer

$$[f] = \{f + c : c \in \mathbb{R}\}$$

Congruence

Recall the *division relation* $|$ from \mathbb{N}_+ to \mathbb{Z} : For $d \in \mathbb{N}_+$ and $n \in \mathbb{Z}$, $d | n$ means that $n = kd$ for some $k \in \mathbb{Z}$. In words, d divides n or equivalently n is a multiple of d . In the previous section, we showed that $|$ is a partial order on \mathbb{N}_+ .

Fix $d \in \mathbb{N}_+$.

1. Define the relation \equiv_d on \mathbb{Z} by $m \equiv_d n$ if and only if $d | (n - m)$. The relation \equiv_d is known as *congruence modulo d* .
2. Let $r_d : \mathbb{Z} \rightarrow \{0, 1, \dots, d-1\}$ be defined so that $r_d(n)$ is the remainder when n is divided by d .

Recall that by the *Euclidean division theorem*, named for Euclid of course, $n \in \mathbb{Z}$ can be written uniquely in the form $n = kd + q$ where $k \in \mathbb{Z}$ and $q \in \{0, 1, \dots, d-1\}$, and then $r_d(n) = q$.

Congruence modulo d .

1. \equiv_d is the equivalence relation associated with the function r_d .
2. There are d distinct equivalence classes, given by $[q]_d = \{q + kd : k \in \mathbb{Z}\}$ for $q \in \{0, 1, \dots, d-1\}$.

Proof

1. Recall that for the equivalence relation associated with r_d , integers m and n are equivalent if and only if $r_d(m) = r_d(n)$. By the division theorem, $m = jd + p$ and $n = kd + q$, where $j, k \in \mathbb{Z}$ and $p, q \in \{0, 1, \dots, d-1\}$, and these representations are unique. Thus $n - m = (k - j)d + (q - p)$, and so $m \equiv_d n$ if and only if $d | (n - m)$ if and only if $p = q$ if and only if $r_d(m) = r_d(n)$.
2. Recall that the equivalence classes are $r_d^{-1}\{q\}$ for $q \in \text{range}(r_d) = \{0, 1, \dots, d-1\}$. By the division theorem, $r_d^{-1}\{q\} = \{kd + q : k \in \mathbb{Z}\}$.

Explicitly give the equivalence classes for \equiv_4 , congruence mod 4.

Answer

- $[0]_4 = \{0, 4, 8, 12, \dots\} \cup \{-4, -8, -12, -16, \dots\}$
- $[1]_4 = \{1, 5, 9, 13, \dots\} \cup \{-3, -7, -11, -15, \dots\}$
- $[2]_4 = \{2, 6, 10, 14, \dots\} \cup \{-2, -6, -10, -14, \dots\}$
- $[3]_4 = \{3, 7, 11, 15, \dots\} \cup \{-1, -5, -9, -13, \dots\}$

Linear Algebra

Linear algebra provides several examples of important and interesting equivalence relations. To set the stage, let $\mathbb{R}^{m \times n}$ denote the set of $m \times n$ matrices with real entries, for $m, n \in \mathbb{N}_+$.

Recall that the following are *row operations* on a matrix:

1. Multiply a row by a non-zero real number.
2. Interchange two rows.
3. Add a multiple of a row to another row.

Row operations are essential for inverting matrices and solving systems of linear equations.

Matrices $A, B \in \mathbb{R}^{m \times n}$ are *row equivalent* if A can be transformed into B by a finite sequence of row operations. Row equivalence is an equivalence relation on $\mathbb{R}^{m \times n}$.

Proof.

If $A \in \mathbb{R}^{m \times n}$, then A is row equivalent to itself: we can simply do nothing, or if you prefer, we can multiply the first row of A by 1. For symmetry, the key is that each row operation can be reversed by another row operation: multiplying a row by $c \neq 0$ is reversed by multiplying the same row of the resulting matrix by $1/c$. Interchanging two rows is reversed by interchanging the same two rows of the resulting matrix. Adding c times row i to row j is reversed by adding $-c$ times row i to row j in the resulting matrix. Thus, if we can transform A into B by a finite sequence of row operations, then we can transform B into A by applying the reversed row operations in the reverse order. Transitivity is clear: If we can transform A into B by a sequence of row operations and B into C by another sequence of row operations, then we can transform A into C by putting the two sequences together.

Our next relation involves similarity, which is very important in the study of linear transformations, change of basis, and the theory of eigenvalues and eigenvectors.

Matrices $A, B \in \mathbb{R}^{n \times n}$ are *similar* if there exists an invertible $P \in \mathbb{R}^{n \times n}$ such that $P^{-1}AP = B$. Similarity is an equivalence relation on $\mathbb{R}^{n \times n}$.

Proof

If $A \in \mathbb{R}^{n \times n}$ then $A = I^{-1}AI$, where I is the $n \times n$ identity matrix, so A is similar to itself. Suppose that $A, B \in \mathbb{R}^{n \times n}$ and that A is similar to B so that $B = P^{-1}AP$ for some invertible $P \in \mathbb{R}^{n \times n}$. Then $A = PBP^{-1} = (P^{-1})^{-1}BP^{-1}$ so B is similar to A . Finally, suppose that $A, B, C \in \mathbb{R}^{n \times n}$ and that A is similar to B and that B is similar to C . Then $B = P^{-1}AP$ and $C = Q^{-1}BQ$ for some invertible $P, Q \in \mathbb{R}^{n \times n}$. Then $C = Q^{-1}P^{-1}APQ = (PQ)^{-1}A(PQ)$, so A is similar to C .

Next recall that for $A \in \mathbb{R}^{m \times n}$, the *transpose* of A is the matrix $A^T \in \mathbb{R}^{n \times m}$ with the property that (i, j) entry of A is the (j, i) entry of A^T , for $i, j \in \{1, 2, \dots, m\}$. Simply stated, A^T is the matrix whose rows are the columns of A . For the theorem that follows, we need to remember that $(AB)^T = B^T A^T$ for $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times k}$, and $(A^T)^{-1} = (A^{-1})^T$ if $A \in \mathbb{R}^{n \times n}$ is invertible.

Matrices $A, B \in \mathbb{R}^{n \times n}$ are *congruent* if there exists an invertible $P \in \mathbb{R}^{n \times n}$ such that $B = P^T A P$. Congruence is an equivalence relation on $\mathbb{R}^{n \times n}$.

Proof

If $A \in \mathbb{R}^{n \times n}$ then $A = I^T A I$, where again I is the $n \times n$ identity matrix, so A is congruent to itself. Suppose that $A, B \in \mathbb{R}^{n \times n}$ and that A is congruent to B so that $B = P^T A P$ for some invertible $P \in \mathbb{R}^{n \times n}$. Then $A = (P^T)^{-1} B P^{-1} = (P^{-1})^T B P^{-1}$ so B is congruent to A . Finally, suppose that $A, B, C \in \mathbb{R}^{n \times n}$ and that A is congruent to B and that B is congruent to C . Then $B = P^T A P$ and $C = Q^T B Q$ for some invertible $P, Q \in \mathbb{R}^{n \times n}$. Then $C = Q^T P^T A P Q = (PQ)^T A (PQ)$, so A is congruent to C .

Congruence is important in the study of orthogonal matrices and change of basis. Of course, the term *congruence* applied to matrices should not be confused with the same term applied to integers.

Number Systems

Equivalence relations play an important role in the construction of complex mathematical structures from simpler ones. Often the objects in the new structure are equivalence classes of objects constructed from the simpler structures, modulo an equivalence relation that captures the essential properties of the new objects.

The construction of number systems is a prime example of this general idea. The next exercise explores the construction of rational numbers from integers.

Define a relation \approx on $\mathbb{Z} \times \mathbb{N}_+$ by $(j, k) \approx (m, n)$ if and only if $jn = km$.

1. \approx is an equivalence relation.
2. Define $\frac{m}{n} = [(m, n)]$, the equivalence class generated by (m, n) , for $m \in \mathbb{Z}$ and $n \in \mathbb{N}_+$. This definition captures the essential properties of the rational numbers.

Proof

1. For $(m, n) \in \mathbb{Z} \times \mathbb{N}_+$, $mn = nm$ of course, so $(m, n) \approx (m, n)$. Hence \approx is reflexive. If $(j, k), (m, n) \in \mathbb{Z} \times \mathbb{N}_+$ and $(j, k) \approx (m, n)$, then $jn = km$ so trivially $mk = nj$, and hence $(m, n) \approx (j, k)$. Thus \approx is symmetric. Finally, suppose that $(j, k), (m, n), (p, q) \in \mathbb{Z} \times \mathbb{N}_+$ and that $(j, k) \approx (m, n)$ and $(m, n) \approx (p, q)$. Then $jn = km$ and $mq = np$, so $jnp = kmp$ which implies $jmq = kmp$, and so $jq = kp$. Hence $(j, k) \approx (p, q)$ so \approx is transitive.
2. Suppose that $\frac{j}{k}$ and $\frac{m}{n}$ are rational numbers in the usual, informal sense, where $j, m \in \mathbb{Z}$ and $k, n \in \mathbb{N}_+$. Then $\frac{j}{k} = \frac{m}{n}$ if and only if $jn = km$ if and only if $(j, k) \approx (m, n)$, so it makes sense to define $\frac{m}{n}$ as the equivalence class generated by (m, n) . Addition and multiplication are defined in the usual way: if $(j, k), (m, n) \in \mathbb{Z} \times \mathbb{N}_+$ then

$$\frac{j}{k} + \frac{m}{n} = \frac{jn + mk}{kn}, \quad \frac{j}{k} \cdot \frac{m}{n} = \frac{jm}{kn} \quad (1.5.2)$$

The definitions are consistent; that is they do not depend on the particular representations of the equivalence classes.

This page titled [1.5: Equivalence Relations](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.6: Cardinality

Basic Theory

Definitions

Suppose that \mathcal{S} is a non-empty collection of sets. We define a relation \approx on \mathcal{S} by $A \approx B$ if and only if there exists a one-to-one function f from A onto B . The relation \approx is an equivalence relation on \mathcal{S} . That is, for all $A, B, C \in \mathcal{S}$,

1. $A \approx A$, the reflexive property
2. If $A \approx B$ then $B \approx A$, the symmetric property
3. If $A \approx B$ and $B \approx C$ then $A \approx C$, the transitive property

Proof

1. The identity function I_A on A , given by $I_A(x) = x$ for $x \in A$, maps A one-to-one onto A . Hence $A \approx A$.
2. If $A \approx B$ then there exists a one-to-one function f from A onto B . But then f^{-1} is a one-to-one function from B onto A , so $B \approx A$.
3. Suppose that $A \approx B$ and $B \approx C$. Then there exists a one-to-one function f from A onto B and a one-to-one function g from B onto C . But then $g \circ f$ is a one-to-one function from A onto C , so $A \approx C$.

A one-to-one function f from A onto B is sometimes called a *bijection*. Thus if $A \approx B$ then A and B are in *one-to-one* correspondence and are said to have the same *cardinality*. The equivalence classes under this equivalence relation capture the notion of having the same number of elements.

Let $\mathbb{N}_0 = \emptyset$, and for $k \in \mathbb{N}_+$, let $\mathbb{N}_k = \{0, 1, \dots, k-1\}$. As always, $\mathbb{N} = \{0, 1, 2, \dots\}$ is the set of all natural numbers.

Suppose that A is a set.

1. A is *finite* if $A \approx \mathbb{N}_k$ for some $k \in \mathbb{N}$, in which case k is the cardinality of A , and we write $\#(A) = k$.
2. A is *infinite* if A is not finite.
3. A is *countably infinite* if $A \approx \mathbb{N}$.
4. A is *countable* if A is finite or countably infinite.
5. A is *uncountable* if A is not countable.

In part (a), think of \mathbb{N}_k as a *reference set* with k elements; any other set with k elements must be equivalent to this one. We will study the cardinality of finite sets in the next two sections on Counting Measure and Combinatorial Structures. In this section, we will concentrate primarily on infinite sets. In part (d), a countable set is one that can be *enumerated* or *counted* by putting the elements into one-to-one correspondence with \mathbb{N}_k for some $k \in \mathbb{N}$ or with all of \mathbb{N} . An uncountable set is one that cannot be so counted. Countable sets play a special role in probability theory, as in many other branches of mathematics. Apriori, it's not clear that there *are* uncountable sets, but we will soon see examples.

Preliminary Examples

If S is a set, recall that $\mathcal{P}(S)$ denotes the power set of S (the set of all subsets of S). If A and B are sets, then A^B is the set of all functions from B into A . In particular, $\{0, 1\}^S$ denotes the set of functions from S into $\{0, 1\}$.

If S is a set then $\mathcal{P}(S) \approx \{0, 1\}^S$.

Proof

The mapping that takes a set $A \in \mathcal{P}(S)$ into its indicator function $\mathbf{1}_A \in \{0, 1\}^S$ is one-to-one and onto. Specifically, if $A, B \in \mathcal{P}(S)$ and $\mathbf{1}_A = \mathbf{1}_B$, then $A = B$, so the mapping is one-to-one. On the other hand, if $f \in \{0, 1\}^S$ then $f = \mathbf{1}_A$ where $A = \{x \in S : f(x) = 1\}$. Hence the mapping is onto.

Next are some examples of countably infinite sets.

The following sets are countably infinite:

1. The set of even natural numbers $E = \{0, 2, 4, \dots\}$
2. The set of integers \mathbb{Z}

Proof

1. The function $f: \mathbb{N} \rightarrow E$ given by $f(n) = 2n$ is one-to-one and onto.
2. The function $g: \mathbb{N} \rightarrow \mathbb{Z}$ given by $g(n) = \frac{n}{2}$ if n is even and $g(n) = -\frac{n+1}{2}$ if n is odd, is one-to-one and onto.

At one level, it might seem that E has only half as many elements as \mathbb{N} while \mathbb{Z} has twice as many elements as \mathbb{N} . as the previous result shows, that point of view is incorrect: \mathbb{N} , E , and \mathbb{Z} all have the same cardinality (and are countably infinite). The next example shows that there are indeed uncountable sets.

If A is a set with at least two elements then $S = A^{\mathbb{N}}$, the set of all functions from \mathbb{N} into A , is uncountable.

Proof

The proof is by contradiction, and uses a nice trick known as the *diagonalization method*. Suppose that S is countably infinite (it's clearly not finite), so that the elements of S can be enumerated: $S = \{f_0, f_1, f_2, \dots\}$. Let a and b denote distinct elements of A and define $g: \mathbb{N} \rightarrow A$ by $g(n) = b$ if $f_n(n) = a$ and $g(n) = a$ if $f_n(n) \neq a$. Note that $g \neq f_n$ for each $n \in \mathbb{N}$, so $g \notin S$. This contradicts the fact that S is the set of *all* functions from \mathbb{N} into A .

Subsets of Infinite Sets

Surely a set must be at least as large as any of its subsets, in terms of cardinality. On the other hand, by example (4), the set of natural numbers \mathbb{N} , the set of even natural numbers E and the set of integers \mathbb{Z} all have exactly the same cardinality, even though $E \subset \mathbb{N} \subset \mathbb{Z}$. In this subsection, we will explore some interesting and somewhat paradoxical results that relate to subsets of infinite sets. Along the way, we will see that the countable infinity is the “smallest” of the infinities.

If S is an infinite set then S has a countable infinite subset.

Proof

Select $a_0 \in S$. It's possible to do this since S is infinite and therefore nonempty. Inductively, having chosen $\{a_0, a_1, \dots, a_{k-1}\} \subseteq S$, select $a_k \in S \setminus \{a_0, a_1, \dots, a_{k-1}\}$. Again, it's possible to do this since S is not finite. Manifestly, $\{a_0, a_1, \dots\}$ is a countably infinite subset of S .

A set S is infinite if and only if S is equivalent to a proper subset of S .

Proof

If S is finite, then S is not equivalent to a proper subset by the “pigeonhole principle”. If S is infinite, then S has countably infinite subset $\{a_0, a_1, a_2, \dots\}$ by the previous result. Define the function $f: S \rightarrow S$ by $f(a_n) = a_{2n}$ for $n \in \mathbb{N}$ and $f(x) = x$ for $x \in S \setminus \{a_0, a_1, a_2, \dots\}$. Then f maps S one-to-one onto $S \setminus \{a_1, a_3, a_5, \dots\}$.

When S was infinite in the proof of the previous result, not only did we map S one-to-one onto a proper subset, we actually threw away a countably infinite subset and still maintained equivalence. Similarly, we can *add* a countably infinite set to an infinite set S without changing the cardinality.

If S is an infinite set and B is a countable set, then $S \approx S \cup B$.

Proof

Consider the most extreme case where B is countably infinite and disjoint from S . Then S has a countably infinite subset $A = \{a_0, a_1, a_2, \dots\}$ by the result above, and B can be enumerated, so $B = \{b_0, b_1, b_2, \dots\}$. Define the function $f: S \rightarrow S \cup B$ by $f(a_n) = a_{n/2}$ if n is even, $f(a_n) = b_{(n-1)/2}$ if n is odd, and $f(x) = x$ if $x \in S \setminus \{a_0, a_1, a_2, \dots\}$. Then f maps S one-to-one onto $S \cup B$.

In particular, if S is uncountable and B is countable then $S \cup B$ and $S \setminus B$ have the same cardinality as S , and in particular are uncountable. In terms of the dichotomies *finite-infinite* and *countable-uncountable*, a set is indeed at least as large as a subset. First we need a preliminary result.

If S is countably infinite and $A \subseteq S$ then A is countable.

Proof

It suffices to show that if A is an infinite subset of S then A is countably infinite. Since S is countably infinite, it can be enumerated: $S = \{x_0, x_1, x_2, \dots\}$. Let n_i be the i th smallest index such that $x_{n_i} \in A$. Then $A = \{x_{n_0}, x_{n_1}, x_{n_2}, \dots\}$ and hence is countably infinite.

Suppose that $A \subseteq B$.

1. If B is finite then A is finite.
2. If A is infinite then B is infinite.
3. If B is countable then A is countable.
4. If A is uncountable then B is uncountable.

Proof

1. This is clear from the definition of a finite set.
2. This is the contrapositive of (a).
3. If A is finite, then A is countable. If A is infinite, then B is infinite by (b) and hence is countably infinite. But then A is countably infinite by (9).
4. This is the contrapositive of (c).

Comparisons by one-to-one and onto functions

We will look deeper at the general question of when one set is “at least as big” as another, in the sense of cardinality. Not surprisingly, this will eventually lead to a partial order on the cardinality equivalence classes.

First note that if there exists a function that maps a set A one-to-one into a set B , then in a sense, there is a copy of A contained in B . Hence B should be at least as large as A .

Suppose that $f : A \rightarrow B$ is one-to-one.

1. If B is finite then A is finite.
2. If A is infinite then B is infinite.
3. If B is countable then A is countable.
4. If A is uncountable then B is uncountable.

Proof

Note that f maps A one-to-one onto $f(A)$. Hence $A \approx f(A)$ and $f(A) \subseteq B$. The results now follow from (10):

1. If B is finite then $f(A)$ is finite and hence A is finite.
2. If A is infinite then $f(A)$ is infinite and hence B is infinite.
3. If B is countable then $f(A)$ is countable and hence A is countable.
4. If A is uncountable then $f(A)$ is uncountable and hence B is uncountable.

On the other hand, if there exists a function that maps a set A onto a set B , then in a sense, there is a copy of B contained in A . Hence A should be at least as large as B .

Suppose that $f : A \rightarrow B$ is onto.

1. If A is finite then B is finite.
2. If B is infinite then A is infinite.
3. If A is countable then B is countable.
4. If B is uncountable then A is uncountable.

Proof

For each $y \in B$, select a specific $x \in A$ with $f(x) = y$ (if you are persnickety, you may need to invoke the axiom of choice). Let C be the set of chosen points. Then f maps C one-to-one onto B , so $C \approx B$ and $C \subseteq A$. The results now follow from (11):

1. If A is finite then C is finite and hence B is finite.
2. If B is infinite then C is infinite and hence A is infinite.
3. If A is countable then C is countable and hence B is countable.
4. If B is uncountable then C is uncountable and hence A is uncountable.

The previous exercise also could be proved from the one before, since if there exists a function f mapping A onto B , then there exists a function g mapping B one-to-one into A . This duality is proven in the discussion of the axiom of choice. A simple and useful corollary of the previous two theorems is that if B is a given countably infinite set, then a set A is countable if and only if there exists a one-to-one function f from A into B , if and only if there exists a function g from B onto A .

If A_i is a countable set for each i in a countable index set I , then $\bigcup_{i \in I} A_i$ is countable.

Proof

Consider the most extreme case in which the index set I is countably infinite. Since A_i is countable, there exists a function f_i that maps \mathbb{N} onto A_i for each $i \in \mathbb{N}$. Let $M = \{2^i 3^j : (i, j) \in \mathbb{N} \times \mathbb{N}\}$. Note that the points in M are distinct, that is, $2^i 3^j \neq 2^m 3^n$ if $(i, j), (m, n) \in \mathbb{N} \times \mathbb{N}$ and $(i, j) \neq (m, n)$. Hence M is infinite, and since $M \subset \mathbb{N}$, M is countably infinite. The function f given by $f(2^i 3^j) = f_i(j)$ for $(i, j) \in \mathbb{N} \times \mathbb{N}$ maps M onto $\bigcup_{i \in I} A_i$, and hence this last set is countable.

If A and B are countable then $A \times B$ is countable.

Proof

There exists a function f that maps \mathbb{N} onto A , and there exists a function g that maps \mathbb{N} onto B . Again, let $M = \{2^i 3^j : (i, j) \in \mathbb{N} \times \mathbb{N}\}$ and recall that M is countably infinite. Define $h : M \rightarrow A \times B$ by $h(2^i 3^j) = (f(i), g(j))$. Then h maps M onto $A \times B$ and hence this last set is countable.

The last result could also be proven from the one before, by noting that

$$A \times B = \bigcup_{a \in A} \{a\} \times B \quad (1.6.1)$$

Both proofs work because the set M is essentially a copy of $\mathbb{N} \times \mathbb{N}$, embedded inside of \mathbb{N} . The last theorem generalizes to the statement that a finite product of countable sets is still countable. But, from (5), a product of *infinitely* many sets (with at least 2 elements each) will be uncountable.

The set of rational numbers \mathbb{Q} is countably infinite.

Proof

The sets \mathbb{Z} and \mathbb{N}_+ are countably infinite and hence the set $\mathbb{Z} \times \mathbb{N}_+$ is countably infinite. The function $f : \mathbb{Z} \times \mathbb{N}_+ \rightarrow \mathbb{Q}$ given by $f(m, n) = \frac{m}{n}$ is onto.

A real number is *algebraic* if it is the root of a polynomial function (of degree 1 or more) with integer coefficients. Rational numbers are algebraic, as are rational roots of rational numbers (when defined). Moreover, the algebraic numbers are closed under addition, multiplication, and division. A real number is *transcendental* if it's not algebraic. The numbers e and π are transcendental, but we don't know very many other transcendental numbers by name. However, as we will see, most (in the sense of cardinality) real numbers are transcendental.

The set of algebraic numbers \mathbb{A} is countably infinite.

Proof

Let $\mathbb{Z}_0 = \mathbb{Z} \setminus \{0\}$ and let $\mathbb{Z}_n = \mathbb{Z}^{n-1} \times \mathbb{Z}_0$ for $n \in \mathbb{N}_+$. The set \mathbb{Z}_n is countably infinite for each n . Let $C = \bigcup_{n=1}^{\infty} \mathbb{Z}_n$. Think of C as the set of coefficients and note that C is countably infinite. Let P denote the set of polynomials of degree 1 or more, with integer coefficients. The function $(a_0, a_1, \dots, a_n) \mapsto a_0 + a_1 x + \dots + a_n x^n$ maps C onto P , and hence P is countable. For $p \in P$, let A_p denote the set of roots of p . A polynomial of degree n in P has at most n roots, by the *fundamental theorem of algebra*, so in particular A_p is finite for each $p \in P$. Finally, note that $\mathbb{A} = \bigcup_{p \in P} A_p$ and so \mathbb{A} is countable. Of course $\mathbb{N} \subset \mathbb{A}$, so \mathbb{A} is countably infinite.

Now let's look at some uncountable sets.

The interval $[0, 1)$ is uncountable.

Proof

Recall that $\{0, 1\}^{\mathbb{N}_+}$ is the set of all functions from \mathbb{N}_+ into $\{0, 1\}$, which in this case, can be thought of as infinite sequences or *bit strings*:

$$\{0, 1\}^{\mathbb{N}_+} = \{\mathbf{x} = (x_1, x_2, \dots) : x_n \in \{0, 1\} \text{ for all } n \in \mathbb{N}_+\} \quad (1.6.2)$$

By (5), this set is uncountable. Let $N = \{\mathbf{x} \in \{0, 1\}^{\mathbb{N}_+} : x_n = 1 \text{ for all but finitely many } n\}$, the set of bit strings that eventually terminate in all 1s. Note that $N = \bigcup_{n=1}^{\infty} N_n$ where $N_n = \{\mathbf{x} \in \{0, 1\}^{\mathbb{N}_+} : x_k = 1 \text{ for all } k \geq n\}$. Clearly N_n is finite for all $n \in \mathbb{N}_+$, so N is countable, and therefore $S = \{0, 1\}^{\mathbb{N}_+} \setminus N$ is uncountable. In fact, $S \approx \{0, 1\}^{\mathbb{N}_+}$. The function

$$\mathbf{x} \mapsto \sum_{n=1}^{\infty} \frac{x_n}{2^n} \quad (1.6.3)$$

maps S one-to-one onto $[0, 1)$. In words every number in $[0, 1)$ has a unique *binary expansion* in the form of a sequence in S . Hence $[0, 1) \approx S$ and in particular, is uncountable. The reason for eliminating the bit strings that terminate in 1s is to ensure uniqueness, so that the mapping is one-to-one. The bit string $x_1 x_2 \cdots x_k 0111 \cdots$ corresponds to the same number in $[0, 1)$ as the bit string $x_1 x_2 \cdots x_k 1000 \cdots$.

The following sets have the same cardinality, and in particular all are uncountable:

1. \mathbb{R} , the set of real numbers.
2. Any interval I of \mathbb{R} , as long as the interval is not empty or a single point.
3. $\mathbb{R} \setminus \mathbb{Q}$, the set of irrational numbers.
4. $\mathbb{R} \setminus \mathbb{A}$, the set of transcendental numbers.
5. $\mathcal{P}(\mathbb{N})$, the power set of \mathbb{N} .

Proof

1. The mapping $x \mapsto \frac{2x-1}{x(1-x)}$ maps $(0, 1)$ one-to-one onto \mathbb{R} so $(0, 1) \approx \mathbb{R}$. But $(0, 1) = [0, 1) \setminus \{0\}$, so $(0, 1) \approx (0, 1) \approx \mathbb{R}$, and all of these sets are uncountable by the previous result.
2. Suppose $a, b \in \mathbb{R}$ and $a < b$. The mapping $x \mapsto a + (b-a)x$ maps $(0, 1)$ one-to-one onto (a, b) and hence $(a, b) \approx (0, 1) \approx \mathbb{R}$. Also, $[a, b) = (a, b) \cup \{a\}$, $(a, b] = (a, b) \cup \{b\}$, and $[a, b] = (a, b) \cup \{a, b\}$, so $(a, b) \approx [a, b) \approx (a, b] \approx [a, b] \approx \mathbb{R}$. The function $x \mapsto e^x$ maps \mathbb{R} one-to-one onto $(0, \infty)$, so $(0, \infty) \approx \mathbb{R}$. For $a \in \mathbb{R}$, the function $x \mapsto a + x$ maps $(0, \infty)$ one-to-one onto (a, ∞) and the mapping $x \mapsto a - x$ maps $(0, \infty)$ one to one onto $(-\infty, a)$ so $(a, \infty) \approx (-\infty, a) \approx (0, \infty) \approx \mathbb{R}$. Next, $[a, \infty) = (a, \infty) \cup \{a\}$ and $(-\infty, a] = (-\infty, a) \cup \{a\}$, so $[a, \infty) \approx (-\infty, a] \approx \mathbb{R}$.
3. \mathbb{Q} is countably infinite, so $\mathbb{R} \setminus \mathbb{Q} \approx \mathbb{R}$.
4. Similarly, \mathbb{A} is countably infinite, so $\mathbb{R} \setminus \mathbb{A} \approx \mathbb{R}$.
5. If S is countably infinite, then by the previous result and (a), $\mathcal{P}(S) \approx \mathcal{P}(\mathbb{N}_+) \approx \{0, 1\}^{\mathbb{N}_+} \approx [0, 1)$.

The Cardinality Partial Order

Suppose that \mathcal{S} is a nonempty collection of sets. We define the relation \preceq on \mathcal{S} by $A \preceq B$ if and only if there exists a one-to-one function f from A into B , if and only if there exists a function g from B onto A . In light of the previous subsection, $A \preceq B$ should capture the notion that B is at least as big as A , in the sense of cardinality.

The relation \preceq is reflexive and transitive.

Proof

For $A \in \mathcal{S}$, the identity function $I_A : A \rightarrow A$ given by $I_A(x) = x$ is one-to-one (and also onto), so $A \preceq A$. Suppose that $A, B, C \in \mathcal{S}$ and that $A \preceq B$ and $B \preceq C$. Then there exist one-to-one functions $f : A \rightarrow B$ and $g : B \rightarrow C$. But then $g \circ f : A \rightarrow C$ is one-to-one, so $A \preceq C$.

Thus, we can use the construction in the section on [Equivalence Relations](#) to first define an equivalence relation on \mathcal{S} , and then extend \preceq to a true partial order on the collection of equivalence classes. The only question that remains is whether the equivalence relation we obtain in this way is the same as the one that we have been using in our study of cardinality. Rephrased, the question is this: *If there exists a one-to-one function from A into B and a one-to-one function from B into A , does there necessarily exist a one-to-one function from A onto B ?* Fortunately, the answer is yes; the result is known as the *Schröder-Bernstein Theorem*, named for Ernst Schröder and Felix Bernstein.

If $A \preceq B$ and $B \preceq A$ then $A \approx B$.

Proof


Set inclusion \subseteq is a partial order on $\mathcal{P}(A)$ (the power set of A) with the property that every subcollection of $\mathcal{P}(A)$ has a supremum (namely the union of the subcollection). Suppose that f maps A one-to-one into B and g maps B one-to-one into A . Define the function $h : \mathcal{P}(A) \rightarrow \mathcal{P}(A)$ by $h(U) = A \setminus g[B \setminus f(U)]$ for $U \subseteq A$. Then h is increasing:

$$U \subseteq V \implies f(U) \subseteq f(V) \implies B \setminus f(V) \subseteq B \setminus f(U) \quad (1.6.4)$$

$$\implies g[B \setminus f(V)] \subseteq g[B \setminus f(U)] \implies A \setminus g[B \setminus f(U)] \subseteq A \setminus g[B \setminus f(V)] \quad (1.6.5)$$

From the fixed point theorem for partially ordered sets, there exists $U \subseteq A$ such that $h(U) = U$. Hence $U = A \setminus g[B \setminus f(U)]$ and therefore $A \setminus U = g[B \setminus f(U)]$. Now define $F : A \rightarrow B$ by $F(x) = f(x)$ if $x \in U$ and $F(x) = g^{-1}(x)$ if $x \in A \setminus U$.

f maps U one-to-one onto $f(U)$; g maps $B \setminus f(U)$ one-to-one onto $A \setminus U$

 Schröder-Bernstein Theorem

Next we show that F is one-to-one. Suppose that $x_1, x_2 \in A$ and $F(x_1) = F(x_2)$. If $x_1, x_2 \in U$ then $f(x_1) = f(x_2)$ so $x_1 = x_2$ since f is one-to-one. If $x_1, x_2 \in A \setminus U$ then $g^{-1}(x_1) = g^{-1}(x_2)$ so $x_1 = x_2$ since g^{-1} is one-to-one. If $x_1 \in U$ and $x_2 \in A \setminus U$. Then $F(x_1) = f(x_1) \in f(U)$ while $F(x_2) = g^{-1}(x_2) \in B \setminus f(U)$, so $F(x_1) = F(x_2)$ is impossible.

Finally we show that F is onto. Let $y \in B$. If $y \in f(U)$ then $y = f(x)$ for some $x \in U$ so $F(x) = y$. If $y \in B \setminus f(U)$ then $x = g(y) \in A \setminus U$ so $F(x) = g^{-1}(x) = y$.

We will write $A \prec B$ if $A \preceq B$, but $A \not\approx B$. That is, there exists a one-to-one function from A into B , but there does not exist a function from A onto B . Note that \prec would have its usual meaning if applied to the equivalence classes. That is, $[A] \prec [B]$ if and only if $[A] \preceq [B]$ but $[A] \neq [B]$. Intuitively, of course, $A \prec B$ means that B is strictly larger than A , in the sense of cardinality.

$A \prec B$ in each of the following cases:

1. A and B are finite and $\#(A) < \#(B)$.
2. A is finite and B is countably infinite.
3. A is countably infinite and B is uncountable.

We close our discussion with the observation that for any set, there is always a larger set.

If S is a set then $S \prec \mathcal{P}(S)$.

Proof

First, it's trivial to map S one-to-one into $\mathcal{P}(S)$; just map x to $\{x\}$. Suppose now that f maps S onto $\mathcal{P}(S)$ and let $R = \{x \in S : x \notin f(x)\}$. Since f is onto, there exists $t \in S$ such that $f(t) = R$. Note that $t \in f(t)$ if and only if $t \notin f(t)$.

The proof that a set cannot be mapped onto its power set is similar to the *Russell paradox*, named for Bertrand Russell.

The *continuum hypothesis* is the statement that there is no set whose cardinality is strictly between that of \mathbb{N} and \mathbb{R} . The continuum hypothesis actually started out as the *continuum conjecture*, until it was shown to be consistent with the usual axioms of the real number system (by Kurt Gödel in 1940), and independent of those axioms (by Paul Cohen in 1963).

Assuming the continuum hypothesis, if S is uncountable then there exists $A \subseteq S$ such that A and A^c are uncountable.

Proof

Under the continuum hypothesis, if S is uncountable then $[0, 1] \preceq S$. Hence there exists a one-to-one function $f : [0, 1] \rightarrow S$. Let $A = f\left[0, \frac{1}{2}\right)$. Then A is uncountable, and since $f\left[\frac{1}{2}, 1\right) \subseteq A^c$, A^c is uncountable.

There is a more complicated proof of the last result, without the continuum hypothesis and just using the axiom of choice.

This page titled [1.6: Cardinality](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.7: Counting Measure

Basic Theory

For our first discussion, we assume that the universal set S is finite. Recall the following definition from the section on [cardinality](#).

For $A \subseteq S$, the *cardinality* of A is the number of elements in A , and is denoted $\#(A)$. The function $\#$ on $\mathcal{P}(S)$ is called *counting measure*.

Counting measure plays a fundamental role in discrete probability structures, and particularly those that involve sampling from a finite set. The set S is typically very large, hence efficient counting methods are essential. The first combinatorial problem is attributed to the Greek mathematician Xenocrates.

In many cases, a set of objects can be counted by establishing a one-to-one correspondence between the given set and some other set. Naturally, the two sets have the same number of elements, but for various reasons, the second set may be easier to count.

The Addition Rule

The *addition rule* of combinatorics is simply the additivity axiom of counting measure.

If $\{A_1, A_2, \dots, A_n\}$ is a collection of disjoint subsets of S then

$$\# \left(\bigcup_{i=1}^n A_i \right) = \sum_{i=1}^n \#(A_i) \quad (1.7.1)$$




Figure 1.7.1: The addition rule

The following counting rules are simple consequences of the [addition rule](#). Be sure to try the proofs yourself before reading the ones in the text.

$\#(A^c) = \#(S) - \#(A)$. This is the *complement rule*.

Proof




Figure 1.7.2: The complement rule

$\#(B \setminus A) = \#(B) - \#(A \cap B)$. This is the *difference rule*.

Proof

Note that $A \cap B$ and $B \setminus A$ are disjoint and their union is B . Hence $\#(A \cap B) + \#(B \setminus A) = \#(B)$.

If $A \subseteq B$ then $\#(B \setminus A) = \#(B) - \#(A)$. This is the *proper difference rule*.

Proof

This follows from the difference rule, since $A \cap B = A$.

If $A \subseteq B$ then $\#(A) \leq \#(B)$.

Proof

This follows from the proper difference rule: $\#(B) = \#(A) + \#(B \setminus A) \geq \#(A)$.

Thus, $\#$ is an increasing function, relative to the subset partial order \subseteq on $\mathcal{P}(S)$, and the ordinary order \leq on \mathbb{N} .

Inequalities

Our next discussion concerns two inequalities that are useful for obtaining bounds on the number of elements in a set. The first is *Boole's inequality* (named after George Boole) which gives an upper bound on the cardinality of a union.

If $\{A_1, A_2, \dots, A_n\}$ is a finite collection of subsets of S then

$$\# \left(\bigcup_{i=1}^n A_i \right) \leq \sum_{i=1}^n \#(A_i) \quad (1.7.2)$$

Proof

Let $B_1 = A_1$ and $B_i = A_i \setminus (A_1 \cup \dots \cup A_{i-1})$ for $i \in \{2, 3, \dots, n\}$. Note that $\{B_1, B_2, \dots, B_n\}$ is a pairwise disjoint collection and has the same union as $\{A_1, A_2, \dots, A_n\}$. From the [increasing property](#), $\#(B_i) \leq \#(A_i)$ for each $i \in \{1, 2, \dots, n\}$. Hence by the [addition rule](#),

$$\# \left(\bigcup_{i=1}^n A_i \right) = \# \left(\bigcup_{i=1}^n B_i \right) \leq \sum_{i=1}^n \#(A_i) \quad (1.7.3)$$

Intuitively, Boole's inequality holds because parts of the union have been counted more than once in the expression on the right. The second inequality is *Bonferroni's inequality* (named after Carlo Bonferroni), which gives a lower bound on the cardinality of an intersection.

If $\{A_1, A_2, \dots, A_n\}$ is a finite collection of subsets of S then

$$\# \left(\bigcap_{i=1}^n A_i \right) \geq \#(S) - \sum_{i=1}^n [\#(S) - \#(A_i)] \quad (1.7.4)$$

Proof

Using the [complement rule](#), [Boole's inequality](#), and DeMorgan's law,

$$\# \left(\bigcap_{i=1}^n A_i \right) = \#(S) - \# \left(\bigcup_{i=1}^n A_i^c \right) \geq \#(S) - \sum_{i=1}^n \#(A_i^c) = \#(S) - \sum_{i=1}^n [\#(S) - \#(A_i)] \quad (1.7.5)$$

The Inclusion-Exclusion Formula

The inclusion-exclusion formula gives the cardinality of a union of sets in terms of the cardinality of the various intersections of the sets. The formula is useful because intersections are often easier to count. We start with the special cases of two sets and three sets. As usual, we assume that the sets are subsets of a finite universal set S .

If A and B are subsets of S then $\#(A \cup B) = \#(A) + \#(B) - \#(A \cap B)$.

Proof




Figure 1.7.3: The inclusion-exclusion theorem for two sets

If A, B, C are subsets of S then $\#(A \cup B \cup C) = \#(A) + \#(B) + \#(C) - \#(A \cap B) - \#(A \cap C) - \#(B \cap C) + \#(A \cap B \cap C)$.




Figure 1.7.4: The inclusion-exclusion theorem for three sets

The inclusion-exclusion rule for two and three sets can be generalized to a union of n sets; the generalization is known as the (general) *inclusion-exclusion formula*.

Suppose that $\{A_i : i \in I\}$ is a collection of subsets of S where I is an index set with $\#(I) = n$. Then

$$\# \left(\bigcup_{i \in I} A_i \right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \# \left(\bigcap_{j \in J} A_j \right) \quad (1.7.6)$$

Proof

The proof is by induction on n . The formula holds for $n = 2$ sets by the result for [two sets](#). Suppose the formula holds for $n \in \{2, 3, \dots\}$, and suppose that $\{A_1, A_2, \dots, A_n, A_{n+1}\}$ is a collection of $n + 1$ subsets of S . Then

$$\bigcup_{i=1}^{n+1} A_i = \left(\bigcup_{i=1}^n A_i \right) \cup \left[A_{n+1} \setminus \left(\bigcup_{i=1}^n A_i \right) \right] \quad (1.7.7)$$

and the two sets connected by the central union are disjoint. Using the [addition rule](#) and the [difference rule](#),

$$\# \left(\bigcup_{i=1}^{n+1} A_i \right) = \# \left(\bigcup_{i=1}^n A_i \right) + \#(A_{n+1}) - \# \left[A_{n+1} \cap \left(\bigcup_{i=1}^n A_i \right) \right] \quad (1.7.8)$$

$$= \# \left(\bigcup_{i=1}^n A_i \right) + \#(A_{n+1}) - \# \left[\bigcup_{i=1}^n (A_i \cap A_{n+1}) \right] \quad (1.7.9)$$

By the induction hypothesis, the formula holds for the two unions of n sets in the last expression. The result then follows by simplification.

The general *Bonferroni inequalities*, named again for Carlo Bonferroni, state that if sum on the right is truncated after k terms ($k < n$), then the truncated sum is an *upper* bound for the cardinality of the union if k is odd (so that the last term has a positive sign) and is a *lower* bound for the cardinality of the union if k is even (so that the last terms has a negative sign).

The Multiplication Rule

The *multiplication rule* of combinatorics is based on the formulation of a *procedure* (or *algorithm*) that generates the objects to be counted.

Suppose that a procedure consists of k steps, performed sequentially, and that for each $j \in \{1, 2, \dots, k\}$, step j can be performed in n_j ways, regardless of the choices made on the previous steps. Then the number of ways to perform the entire procedure is $n_1 n_2 \cdots n_k$.

The key to a successful application of the multiplication rule to a counting problem is the clear formulation of an algorithm that generates the objects being counted, so that each object is generated once and only once. That is, we must neither over count nor under count. It's also important to notice that the set of choices available at step j may well depend on the previous steps; the assumption is only that the *number* of choices available does not depend on the previous steps.

The first two results below give equivalent formulations of the multiplication principle.

Suppose that S is a set of sequences of length k , and that we denote a generic element of S by (x_1, x_2, \dots, x_k) . Suppose that for each $j \in \{1, 2, \dots, k\}$, x_j has n_j different values, regardless of the values of the previous coordinates. Then $\#(S) = n_1 n_2 \cdots n_k$.

Proof

A procedure that generates the sequences in S consists of k steps. Step j is to select the j th coordinate.

Suppose that T is an ordered tree with depth k and that each vertex at level $i - 1$ has n_i children for $i \in \{1, 2, \dots, k\}$. Then the number of endpoints of the tree is $n_1 n_2 \cdots n_k$.

Proof

Each endpoint of the tree is uniquely associated with the path from the root vertex to the endpoint. Each such path is a sequence of length k , in which there are n_j values for coordinate j for each $j \in \{1, 2, \dots, k\}$. Hence the result follows from the result above on sequences.

Product Sets

If S_i is a set with n_i elements for $i \in \{1, 2, \dots, k\}$ then

$$\#(S_1 \times S_2 \times \cdots \times S_k) = n_1 n_2 \cdots n_k \quad (1.7.10)$$

Proof

This is a corollary of the result above on sequences.

If S is a set with n elements, then S^k has n^k elements.

Proof

This is a corollary of the previous result.

In (16), note that the elements of S^k can be thought of as ordered samples of size k that can be chosen with replacement from a population of n objects. Elements of $\{0, 1\}^n$ are sometimes called *bit strings* of length n . Thus, there are 2^n bit strings of length n .

Functions

The number of functions from a set A of m elements into a set B of n elements is n^m .

Proof

An algorithm for constructing a function $f : A \rightarrow B$ is to choose the value of $f(x) \in B$ for each $x \in A$. There are n choices for each of the m elements in the domain.

Recall that the set of functions from a set A into a set B (regardless of whether the sets are finite or infinite) is denoted B^A . This theorem is motivation for the notation. Note also that if S is a set with n elements, then the elements in the Cartesian power set S^k can be thought of as functions from $\{1, 2, \dots, k\}$ into S . So the counting formula for sequences can be thought of as a corollary of counting formula for functions.

Subsets

Suppose that S is a set with n elements, where $n \in \mathbb{N}$. There are 2^n subsets of S .

Proof from the multiplication principle

An algorithm for constructing $A \subseteq S$, is to decide whether $x \in A$ or $x \notin A$ for each $x \in S$. There are 2 choices for each of the n elements of S .

Proof using indicator functions

Recall that there is a one-to-one correspondence between subsets of S and indicator functions on S . An indicator function is simply a function from S into $\{0, 1\}$, and the number of such functions is 2^n by the previous result.

Suppose that $\{A_1, A_2, \dots, A_k\}$ is a collection of k subsets of a set S , where $k \in \mathbb{N}_+$. There are 2^{2^k} different (in general) sets that can be constructed from the k given sets, using the operations of union, intersection, and complement. These sets form the algebra generated by the given sets.

Proof

First note that there are 2^k pairwise disjoint sets of the form $B_1 \cap B_2 \cap \cdots \cap B_k$ where $B_i = A_i$ or $B_i = A_i^c$ for each i . Next, note that every set that can be constructed from $\{A_1, A_2, \dots, A_n\}$ is a union of some (perhaps all, perhaps none) of these intersection sets.

Open the Venn diagram app.

1. Select each of the 4 disjoint sets $A \cap B$, $A \cap B^c$, $A^c \cap B$, $A^c \cap B^c$.
2. Select each of the 12 other subsets of S . Note how each is a union of some of the sets in (a).

Suppose that S is a set with n elements and that A is a subset of S with k elements, where $n, k \in \mathbb{N}$ and $k \leq n$. The number of subsets of S that contain A is 2^{n-k} .

Proof

Note that subset B of S that contains A can be written uniquely in the form $B = A \cup C$ where $C \subseteq A^c$. A^c has $n - k$ elements and hence there are 2^{n-k} subsets of A^c by the general [subset result](#).

Our last result in this discussion generalizes the basic [subset result](#) above.

Suppose that $n, k \in \mathbb{N}$ and that S is a set with n elements. The number of sequences of subsets (A_1, A_2, \dots, A_k) with $A_1 \subseteq A_2 \subseteq \cdots \subseteq A_k \subseteq S$ is $(k+1)^n$.

Proof

To construct a sequence of the type in the theorem, we can use the following algorithm: For each $x \in S$, either x is not in the sets, or x occurs for the first time in set A_i where $i \in \{1, 2, \dots, k\}$. (That is, $x \notin A_j$ for $j \in \{1, \dots, i-1\}$ and $x \in A_i$ for $j \in \{i, \dots, k\}$.) So there are $k+1$ choices for each of the n elements of S .

When $k = 1$ we get 2^n as the number of subsets of S , as before.

Computational Exercises

Identification Numbers

A license number consists of two letters (uppercase) followed by five digits. How many different license numbers are there?

Answer

$$26^2 \cdot 10^5 = 67\,600\,000$$

Suppose that a Personal Identification Number (PIN) is a four-symbol code word in which each entry is either a letter (uppercase) or a digit. How many PINs are there?

Answer

$$36^4 = 1\,679\,616$$

Cards, Dice, and Coins

In the board game [Clue](#), Mr. Boddy has been murdered. There are 6 suspects, 6 possible weapons, and 9 possible rooms for the murder.

1. The game includes a card for each suspect, each weapon, and each room. How many cards are there?
2. The outcome of the game is a sequence consisting of a suspect, a weapon, and a room (for example, *Colonel Mustard with the knife in the billiard room*). How many outcomes are there?
3. Once the three cards that constitute the outcome have been randomly chosen, the remaining cards are dealt to the players. Suppose that you are dealt 5 cards. In trying to guess the outcome, what hand of cards would be best?

Answer

1. $6 + 6 + 9 = 21$ cards

2. $6 \cdot 6 \cdot 9 = 324$ outcomes

3. The best hand would be the 5 remaining weapons or the 5 remaining suspects.

An experiment consists of rolling a standard die, drawing a card from a standard deck, and tossing a standard coin. How many outcomes are there?

Answer

$$6 \cdot 52 \cdot 2 = 624$$

A standard die is rolled 5 times and the sequence of scores recorded. How many outcomes are there?

Answer

$$6^5 = 7776$$

In the card game *Set*, each card has 4 properties: *number* (one, two, or three), *shape* (diamond, oval, or squiggle), *color* (red, blue, or green), and *shading* (solid, open, or stripped). The deck has one card of each (number, shape, color, shading) configuration. A *set* in the game is defined as a set of three cards which, for each property, the cards are either all the same or all different.

1. How many cards are in a deck?

2. How many sets are there?

Answer

1. $3^4 = 81$

2. 1080

A coin is tossed 10 times and the sequence of scores recorded. How many sequences are there?

Answer

$$2^{10} = 1024$$

The die-coin experiment consists of rolling a die and then tossing a coin the number of times shown on the die. The sequence of coin results is recorded.

1. How many outcomes are there?

2. How many outcomes are there with all heads?

3. How many outcomes are there with exactly one head?

Answer

1. $\sum_{k=1}^6 2^k = 126$

2. 6

3. $\sum_{k=1}^6 k = 21$

Run the die-coin experiment 100 times and observe the outcomes.

Consider a deck of cards as a set D with 52 elements.

1. How many subsets of D are there?

2. How many functions are there from D into the set $\{1, 2, 3, 4\}$?

Answer

1. $2^{52} = 4\,503\,599\,627\,370\,496$

2. $4^{52} = 20\,282\,409\,603\,651\,670\,423\,947\,251\,286\,016$

Birthdays

Consider a group of 10 persons.

1. If we record the birth month of each person, how many outcomes are there?
2. If we record the birthday of each person (ignoring leap day), how many outcomes are there?

Answer

1. $12^{10} = 61\,917\,364\,224$
2. $365^{10} = 41\,969\,002\,243\,198\,805\,166\,015\,625$

Reliability

In the usual model of *structural reliability*, a system consists of *components*, each of which is either *working* or *defective*. The system as a whole is also either working or defective, depending on the states of the components and how the components are connected.

A string of lights has 20 bulbs, each of which may be good or defective. How many configurations are there?

Answer

$$2^{20} = 1\,048\,576$$

If the components are connected in *series*, then the system as a whole is working if and only if each component is working. If the components are connected *parallel*, then the system as a whole is working if and only if at least one component is working.

A system consists of three subsystems with 6, 5, and 4 components, respectively. Find the number of component states for which the system is working in each of the following cases:

1. The components in each subsystem are in parallel and the subsystems are in series.
2. The components in each subsystem are in series and the subsystems are in parallel.

Answer

1. $(2^6 - 1)(2^5 - 1)(2^4 - 1) = 29\,295$
2. $2^3 - 1 = 7$

Menus

Suppose that a sandwich at a restaurant consists of bread, meat, cheese, and various toppings. There are 4 choices for the bread, 3 choices for the meat, 5 choices for the cheese, and 10 different toppings (each of which may be chosen). How many sandwiches are there?

Answer

$$4 \cdot 3 \cdot 5 \cdot 2^{10} = 61\,440$$

At a wedding dinner, there are three choices for the entrée, four choices for the beverage, and two choices for the dessert.

1. How many different meals are there?
2. If there are 50 guests at the wedding and we record the meal requested for each guest, how many possible outcomes are there?

Answer

1. $3 \cdot 4 \cdot 2 = 24$
2. $24^{50} \approx 1.02462 \times 10^{69}$

Braille

Braille is a tactile writing system used by people who are visually impaired. The system is named for the French educator Louis Braille and uses raised dots in a 3×2 grid to encode characters. How many meaningful Braille configurations are there?

Answer



Figure 1.7.5: The Braille encoding of the number 2 and the letter *b*

Personality Typing

The *Meyers-Briggs* personality typing is based on four dichotomies: A person is typed as either *extroversion* (E) or *introversion* (I), either *sensing* (S) or *intuition* (I), either *thinking* (T) or *feeling* (F), and either *judgement* (J) or *perception* (P).

1. How many Meyers-Briggs personality types are there? List them.
2. Suppose that we list the personality types of 10 persons. How many possible outcomes are there?

Answer

1. 16
2. $16^{10} = 1\,099\,511\,627\,776$

The Galton Board

The *Galton Board*, named after Francis Galton, is a triangular array of pegs. Galton, apparently too modest to name the device after himself, called it a *quincunx* from the Latin word for five twelfths (go figure). The rows are numbered, from the top down, by $(0, 1, \dots)$. Row n has $n + 1$ pegs that are labeled, from left to right by $(0, 1, \dots, n)$. Thus, a peg can be uniquely identified by an ordered pair (n, k) where n is the row number and k is the peg number in that row.

A ball is dropped onto the top peg $(0, 0)$ of the Galton board. In general, when the ball hits peg (n, k) , it either bounces to the left to peg $(n + 1, k)$ or to the right to peg $(n + 1, k + 1)$. The sequence of pegs that the ball hits is a *path* in the Galton board.

There is a one-to-one correspondence between each pair of the following three collections:

1. Bit strings of length n
2. Paths in the Galton board from $(0, 0)$ to any peg in row n .
3. Subsets of a set with n elements.

Thus, each of these collections has 2^n elements.

Open the Galton board app.

1. Move the ball from $(0, 0)$ to $(10, 6)$ along a path of your choice. Note the corresponding bit string and subset.
2. Generate the bit string 0111001010 Note the corresponding subset and path.
3. Generate the subset $\{2, 4, 5, 9, 10\}$ Note the corresponding bit string and path.
4. Generate all paths from $(0, 0)$ to $(4, 2)$. How many paths are there?

Answer

4. 6

This page titled [1.7: Counting Measure](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.8: Combinatorial Structures

The purpose of this section is to study several combinatorial structures that are of basic importance in probability.

Permutations

Suppose that D is a set with $n \in \mathbb{N}$ elements. A *permutation* of length $k \in \{0, 1, \dots, n\}$ from D is an ordered sequence of k distinct elements of D ; that is, a sequence of the form (x_1, x_2, \dots, x_k) where $x_i \in D$ for each i and $x_i \neq x_j$ for $i \neq j$.

Statistically, a permutation of length k from D corresponds to an *ordered sample* of size k chosen *without replacement* from the population D .

The number of permutations of length k from an n element set is

$$n^{(k)} = n(n-1) \cdots (n-k+1) \quad (1.8.1)$$

Proof

This follows easily from the multiplication principle. There are n ways to choose the first element, $n-1$ ways to choose the second element, and so forth.

By convention, $n^{(0)} = 1$. Recall that, in general, a product over an empty index set is 1. Note that $n^{(k)}$ has k factors, starting at n , and with each subsequent factor one less than the previous factor. Some alternate notations for the number of permutations of size k from a set of n objects are $P(n, k)$, $P_{n,k}$, and ${}_nP_k$.

The number of permutations of length n from the n element set D (these are called simply *permutations* of D) is

$$n! = n^{(n)} = n(n-1) \cdots (1) \quad (1.8.2)$$

The function on \mathbb{N} given by $n \mapsto n!$ is the *factorial function*. The general permutation formula in (2) can be written in terms of factorials:

For $n \in \mathbb{N}$ and $k \in \{0, 1, \dots, n\}$

$$n^{(k)} = \frac{n!}{(n-k)!} \quad (1.8.3)$$

Although this formula is succinct, it's not always a good idea numerically. If n and $n-k$ are large, $n!$ and $(n-k)!$ are enormous, and division of the first by the second can lead to significant round-off errors.

Note that the basic permutation formula in (2) is defined for every real number n and nonnegative integer k . This extension is sometimes referred to as the *generalized permutation formula*. Actually, we will sometimes need an even more general formula of this type (particularly in the sections on [Pólya's urn](#) and the [beta-Bernoulli process](#)).

For $a \in \mathbb{R}$, $s \in \mathbb{R}$, and $k \in \mathbb{N}$, define

$$a^{(s,k)} = a(a+s)(a+2s) \cdots [a+(k-1)s] \quad (1.8.4)$$

1. $a^{(0,k)} = a^k$
2. $a^{(-1,k)} = a^{(k)}$
3. $a^{(1,k)} = a(a+1) \cdots (a+k-1)$
4. $1^{(1,k)} = k!$

The product $a^{(-1,k)} = a^{(k)}$ (our ordinary permutation formula) is sometimes called the *falling power* of a of order k , while $a^{(1,k)}$ is called the *rising power* of a of order k , and is sometimes denoted $a^{[k]}$. Note that $a^{(0,k)}$ is the ordinary k th power of a . In general, note that $a^{(s,k)}$ has k factors, starting at a and with each subsequent factor obtained by adding s to the previous factor.

Combinations

Consider again a set D with $n \in \mathbb{N}$ elements. A *combination* of size $k \in \{0, 1, \dots, n\}$ from D is an (unordered) subset of k distinct elements of D . Thus, a combination of size k from D has the form $\{x_1, x_2, \dots, x_k\}$, where $x_i \in D$ for each i and $x_i \neq x_j$ for $i \neq j$.

Statistically, a combination of size k from D corresponds to an *unordered sample* of size k chosen *without replacement* from the population D . Note that for each combination of size k from D , there are $k!$ distinct orderings of the elements of that combination. Each of these is a permutation of length k from D . The number of combinations of size k from an n -element set is denoted by $\binom{n}{k}$. Some alternate notations are $C(n, k)$, $C_{n,k}$, and ${}_nC_k$.

The number of combinations of size k from an n element set is

$$\binom{n}{k} = \frac{n^{(k)}}{k!} = \frac{n!}{k!(n-k)!} \quad (1.8.5)$$

Proof

An algorithm for generating all permutations of size k from D is to first select a combination of size k and then to select an ordering of the elements. From the multiplication principle it follows that $n^{(k)} = \binom{n}{k}k!$. Hence $\binom{n}{k} = n^{(k)} / k! = n! / [k!(n-k)!]$.

The number $\binom{n}{k}$ is called a *binomial coefficient*. Note that the formula makes sense for any real number n and nonnegative integer k since this is true of the generalized permutation formula $n^{(k)}$. With this extension, $\binom{n}{k}$ is called the *generalized binomial coefficient*. Note that if n and k are positive integers and $k > n$ then $\binom{n}{k} = 0$. By convention, we will also define $\binom{n}{k} = 0$ if $k < 0$. This convention sometimes simplifies formulas.

Properties of Binomial Coefficients

For some of the identities below, there are two possible proofs. An algebraic proof, of course, should be based on (5). A *combinatorial proof* is constructed by showing that the left and right sides of the identity are two different ways of counting the same collection.

$$\binom{n}{n} = \binom{n}{0} = 1.$$

Algebraically, the last result is trivial. It also makes sense combinatorially: There is only one way to select a subset of D with n elements (D itself), and there is only one way to select a subset of size 0 from D (the empty set \emptyset).

If $n, k \in \mathbb{N}$ with $k \leq n$ then

$$\binom{n}{k} = \binom{n}{n-k} \quad (1.8.6)$$

Combinatorial Proof

Note that if we select a subset of size k from a set of size n , then we leave a subset of size $n - k$ behind (the complement). Thus $A \mapsto A^c$ is a one-to-one correspondence between subsets of size k and subsets of size $n - k$.

The next result is one of the most famous and most important. It's known as *Pascal's rule* and is named for Blaise Pascal.

If $n, k \in \mathbb{N}_+$ with $k \leq n$ then

$$\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k} \quad (1.8.7)$$

Combinatorial Proof

Suppose that we have n persons, one named Fred, and that we want to select a committee of size k . There are $\binom{n}{k}$ different committees. On the other hand, there are $\binom{n-1}{k-1}$ committees with Fred as a member, and $\binom{n-1}{k}$ committees without Fred as a member. The sum of these two numbers is also the number of committees.

Recall that the *Galton board* is a triangular array of pegs: the rows are numbered $n = 0, 1, \dots$ and the pegs in row n are numbered $k = 0, 1, \dots, n$. If each peg in the Galton board is replaced by the corresponding binomial coefficient, the resulting table of numbers is known as *Pascal's triangle*, named again for Pascal. By (8), each interior number in Pascal's triangle is the sum of the two numbers directly above it.

The following result is the *binomial theorem*, and is the reason for the term *binomial coefficient*.

If $a, b \in \mathbb{R}$ and $n \in \mathbb{N}$ is a positive integer, then

$$(a + b)^n = \sum_{k=0}^n \binom{n}{k} a^k b^{n-k} \quad (1.8.8)$$

Combinatorial Proof

Note that to get the term $a^k b^{n-k}$ in the expansion of $(a + b)^n$, we must select a from k of the factors and b from the remaining $n - k$ factors. The number of ways to do this is $\binom{n}{k}$.

If $j, k, n \in \mathbb{N}_+$ with $j \leq k \leq n$ then

$$k^{(j)} \binom{n}{k} = n^{(j)} \binom{n-j}{k-j} \quad (1.8.9)$$

Combinatorial Proof

Consider two procedures for selecting a committee of size k from a group of n persons, with j distinct members of the committee as officers (chair, vice chair, etc.). For the first procedure, select the committee from the population and then select the member of the committee to be the officers. The number of ways to perform the first step is $\binom{n}{k}$ and the number of ways to perform the second step is $k^{(j)}$. So by the multiplication principle, the number of ways to choose the committee is the left side of the equation. For the second procedure, select the officers of the committee from the population and then select $k - j$ other committee members from the remaining $n - j$ members of the population. The number of ways to perform the first step is $n^{(j)}$ and the number of ways to perform the second step is $\binom{n-j}{k-j}$. So by the multiplication principle, the number of committees is the right side of the equation.

The following result is known as *Vandermonde's identity*, named for Alexandre-Théophile Vandermonde.

If $m, n, k \in \mathbb{N}$ with $k \leq m + n$, then

$$\sum_{j=0}^k \binom{m}{j} \binom{n}{k-j} = \binom{m+n}{k} \quad (1.8.10)$$

Combinatorial Proof

Suppose that a committee of size k is chosen from a group of $m + n$ persons, consisting of m men and n women. The number of committees with exactly j men and $k - j$ women is $\binom{m}{j} \binom{n}{k-j}$. The sum of this product over j is the total number of committees, which is $\binom{m+n}{k}$.

The next result is a general identity for the sum of binomial coefficients.

If $m, n \in \mathbb{N}$ with $n \leq m$ then

$$\sum_{j=n}^m \binom{j}{n} = \binom{m+1}{n+1} \quad (1.8.11)$$

Combinatorial Proof

Suppose that we pick a subset of size $n + 1$ from the set $\{1, 2, \dots, m + 1\}$. For $j \in \{n, n + 1, \dots, m\}$, the number of subsets in which the largest element is $j + 1$ is $\binom{j}{n}$. Hence the sum of these numbers over j is the total number of subsets of size $n + 1$, which is also $\binom{m+1}{n+1}$.

For an even more general version of the last result, see the section on [Order Statistics](#) in the chapter on Finite Sampling Models. The following identity for the sum of the first m positive integers is a special case of the last result.

If $m \in \mathbb{N}$ then

$$\sum_{j=1}^m j = \binom{m+1}{2} = \frac{(m+1)m}{2} \quad (1.8.12)$$

Proof

Let $n = 1$ in previous result.

There is a one-to-one correspondence between each pair of the following collections. Hence the number objects in each of these collection is $\binom{n}{k}$.

1. Subsets of size k from a set of n elements.
2. Bit strings of length n with exactly k 1's.
3. Paths in the Galton board from $(0, 0)$ to (n, k) .

Proof

Let $S = \{x_1, x_2, \dots, x_n\}$ be a set with n elements. A one-to-one correspondence between the subsets A of S with k elements and the bit strings $\mathbf{b} = b_1 b_2 \dots b_n$ of length n with k 1's can be constructed by the rule that $x_i \in A$ if and only if $b_i = 1$. In turn, a one-to-one correspondence between the bit strings \mathbf{b} in part (b) and the paths in Galton board in part (c) can be constructed by the rule that in row $i \in \{0, 1, \dots, n-1\}$, turn right if $b_{i+1} = 1$ and turn left if $b_{i+1} = 0$.

The following identity is known as the *alternating sum identity* for binomial coefficients. It turns out to be useful in the [Irwin-Hall probability distribution](#). We give the identity in two equivalent forms, one for falling powers and one for ordinary powers.

If $n \in \mathbb{N}_+$, $j \in \{0, 1, \dots, n-1\}$ then

$$1. \quad \sum_{k=0}^n \binom{n}{k} (-1)^k k^{(j)} = 0 \quad (1.8.13)$$

$$2. \quad \sum_{k=0}^n \binom{n}{k} (-1)^k k^j = 0 \quad (1.8.14)$$

Proof

1. We use the [identity](#) above and the binomial theorem [binomial theorem](#):

$$\begin{aligned} \sum_{k=0}^n (-1)^k k^{(j)} \binom{n}{k} &= \sum_{k=j}^n (-1)^k k^{(j)} \binom{n}{k} = \sum_{k=j}^n (-1)^k n^{(j)} \binom{n-j}{k-j} \\ &= n^{(j)} (-1)^j \sum_{k=j}^n (-1)^{k-j} \binom{n-j}{k-j} = n^{(j)} (-1)^j \sum_{i=0}^{n-j} (-1)^i \binom{n-j}{i} \\ &= n^{(j)} (-1)^j (-1+1)^{n-j} = 0. \end{aligned}$$

Note that it's the last step where we need $j < n$.

2. This follows from (a), since k^j is a linear combination of $k^{(i)}$ for $i \in \{0, 1, \dots, j\}$. That is, there exists $c_i \in \mathbb{R}$ for $i \in \{0, 1, \dots, j\}$ such that $k^j = \sum_{i=0}^j c_i k^{(i)}$. Hence

$$\sum_{k=0}^n (-1)^k k^j \binom{n}{k} = \sum_{i=0}^j c_i \sum_{k=0}^n (-1)^k k^{(i)} \binom{n}{k} = 0 \quad (1.8.15)$$

Our next identity deals with a generalized binomial coefficient.

If $n, k \in \mathbb{N}$ then

$$\binom{-n}{k} = (-1)^k \binom{n+k-1}{k} \quad (1.8.16)$$

Proof

Note that

$$\binom{-n}{k} = \frac{(-n)^{(k)}}{k!} = \frac{(-n)(-n-1)\cdots(-n-k+1)}{k!} \quad (1.8.17)$$

$$= (-1)^k \frac{(n+k-1)(n+k-2)\cdots(n)}{k!} = (-1)^k \frac{(n+k-1)^{(k)}}{k!} = (-1)^k \binom{n+k-1}{k} \quad (1.8.18)$$

In particular, note that $\binom{-1}{k} = (-1)^k$. Our last result in this discussion concerns the binomial operator and its inverse.

The *binomial operator* takes a sequence of real numbers $\mathbf{a} = (a_0, a_1, a_2, \dots)$ and returns the sequence of real numbers $\mathbf{b} = (b_0, b_1, b_2, \dots)$ by means of the formula

$$b_n = \sum_{k=0}^n \binom{n}{k} a_k, \quad n \in \mathbb{N} \quad (1.8.19)$$

The *inverse binomial operator* recovers the sequence \mathbf{a} from the sequence \mathbf{b} by means of the formula

$$a_n = \sum_{k=0}^n (-1)^{n-k} \binom{n}{k} b_k, \quad n \in \mathbb{N} \quad (1.8.20)$$

Proof

Exponential generating functions can be used for an elegant proof. Exponential generating functions are the combinatorial equivalent of moment generating functions for discrete probability distributions on \mathbb{N} . So let G and H denote the exponential generating functions of the sequences \mathbf{a} and \mathbf{b} , respectively. Then

$$\begin{aligned} H(x) &= \sum_{n=0}^{\infty} \frac{b_n}{n!} x^n = \sum_{n=0}^{\infty} \frac{1}{n!} \sum_{k=0}^n \binom{n}{k} a_k x^n = \sum_{k=0}^{\infty} \sum_{n=k}^{\infty} \frac{1}{n!} \frac{n!}{k!(n-k)!} a_k x^n \\ &= \sum_{k=0}^{\infty} \frac{1}{k!} a_k x^k \sum_{n=k}^{\infty} \frac{1}{(n-k)!} x^{n-k} = e^x \sum_{k=0}^{\infty} \frac{1}{k!} a_k x^k = e^x G(x) \end{aligned}$$

So it follows that

$$\begin{aligned} G(x) &= e^{-x} H(x) = \sum_{k=0}^{\infty} \frac{1}{k!} b_k x^k \sum_{n=k}^{\infty} \frac{1}{(n-k)!} (-1)^{n-k} x^{n-k} \\ &= \sum_{k=0}^{\infty} \sum_{n=k}^{\infty} \frac{1}{n!} \frac{n!}{k!(n-k)!} (-1)^{n-k} b_k x^n = \sum_{n=0}^{\infty} \frac{1}{n!} \sum_{k=0}^n \binom{n}{k} (-1)^{n-k} b_k x^n \end{aligned}$$

But by definition,

$$G(x) = \sum_{n=0}^{\infty} \frac{a_n}{n!} x^n \quad (1.8.21)$$

and so the inverse formula follows.

Samples

The experiment of drawing a sample from a population is basic and important. There are two essential attributes of samples: whether or not *order* is important, and whether or not a sampled object is *replaced* in the population before the next draw. Suppose now that the population D contains n objects and we are interested in drawing a sample of k objects. Let's review what we know so far:

- If order is important and sampled objects are replaced, then the samples are just elements of the product set D^k . Hence, the number of samples is n^k .
- If order is important and sample objects are not replaced, then the samples are just permutations of size k chosen from D . Hence the number of samples is $n^{(k)}$.
- If order is not important and sample objects are not replaced, then the samples are just combinations of size k chosen from D . Hence the number of samples is $\binom{n}{k}$.

Thus, we have one case left to consider.

Unordered Samples With Replacement

An unordered sample chosen *with* replacement from D is called a *multiset*. A multiset is like an ordinary set except that elements may be repeated.

There is a one-to-one correspondence between each pair of the following collections:

1. Multisets of size k from a population D of n elements.
2. Bit strings of length $n + k - 1$ with exactly k 1s.
3. Nonnegative integer solutions (x_1, x_2, \dots, x_n) of the equation $x_1 + x_2 + \dots + x_n = k$.

Each of these collections has $\binom{n+k-1}{k}$ members.

Proof

Suppose that $D = \{d_1, d_2, \dots, d_n\}$. Consider a multiset of size k . Since order does not matter, we can list all of the occurrences of d_1 (if any) first, then the occurrences of d_2 (if any), and so forth, until we at last list the occurrences of d_n (if any). If we know we are using this data structure, we don't actually have to list the actual elements, we can simply use 1 as a placeholder with 0 as a separator. In the resulting bit string, 1 occurs k times and 0 occurs $n - 1$ times. Conversely, any such bit string uniquely defines a multiset of size k . Next, given a multiset of size k from D , let x_i denote the number of times that d_i occurs, for $i \in \{1, 2, \dots, n\}$. Then (x_1, x_2, \dots, x_n) satisfies the conditions in (c). Conversely, any solution to the equation in (c) uniquely defines a multiset of size k from D . We already know how to count the collection in (b): the number of bit strings of length $n + k - 1$ with 1 occurring k times is $\binom{n+k-1}{k} = \binom{n+k-1}{n-1}$.

Summary of Sampling Formulas

The following table summarizes the formulas for the number of samples of size k chosen from a population of n elements, based on the criteria of order and replacement.

Sampling formulas

Number of Samples	With order	Without
With replacement	n^k	$\binom{n+k-1}{k}$
Without	$n^{(k)}$	$\binom{n}{k}$

Multinomial Coefficients

Partitions of a Set

Recall that the binomial coefficient $\binom{n}{j}$ is the number of subsets of size j from a set S of n elements. Note also that when we select a subset A of size j from S , we effectively *partition* S into two disjoint subsets of sizes j and $n - j$, namely A and A^c . A natural generalization is to *partition* S into a union of k distinct, pairwise disjoint subsets (A_1, A_2, \dots, A_k) where $\#(A_i) = n_i$ for each $i \in \{1, 2, \dots, k\}$. Of course we must have $n_1 + n_2 + \dots + n_k = n$.

The number of ways to partition a set of n elements into a sequence of k sets of sizes (n_1, n_2, \dots, n_k) is

$$\binom{n}{n_1} \binom{n-n_1}{n_2} \dots \binom{n-n_1-\dots-n_{k-1}}{n_k} = \frac{n!}{n_1! n_2! \dots n_k!} \quad (1.8.22)$$

Proof

The left side follows from the multiplication rule. There are $\binom{n}{n_1}$ ways to select the first set in the partition, $\binom{n-n_1}{n_2}$ ways to select the second set in the partition, and so forth. The right side follows by writing the binomial coefficients on the left in terms of factorials and simplifying.

The number in (18) is called a *multinomial coefficient* and is denoted by

$$\binom{n}{n_1, n_2, \dots, n_k} = \frac{n!}{n_1! n_2! \dots n_k!} \quad (1.8.23)$$

If $n, k \in \mathbb{N}$ with $k \leq n$ then

$$\binom{n}{k, n-k} = \binom{n}{k} \quad (1.8.24)$$

Combinatorial Proof

As noted before, if we select a subset of size k from an n element set, then we partition the set into two subsets of sizes k and $n-k$.

Sequences

Consider now the set $T = \{1, 2, \dots, k\}^n$. Elements of this set are sequences of length n in which each coordinate is one of k values. Thus, these sequences generalize the bit strings of length n . Again, let (n_1, n_2, \dots, n_k) be a sequence of nonnegative integers with $\sum_{i=1}^k n_i = n$.

There is a one-to-one correspondence between the following collections:

1. Partitions of S into pairwise disjoint subsets (A_1, A_2, \dots, A_k) where $\#(A_j) = n_j$ for each $j \in \{1, 2, \dots, k\}$.
2. Sequences in $\{1, 2, \dots, k\}^n$ in which j occurs n_j times for each $j \in \{1, 2, \dots, k\}$.

Proof

Suppose that $S = \{s_1, s_2, \dots, s_n\}$. A one-to-one correspondence between a partition (A_1, A_2, \dots, A_k) of the type in (a) and a sequence $\mathbf{x} = (x_1, x_2, \dots, x_n)$ of the type in (b) can be constructed by the rule that $s_i \in A_j$ if and only if $x_i = j$.

It follows that the number of elements in both of these collections is

$$\binom{n}{n_1, n_2, \dots, n_k} = \frac{n!}{n_1! n_2! \dots n_k!} \quad (1.8.25)$$

Permutations with Indistinguishable Objects

Suppose now that we have n object of k different types, with n_i elements of type i for each $i \in \{1, 2, \dots, k\}$. Moreover, objects of a given type are considered identical. There is a one-to-one correspondence between the following collections:

1. Sequences in $\{1, 2, \dots, k\}^n$ in which j occurs n_j times for each $j \in \{1, 2, \dots, k\}$.
2. Distinguishable permutations of the n objects.

Proof

A one-to-one correspondence between a sequence $\mathbf{x} = (x_1, x_2, \dots, x_n)$ of the type in (a) and a permutation of the n objects can be constructed by the rule that we put an object of type j in position i if and only if $x_i = j$.

Once again, it follows that the number of elements in both collections is

$$\binom{n}{n_1, n_2, \dots, n_k} = \frac{n!}{n_1! n_2! \dots n_k!} \quad (1.8.26)$$

The Multinomial Theorem

The following result is the *multinomial theorem* which is the reason for the name of the coefficients.

If $x_1, x_2, \dots, x_k \in \mathbb{R}$ and $n \in \mathbb{N}$ then

$$(x_1 + x_2 + \cdots + x_k)^n = \sum \binom{n}{n_1, n_2, \dots, n_k} x_1^{n_1} x_2^{n_2} \cdots x_k^{n_k} \quad (1.8.27)$$

The sum is over sequences of nonnegative integers (n_1, n_2, \dots, n_k) with $n_1 + n_2 + \cdots + n_k = n$. There are $\binom{n+k-1}{n}$ terms in this sum.

Combinatorial Proof

Note that to get $x_1^{n_1} x_2^{n_2} \cdots x_k^{n_k}$ in the expansion of $(x_1 + x_2 + \cdots + x_k)^n$, we must choose x_i in n_i of the factors, for each i . The number of ways to do this is the multinomial coefficient $\binom{n}{n_1, n_2, \dots, n_k}$. The number of terms in the sum follows from the [formula above](#).

Computational Exercises

Arrangements

In a race with 10 horses, the first, second, and third place finishers are noted. How many outcomes are there?

Answer

720

Eight persons, consisting of four male-female couples, are to be seated in a row of eight chairs. How many seating arrangements are there in each of the following cases:

1. There are no other restrictions.
2. The men must sit together and the women must sit together.
3. The men must sit together.
4. Each couple must sit together.

Answer

1. 40 320
2. 1152
3. 2880
4. 384

Suppose that n people are to be seated at a round table. How many seating arrangements are there? The mathematical significance of a round table is that there is no dedicated *first* chair.

Answer

$(n - 1)!$. Seat one, distinguished person arbitrarily. Every seating arrangement can then be specified by giving the position of a person (say clockwise) relative to the distinguished person.

Twelve books, consisting of 5 math books, 4 science books, and 3 history books are arranged on a bookshelf. Find the number of arrangements in each of the following cases:

1. There are no restrictions.
2. The books of each type must be together.
3. The math books must be together.

Answer

1. 479 001 600
2. 103 680
3. 4 838 400

Find the number of distinct arrangements of the letters in each of the following words:

1. statistics
2. probability

3. mississippi
4. tennessee
5. alabama

Answer

1. 50 400
2. 9 979 200
3. 34 650
4. 3780
5. 210

A child has 12 blocks; 5 are red, 4 are green, and 3 are blue. In how many ways can the blocks be arranged in a line if blocks of a given color are considered identical?

Answer

27 720

Code Words

A license tag consists of 2 capital letters and 5 digits. Find the number of tags in each of the following cases:

1. There are no other restrictions
2. The letters and digits are all different.

Answer

1. 67 600 000
2. 19 656 000

Committees

A club has 20 members; 12 are women and 8 are men. A committee of 6 members is to be chosen. Find the number of different committees in each of the following cases:

1. There are no other restrictions.
2. The committee must have 4 women and 2 men.
3. The committee must have at least 2 women and at least 2 men.

Answer

1. 38 760
2. 13 860
3. 30 800

Suppose that a club with 20 members plans to form 3 distinct committees with 6, 5, and 4 members, respectively. In how many ways can this be done.

Answer

9 777 287 520 Note that the members *not* on a committee also form one of the sets in the partition.

Cards

A standard *card deck* can be modeled by the Cartesian product set

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (1.8.28)$$

where the first coordinate encodes the *denomination* or *kind* (ace, 2-10, jack, queen, king) and where the second coordinate encodes the *suit* (clubs, diamonds, hearts, spades). Sometimes we represent a card as a *string* rather than an ordered pair (for example $q\heartsuit$).

A *poker hand* (in draw poker) consists of 5 cards dealt without replacement and without regard to order from a deck of 52 cards. Find the number of poker hands in each of the following cases:

1. There are no restrictions.
2. The hand is a *full house* (3 cards of one kind and 2 of another kind).
3. The hand has *4 of a kind*.
4. The cards are all in the same suit (so the hand is a *flush* or a *straight flush*).

Answer

1. 2 598 960
2. 3744
3. 624
4. 5148

The game of **poker** is studied in detail in the chapter on Games of Chance.

A *bridge hand* consists of 13 cards dealt without replacement and without regard to order from a deck of 52 cards. Find the number of bridge hands in each of the following cases:

1. There are no restrictions.
2. The hand has exactly 4 spades.
3. The hand has exactly 4 spades and 3 hearts.
4. The hand has exactly 4 spades, 3 hearts, and 2 diamonds.

Answer

1. 635 013 559 600
2. 151 519 319 380
3. 47 079 732 700
4. 11 404 407 300

A hand of cards that has no cards in a particular suit is said to be *void* in that suit. Use the inclusion-exclusion formula to find each of the following:

1. The number of poker hands that are void in at least one suit.
2. The number of bridge hands that are void in at least one suit.

Answer

1. 1 913 496
2. 32 427 298 180

A bridge hand that has no *honor cards* (cards of denomination 10, jack, queen, king, or ace) is said to be a *Yarborough*, in honor of the Second Earl of Yarborough. Find the number of Yarboroughs.

Answer

347 373 600

A *bridge deal* consists of dealing 13 cards (a *bridge hand*) to each of 4 distinct players (generically referred to as *north*, *south*, *east*, and *west*) from a standard deck of 52 cards. Find the number of bridge deals.

Answer

$53\,644\,737\,765\,488\,792\,839\,237\,440\,000 \approx 5.36 \times 10^{28}$

This staggering number is about the same order of magnitude as the number of atoms in your body, and is one of the reasons that bridge is a rich and interesting game.

Find the number of permutations of the cards in a standard deck.

Answer

$52! \approx 8.0658 \times 10^{67}$

This number is even more staggering. Indeed if you perform the experiment of dealing all 52 cards from a well-shuffled deck, you may well generate a pattern of cards that has never been generated before, thereby ensuring your immortality. Actually, this experiment shows that, in a sense, rare events can be very common. By the way, Persi Diaconis has shown that it takes about seven standard riffle shuffles to thoroughly randomize a deck of cards.

Dice and Coins

Suppose that 5 distinct, standard dice are rolled and the sequence of scores recorded.

1. Find the number of sequences.
2. Find the number of sequences with the scores all different.

Answer

1. 7776
2. 720

Suppose that 5 identical, standard dice are rolled. How many outcomes are there?

Answer

252

A coin is tossed 10 times and the outcome is recorded as a bit string (where 1 denotes heads and 0 tails).

1. Find the number of outcomes.
2. Find the number of outcomes with exactly 4 heads.
3. Find the number of outcomes with at least 8 heads.

Answer

1. 1024
2. 210
3. 56

Polynomial Coefficients

Find the coefficient of $x^3 y^4$ in $(2x - 4y)^7$.

Answer

71 680

Find the coefficient of x^5 in $(2 + 3x)^8$.

Answer

108 864

Find the coefficient of $x^3 y^7 z^5$ in $(x + y + z)^{15}$.

Answer

360 360

The Galton Board

In the Galton board game,

1. Move the ball from $(0, 0)$ to $(10, 6)$ along a path of your choice. Note the corresponding bit string and subset.
2. Generate the bit string 0011101001. Note the corresponding subset and path.
3. Generate the subset $\{1, 4, 5, 7, 8, 10\}$. Note the corresponding bit string and path.
4. Generate all paths from $(0, 0)$ to $(5, 3)$. How many paths are there?

Answer

4. 10

Generate Pascal's triangle up to $n = 10$.

Samples

A shipment contains 12 good and 8 defective items. A sample of 5 items is selected. Find the number of samples that contain exactly 3 good items.

Answer

6160

In the (n, k) lottery, k numbers are chosen without replacement from the set of integers from 1 to n (where $n, k \in \mathbb{N}_+$ and $k < n$). Order does not matter.

1. Find the number of outcomes in the general (n, k) lottery.
2. Explicitly compute the number of outcomes in the $(44, 6)$ lottery (a common format).

Answer

1. $\binom{n}{k}$
2. 7 059 052

For more on this topic, see the section on [Lotteries](#) in the chapter on Games of Chance.

Explicitly compute each formula in the [sampling table](#) above when $n = 10$ and $k = 4$.

Answer

1. Ordered samples with replacement: 10 000
2. Ordered samples without replacement: 5040
3. Unordered samples with replacement: 715
4. Unordered samples without replacement: 210

Greetings

Suppose there are n people who shake hands with each other. How many handshakes are there?

Answer

$\binom{n}{2}$. Note that a handshake can be thought of as a subset of size 2 from the set of n people.

There are m men and n women. The men shake hands with each other; the women hug each other; and each man bows to each woman.

1. How many handshakes are there?
2. How many hugs are there?
3. How many bows are there?
4. How many greetings are there?

Answer

1. $\binom{m}{2}$
2. $\binom{n}{2}$
3. mn
4. $\binom{m}{2} + \binom{n}{2} + mn = \binom{m+n}{2}$

Integer Solutions

Find the number of integer solutions of $x_1 + x_2 + x_3 = 10$ in each of the following cases:

1. $x_i \geq 0$ for each i .

2. $x_i > 0$ for each i .

Answer

1. 66
2. 36

Generalized Coefficients

Compute each of the following:

1. $(-5)^{(3)}$
2. $(\frac{1}{2})^{(4)}$
3. $(-\frac{1}{3})^{(5)}$

Answer

1. -210
2. $-\frac{15}{16}$
3. $-\frac{3640}{243}$

Compute each of the following:

1. $(\frac{1/2}{3})$
2. $(\frac{-5}{4})$
3. $(\frac{-1/3}{5})$

Answer

1. $\frac{1}{16}$
2. 70
3. $-\frac{91}{729}$

Birthdays

Suppose that n persons are selected and their birthdays noted. (Ignore leap years, so that a year has 365 days.)

1. Find the number of outcomes.
2. Find the number of outcomes with distinct birthdays.

Answer

1. 365^n .
2. $365^{(n)}$.

Chess

Note that the squares of a chessboard are distinct, and in fact are often identified with the Cartesian product set

$$\{a, b, c, d, e, f, g, h\} \times \{1, 2, 3, 4, 5, 6, 7, 8\} \quad (1.8.29)$$

Find the number of ways of placing 8 rooks on a chessboard so that no rook can capture another in each of the following cases.

1. The rooks are distinguishable.
2. The rooks are indistinguishable.

Answer

1. 1 625 702 400
2. 40 320

Gifts

Suppose that 20 identical candies are distributed to 4 children. Find the number of distributions in each of the following cases:

1. There are no restrictions.
2. Each child must get at least one candy.

Answer

1. 1771
2. 969

In the song *The Twelve Days of Christmas*, find the number of gifts given to the singer by her true love. (Note that the singer starts afresh with gifts each day, so that for example, the true love gets a new partridge in a pear tree each of the 12 days.)

Answer

364

Teams

Suppose that 10 kids are divided into two teams of 5 each for a game of basketball. In how many ways can this be done in each of the following cases:

1. The teams are distinguishable (for example, one team is labeled “Alabama” and the other team is labeled “Auburn”).
2. The teams are not distinguishable.

Answer

1. 252
2. 126

This page titled [1.8: Combinatorial Structures](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.9: Topological Spaces

Topology is one of the major branches of mathematics, along with other such branches as algebra (in the broad sense of algebraic structures), and analysis. Topology deals with spatial concepts involving distance, closeness, separation, convergence, and continuity. Needless to say, entire series of books have been written about the subject. Our goal in this section and the next is simply to review the basic definitions and concepts of topology that we will need for our study of probability and stochastic processes. You may want to refer to this section as needed.

Basic Theory

Definitions

A *topological space* consists of a nonempty set S and a collection \mathcal{S} of subsets of S that satisfy the following properties:

1. $S \in \mathcal{S}$ and $\emptyset \in \mathcal{S}$
2. If $\mathcal{A} \subseteq \mathcal{S}$ then $\bigcup \mathcal{A} \in \mathcal{S}$
3. If $\mathcal{A} \subseteq \mathcal{S}$ and \mathcal{A} is finite, then $\bigcap \mathcal{A} \in \mathcal{S}$

If $A \in \mathcal{S}$, then A is said to be *open* and A^c is said to be *closed*. The collection \mathcal{S} of open sets is a *topology* on S .

So the union of an arbitrary number of open sets is still open, as is the intersection of a finite number of open sets. The universal set S and the empty set \emptyset are both open and closed. There may or may not exist other subsets of S with this property.

Suppose that S is a nonempty set, and that \mathcal{S} and \mathcal{T} are topologies on S . If $\mathcal{S} \subseteq \mathcal{T}$ then \mathcal{T} is *finer* than \mathcal{S} , and \mathcal{S} is *coarser* than \mathcal{T} . *Coarser than* defines a partial order on the collection of topologies on S . That is, if \mathcal{R} , \mathcal{S} , \mathcal{T} are topologies on S then

1. \mathcal{R} is coarser than \mathcal{R} , the *reflexive property*.
2. If \mathcal{R} is coarser than \mathcal{S} and \mathcal{S} is coarser than \mathcal{R} then $\mathcal{R} = \mathcal{S}$, the *anti-symmetric property*.
3. If \mathcal{R} is coarser than \mathcal{S} and \mathcal{S} is coarser than \mathcal{T} then \mathcal{R} is coarser than \mathcal{T} , the *transitive property*.

A topology can be characterized just as easily by means of closed sets as open sets.

Suppose that S is a nonempty set. A collection of subsets \mathcal{C} is the collection of closed sets for a topology on S if and only if

1. $S \in \mathcal{C}$ and $\emptyset \in \mathcal{C}$
2. If $\mathcal{A} \subseteq \mathcal{C}$ then $\bigcap \mathcal{A} \in \mathcal{C}$.
3. If $\mathcal{A} \subseteq \mathcal{C}$ and \mathcal{A} is a finite then $\bigcup \mathcal{A} \in \mathcal{C}$.

Proof

The set $\mathcal{S} = \{A^c : A \in \mathcal{C}\}$ must satisfy the axioms of a topology. So the result follows DeMorgan's laws: if \mathcal{A} is a collection of subsets of S then

$$\begin{aligned} \left(\bigcup \mathcal{A}\right)^c &= \bigcap \{A^c : A \in \mathcal{A}\} \\ \left(\bigcap \mathcal{A}\right)^c &= \bigcup \{A^c : A \in \mathcal{A}\} \end{aligned}$$

Suppose that (S, \mathcal{S}) is a topological space, and that $x \in S$. A set $A \subseteq S$ is a *neighborhood* of x if there exists $U \in \mathcal{S}$ with $x \in U \subseteq A$.

So a neighborhood of a point $x \in S$ is simply a set with an open subset that contains x . The idea is that points in a “small” neighborhood of x are “close” to x in a sense. An open set can be defined in terms of the neighborhoods of the points in the set.

Suppose again that (S, \mathcal{S}) is a topological space. A set $U \subseteq S$ is open if and only if U is a neighborhood of every $x \in U$

Proof

If U is open, then clearly U is a neighborhood of every point $x \in U$ and clearly satisfies the condition in the theorem. Conversely, suppose that U is a neighborhood of every $x \in U$. Then by definition of neighborhood, for every $x \in U$ there exists an open set U_x with $x \in U_x \subseteq U$. But then $\bigcup_{x \in U} U_x$ is open, and clearly this set is U .

Although the proof seems trivial, the neighborhood concept is how you should think of openness. A set U is open if every point in U has a set of “nearby points” that are also in U .

Our next three definitions deal with *topological* sets that are naturally associated with a given subset.

Suppose again that (S, \mathcal{S}) is a topological space and that $A \subseteq S$. The *closure* of A is the set

$$\text{cl}(A) = \bigcap \{B \subseteq S : B \text{ is closed and } A \subseteq B\} \quad (1.9.1)$$

This is the smallest closed set containing A :

1. $\text{cl}(A)$ is closed.
2. $A \subseteq \text{cl}(A)$.
3. If B is closed and $A \subseteq B$ then $\text{cl}(A) \subseteq B$

Proof

Note that $\mathcal{B} = \{B \subseteq S : B \text{ is closed and } A \subseteq B\}$ is nonempty since $S \in \mathcal{B}$.

1. The sets in \mathcal{B} are closed so $\bigcap \mathcal{B}$ is closed.
2. By definition, $A \subseteq B$ for each $B \in \mathcal{B}$. Hence $A \subseteq \bigcap \mathcal{B}$.
3. If B is closed and $A \subseteq B$ then $B \in \mathcal{B}$ so $\bigcap \mathcal{B} \subseteq B$.

Of course, if A is closed then $A = \text{cl}(A)$. Complementary to the closure of a set is the interior of the set.

Suppose again that (S, \mathcal{S}) is a topological space and that $A \subseteq S$. The *interior* of A is the set

$$\text{int}(A) = \bigcup \{U \subseteq S : U \text{ is open and } U \subseteq A\} \quad (1.9.2)$$

This set is the largest open subset of A :

1. $\text{int}(A)$ is open.
2. $\text{int}(A) \subseteq A$.
3. If U is open and $U \subseteq A$ then $U \subseteq \text{int}(A)$

Proof

Note that $\mathcal{U} = \{U \subseteq S : U \text{ is open and } U \subseteq A\}$ is nonempty since $\emptyset \in \mathcal{U}$.

1. The sets in \mathcal{U} are open so $\bigcup \mathcal{U}$ is open.
2. By definition, $U \subseteq A$ for each $U \in \mathcal{U}$. Hence $\bigcup \mathcal{U} \subseteq A$.
3. If U is open and $U \subseteq A$ then $U \in \mathcal{U}$ so $U \subseteq \bigcup \mathcal{U}$.

Of course, if A is open then $A = \text{int}(A)$. The boundary of a set is the set difference between the closure and the interior.

Suppose again that (S, \mathcal{S}) is a topological space. The *boundary* of A is $\partial(A) = \text{cl}(A) \setminus \text{int}(A)$. This set is closed.

Proof

By definition, $\partial(A) = \text{cl}(A) \cap [\text{int}(A)]^c$, the intersection of two closed sets.

A topology on a set induces a natural topology on any subset of the set.

Suppose that (S, \mathcal{S}) is a topological space and that R is a nonempty subset of S . Then $\mathcal{R} = \{A \cap R : A \in \mathcal{S}\}$ is a topology on R , known as the *relative topology* induced by \mathcal{S} .

Proof

First $S \in \mathcal{S}$ and $S \cap R = R$, so $R \in \mathcal{R}$. Next, $\emptyset \in \mathcal{S}$ and $\emptyset \cap R = \emptyset$ so $\emptyset \in \mathcal{R}$. Suppose that $\mathcal{B} \subseteq \mathcal{R}$. For each $B \in \mathcal{B}$, select $A \in \mathcal{S}$ such that $B = A \cap R$. Let \mathcal{A} denote the collection of sets selected (we need the axiom of choice to do this).

Then $\bigcup \mathcal{A} \in \mathcal{S}$ and $\bigcup \mathcal{B} = (\bigcup \mathcal{A}) \cap R$, so $\bigcup \mathcal{B} \in \mathcal{R}$. Finally, suppose that $\mathcal{B} \subseteq \mathcal{R}$ is finite. Once again, for each $B \in \mathcal{B}$ there exists $A \in \mathcal{S}$ with $A \cap R = B$. Let \mathcal{A} denote the collection of sets selected. Then \mathcal{A} is finite so $\bigcap \mathcal{A} \in \mathcal{S}$. But $\bigcap \mathcal{B} = (\bigcap \mathcal{A}) \cap R$ so $\bigcap \mathcal{B} \in \mathcal{R}$.

In the context of the previous result, note that if R is itself open, then the relative topology is $\mathcal{R} = \{A \in \mathcal{S} : A \subseteq R\}$, the subsets of R that are open in the original topology.

Separation Properties

Separation properties refer to the ability to separate points or sets with disjoint open sets. Our first definition deals with separating two points.

Suppose that (S, \mathcal{S}) is a topological space and that x, y are distinct points in S . Then x and y can be *separated* if there exist disjoint open sets U and V with $x \in U$ and $y \in V$. If every pair of distinct points in S can be separated, then (S, \mathcal{S}) is called a *Hausdorff space*.

Hausdorff spaces are named for the German mathematician Felix Hausdorff. There are weaker separation properties. For example, there could be an open set U that contains x but not y , and an open set V that contains y but not x , but no disjoint open sets that contain x and y . Clearly if every open set that contains one of the points also contains the other, then the points are indistinguishable from a topological viewpoint. In a Hausdorff space, singletons are closed.

Suppose that (S, \mathcal{S}) is a Hausdorff space. Then $\{x\}$ is closed for each $x \in S$.

Proof

The definition shows immediately that $\{x\}^c$ is open: if $y \in \{x\}^c$, there exists an open set V with $y \in V \subseteq \{x\}^c$.

Our next definition deals with separating a point from a closed set.

Suppose again that (S, \mathcal{S}) is a topological space. A nonempty closed set $A \subseteq S$ and a point $x \in A^c$ can be *separated* if there exist disjoint open sets U and V with $A \subseteq U$ and $x \in V$. If every nonempty closed set A and point $x \in A^c$ can be separated, then the space (S, \mathcal{S}) is *regular*.

Clearly if (S, \mathcal{S}) is a regular space and singleton sets are closed, then (S, \mathcal{S}) is a Hausdorff space.

Bases

Topologies, like other set structures, are often defined by first giving some basic sets that should belong to the collection, and then extending the collection so that the defining axioms are satisfied. This idea is motivation for the following definition:

Suppose again that (S, \mathcal{S}) is a topological space. A collection $\mathcal{B} \subseteq \mathcal{S}$ is a *base* for \mathcal{S} if every set in \mathcal{S} can be written as a union of sets in \mathcal{B} .

So, a base is a smaller collection of open sets with the property that every other open set can be written as a union of basic open sets. But again, we often want to start with the basic open sets and extend this collection to a topology. The following theorem gives the conditions under which this can be done.

Suppose that S is a nonempty set. A collection \mathcal{B} of subsets of S is a base for a topology on S if and only if

1. $S = \bigcup \mathcal{B}$
2. If $A, B \in \mathcal{B}$ and $x \in A \cap B$, there exists $C \in \mathcal{B}$ with $x \in C \subseteq A \cap B$

Proof

Suppose that \mathcal{B} is a base for a topology \mathcal{S} on S . Since S is open, S is a union of sets in \mathcal{B} . Since every set in \mathcal{B} is a subset of S , we must have $S = \bigcup \mathcal{B}$. Suppose that $A, B \in \mathcal{B}$ and that $x \in A \cap B$. Since $A \cap B$ is open, it's a union of sets in \mathcal{B} . The point x must be in one of those sets, so there exists $C \in \mathcal{B}$ with $x \in C \subseteq A \cap B$.

Suppose now that \mathcal{B} satisfies the two conditions in the theorem. Let \mathcal{S} be the collection of all unions of sets in \mathcal{B} . Then $S \in \mathcal{S}$ by condition (a), and $\emptyset \in \mathcal{S}$ by taking a vacuous union. Suppose that $U_i \in \mathcal{S}$ for $i \in I$ where I is an arbitrary index

set. Then for each $i \in I$, there exists an index set J_i such that $U_i = \bigcup_{j \in J_i} B_{i,j}$ where $B_{i,j} \in \mathcal{B}$ for each $j \in J_i$. But then

$$\bigcup_{i \in I} U_i = \bigcup_{i \in I} \bigcup_{j \in J_i} B_{i,j} \in \mathcal{S} \quad (1.9.3)$$

Finally, suppose that $U, V \in \mathcal{S}$. Then there exist index sets I and J with $U = \bigcup_{i \in I} A_i$ and $V = \bigcup_{j \in J} B_j$ where $A_i \in \mathcal{B}$ for all $i \in I$ and $B_j \in \mathcal{B}$ for all $j \in J$. Then

$$U \cap V = \bigcup_{i \in I, j \in J} (A_i \cap B_j) \quad (1.9.4)$$

By condition (b), for each $i \in I, j \in J$, and $x \in A_i \cap B_j$ there exists $C_{x,i,j} \in \mathcal{B}$ with $x \in C_{x,i,j} \subseteq A_i \cap B_j$. But then clearly

$$U \cap V = \bigcup \{C_{x,i,j} : i \in I, j \in J, x \in A_i \cap B_j\} \in \mathcal{S} \quad (1.9.5)$$

Here is a slightly weaker condition, but one that is often satisfied in practice.

Suppose that S is a nonempty set. A collection \mathcal{B} of subsets of S that satisfies the following properties is a base for a topology on S :

1. $S = \bigcup \mathcal{B}$
2. If $A, B \in \mathcal{B}$ then $A \cap B \in \mathcal{B}$

Part (b) means that \mathcal{B} is closed under finite intersections.

Compactness

Our next discussion considers another very important type of set. Some additional terminology will make the discussion easier. Suppose that S is a set and $A \subseteq S$. A collection of subsets \mathcal{A} of S is said to *cover* A if $A \subseteq \bigcup \mathcal{A}$. So the word *cover* simply means a collection of sets whose union contains a given set. In a topological space, we can have open an open cover (that is, a cover with open sets), a closed cover (that is, a cover with closed sets), and so forth.

Suppose again that (S, \mathcal{S}) is a topological space. A set $C \subseteq S$ is *compact* if every open cover of C has a finite sub-cover. That is, if $\mathcal{A} \subseteq \mathcal{S}$ with $C \subseteq \bigcup \mathcal{A}$ then there exists a finite $\mathcal{B} \subseteq \mathcal{A}$ with $C \subseteq \bigcup \mathcal{B}$.

So intuitively, a compact set is *compact* in the ordinary sense of the word. No matter how “small” are the open sets in the covering of C , there will always exist a finite number of the open sets that cover C .

Suppose again that (S, \mathcal{S}) is a topological space and that $C \subseteq S$ is a compact. If $B \subseteq C$ is closed, then B is also compact.

Proof

Suppose that \mathcal{A} is an open cover of B . Since B is closed, B^c is open, so $\mathcal{A} \cup \{B^c\}$ is an open cover of C . Since C is compact, this last collection has a finite sub-cover of C , which is also a finite sub-cover of B .

Compactness is also preserved under finite unions.

Suppose again that (S, \mathcal{S}) is a topological space, and that $C_i \subseteq S$ is compact for each i in a finite index set I . Then $C = \bigcup_{i \in I} C_i$ is compact.

Proof

Suppose that \mathcal{A} is an open cover of C . Then trivially, \mathcal{A} is also an open cover of C_i for each $i \in I$. Hence there exists a finite subcover $\mathcal{A}_i \subseteq \mathcal{A}$ of C_i for each $i \in I$. But then $\bigcup_{i \in I} \mathcal{A}_i$ is also finite and is a covering of C .

As we saw above, closed subsets of a compact set are themselves compact. In a Hausdorff space, a compact set is itself closed.

Suppose that (S, \mathcal{S}) is a Hausdorff space. If $C \subseteq S$ is compact then C is closed.

Proof

We will show that C^c is open, so fix $x \in C^c$. For each $y \in C$, the points x and y can be separated, so there exist disjoint open sets U_y and V_y such that $x \in U_y$ and $y \in V_y$. Trivially, the collection $\{V_y : y \in C\}$ is an open cover of C , and hence there exist a finite subset $B \subseteq C$ such that $\{V_y : y \in B\}$ covers C . But then $U = \bigcap_{y \in B} U_y$ is open and is disjoint from $\bigcup_{y \in B} V_y$. Hence also U is disjoint from C . So to summarize, U is open and $x \in U \subseteq C^c$.

Also in a Hausdorff space, a point can be separated from a compact set that does not contain the point.

Suppose that (S, \mathcal{S}) is a Hausdorff space. If $x \in S$, $C \subseteq S$ is compact, and $x \notin C$, then there exist disjoint open sets U and V with $x \in U$ and $C \subseteq V$.

Proof

Since the space is Hausdorff, for each $y \in C$ there exist disjoint open sets U_y and V_y with $x \in U_y$ and $y \in V_y$. The collection $\{V_y : y \in C\}$ is an open cover of C , and hence there exists a finite set $B \subseteq C$ such that $\{V_y : y \in B\}$ covers C . Thus let $U = \bigcap_{y \in B} U_y$ and $V = \bigcup_{y \in B} V_y$. Then U is open, since B is finite, and V is open. Moreover U and V are disjoint, and $x \in U$ and $C \subseteq V$.

In a Hausdorff space, if a point has a neighborhood with a compact boundary, then there is a smaller, closed neighborhood.

Suppose again that (S, \mathcal{S}) is a Hausdorff space. If $x \in S$ and A is a neighborhood of x with $\partial(A)$ compact, then there exists a closed neighborhood B of x with $B \subseteq A$.

Proof

By (20), there exist disjoint open sets U and V with $x \in U$ and $\partial(A) \subseteq V$. Hence $\text{cl}(U)$ and $\partial(A)$ are disjoint. Let $B = \text{cl}(A \cap U)$. Note that B is closed, and is a neighborhood of x since U and A are neighborhoods of x . Moreover,

$$B \subseteq \text{cl}(A) \cap \text{cl}(U) = [A \cup \partial(A)] \cap \text{cl}(U) = [A \cap \text{cl}(U)] \cup [\partial(A) \cap \text{cl}(U)] = A \cap \text{cl}(U) \subseteq A \quad (1.9.6)$$

Generally, *local* properties in a topological space refer to properties that hold on the neighborhoods of a point $x \in S$.

A topological space (S, \mathcal{S}) is *locally compact* if every point $x \in S$ has a compact neighborhood.

This definition is important because many of the topological spaces that occur in applications (like probability) are not compact, but are locally compact. Locally compact Hausdorff spaces have a number of nice properties. In particular, in a locally compact Hausdorff space, there are arbitrarily “small” compact neighborhoods of a point.

Suppose that (S, \mathcal{S}) is a locally compact Hausdorff space. If $x \in S$ and A is a neighborhood of x , then there exists a compact neighborhood B of x with $B \subseteq A$.

Proof

Since S is locally compact, there exists a compact neighborhood C of x . Hence $A \cap C$ is a neighborhood of x . Moreover, $\partial(A \cap C)$ is closed and is a subset of C and hence is compact. From (21), there exists a closed neighborhood B of x with $B \subseteq A \cap C$. Since B is closed and $B \subseteq C$, B is compact. Of course also, $B \subseteq A$.

Countability Axioms

Our next discussion concerns topologies that can be “countably constructed” in a certain sense. Such axioms limit the “size” of the topology in a way, and are often satisfied by important topological spaces that occur in applications. We start with an important preliminary definition.

Suppose that (S, \mathcal{S}) is a topological space. A set $D \subseteq S$ is *dense* if $U \cap D$ is nonempty for every nonempty $U \in \mathcal{S}$.

Equivalently, D is dense if every neighborhood of a point $x \in S$ contains an element of D . So in this sense, one can find elements of D “arbitrarily close” to a point $x \in S$. Of course, the entire space S is dense, but we are usually interested in topological spaces that have dense sets of limited cardinality.

Suppose again that (S, \mathcal{S}) is a topological space. A set $D \subseteq S$ is dense if and only if $\text{cl}(D) = S$.

Proof

Suppose that D is dense. Since $\text{cl}(D)$ is closed, $[\text{cl}(D)]^c$ is open. If this set is nonempty, it must contain a point in D . But that's clearly a contradiction since $D \subseteq \text{cl}(D)$. Conversely, suppose that $\text{cl}(D) = S$. Suppose that U is a nonempty, open set. Then U^c is closed, and $U^c \neq S$. If $D \cap U = \emptyset$, then $D \subseteq U^c$. But then $\text{cl}(D) \subseteq U^c$ so $\text{cl}(D) \neq S$.

Here is our first countability axiom:

A topological space (S, \mathcal{S}) is *separable* if there exists a countable dense subset.

So in a separable space, there is a *countable* set D with the property that there are points in D “arbitrarily close” to every $x \in S$. Unfortunately, the term *separable* is similar to *separating points* that we discussed above in the definition of a Hausdorff space. But clearly the concepts are very different. Here is another important countability axiom.

A topological space (S, \mathcal{S}) is *second countable* if it has a countable base.

So in a second countable space, there is a countable collection of open sets \mathcal{B} with the property that every other open set is a union of sets in \mathcal{B} . Here is how the two properties are related:

If a topological space (S, \mathcal{S}) is second countable then it is separable.

Proof

Suppose that $\mathcal{B} = \{U_i : i \in I\}$ is a base for \mathcal{S} , where I is a countable index set. Select $x_i \in U_i$ for each $i \in I$, and let $D = \{x_i : i \in I\}$. Of course, D is countable. If U is open and nonempty, then $U = \bigcup_{j \in J} U_j$ for some nonempty $J \subseteq I$. But then $\{x_j : j \in J\} \subseteq U$, so D is dense.

As the terminology suggests, there are other axioms of countability (such as *first countable*), but the two we have discussed are the most important.

Connected and Disconnected Spaces

This discussion deals with the situation in which a topological space falls into two or more separated pieces, in a sense.

A topological space (S, \mathcal{S}) is *disconnected* if there exist nonempty, disjoint, open sets U and V with $S = U \cup V$. If (S, \mathcal{S}) is not disconnected, then it is *connected*.

Since $U = V^c$, it follows that U and V are also closed. So the space is disconnected if and only if there exists a proper subset U that is open and closed (sadly, such sets are sometimes called *clopen*). If S is disconnected, then S consists of two pieces U and V , and the points in U are not “close” to the points in V , in a sense. To study S topologically, we could simply study U and V separately, with their relative topologies.

Convergence

There is a natural definition for a convergent sequence in a topological space, but the concept is not as useful as one might expect.

Suppose again that (S, \mathcal{S}) is a topological space. A sequence of points $(x_n : n \in \mathbb{N}_+)$ in S *converges* to $x \in S$ if for every neighborhood A of x there exists $m \in \mathbb{N}_+$ such that $x_n \in A$ for $n > m$. We write $x_n \rightarrow x$ as $n \rightarrow \infty$.

So for every neighborhood of x , regardless of how “small”, all but finitely many of the terms of the sequence will be in the neighborhood. One would naturally hope that limits, when they exist, are unique, but this will only be the case if points in the space can be separated.

Suppose that (S, \mathcal{S}) is a Hausdorff space. If $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in S with $x_n \rightarrow x \in S$ as $n \rightarrow \infty$ and $x_n \rightarrow y \in S$ as $n \rightarrow \infty$, then $x = y$.

Proof

If $x \neq y$, there exist disjoint neighborhoods A and B of x and y , respectively. There exist $k, m \in \mathbb{N}_+$ such that $x_n \in A$ for all $n > k$ and $x_n \in B$ for all $n > m$. But then if $n > \max\{k, m\}$, $x_n \in A$ and $x_n \in B$, a contradiction.

On the other hand, if distinct points $x, y \in S$ cannot be separated, then any sequence that converges to x will also converge to y .

Continuity

Continuity of functions is one of the most important concepts to come out of general topology. The idea, of course, is that if two points are close together in the domain, then the functional values should be close together in the range. The abstract topological definition, based on inverse images is very simple, but not very intuitive at first.

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces. A function $f : S \rightarrow T$ is *continuous* if $f^{-1}(A) \in \mathcal{S}$ for every $A \in \mathcal{T}$.

So a continuous function has the property that the inverse image of an open set (in the range space) is also open (in the domain space). Continuity can equivalently be expressed in terms of closed subsets.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces. A function $f : S \rightarrow T$ is continuous if and only if $f^{-1}(A)$ is a closed subset of S for every closed subset A of T .

Proof

Recall that $f^{-1}(A^c) = [f^{-1}(A)]^c$ for $A \subseteq T$. The result follows directly from the definition and the fact that a set is open if and only if its complement is closed.

Continuity preserves limits.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces, and that $f : S \rightarrow T$ is continuous. If $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in S with $x_n \rightarrow x \in S$ as $n \rightarrow \infty$, then $f(x_n) \rightarrow f(x)$ as $n \rightarrow \infty$.

Proof

Suppose that $V \subseteq T$ is open and $f(x) \in V$. Then $f^{-1}(V)$ is open in S and $x \in f^{-1}(V)$. Hence there exists $m \in \mathbb{N}_+$ such that $x_n \in f^{-1}(V)$ for every $n > m$. But then $f(x_n) \in V$ for $n > m$. So $f(x_n) \rightarrow f(x)$ as $n \rightarrow \infty$.

The converse of the last result is not true, so continuity of functions in a general topological space cannot be characterized in terms of convergent sequences. There are objects like sequences but more general, known as *nets*, that do characterize continuity, but we will not study these. Composition, the most important way to combine functions, preserves continuity.

Suppose that (S, \mathcal{S}) , (T, \mathcal{T}) , and (U, \mathcal{U}) are topological spaces. If $f : S \rightarrow T$ and $g : T \rightarrow U$ are continuous, then $g \circ f : S \rightarrow U$ is continuous.

Proof

If A is open in U then $g^{-1}(A)$ is open in T and therefore $f^{-1}[g^{-1}(A)] = (f^{-1} \circ g^{-1})(A)$ is open in S . But $(g \circ f)^{-1} = f^{-1} \circ g^{-1}$.

The next definition is very important. A recurring theme in mathematics is to recognize when two mathematical structures of a certain type are fundamentally the same, even though they may appear to be different.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces. A one-to-one function f that maps S onto T with both f and f^{-1} continuous is a *homeomorphism* from (S, \mathcal{S}) to (T, \mathcal{T}) . When such a function exists, the topological spaces are said to be *homeomorphic*.

Note that in this definition, f^{-1} refers to the inverse function, not the mapping of inverse images. If f is a homeomorphism, then A is open in S if and only if $f(A)$ is open in T . It follows that the topological spaces are essentially equivalent: any purely topological property can be characterized in terms of open sets and therefore any such property is shared by the two spaces.

Being homeomorphic is an equivalence relation on the collection of topological spaces. That is, for spaces (S, \mathcal{S}) , (T, \mathcal{T}) , and (U, \mathcal{U}) ,

1. (S, \mathcal{S}) is homeomorphic to (S, \mathcal{S}) (the *reflexive property*).
2. If (S, \mathcal{S}) is homeomorphic to (T, \mathcal{T}) then (T, \mathcal{T}) is homeomorphic to (S, \mathcal{S}) (the *symmetric property*).

3. If (S, \mathcal{S}) is homeomorphic to (T, \mathcal{T}) and (T, \mathcal{T}) is homeomorphic to (U, \mathcal{U}) then (S, \mathcal{S}) is homeomorphic to (U, \mathcal{U}) (the *transitive property*).

Proof

1. The identity function $I : S \rightarrow S$ defined by $I(x) = x$ for $x \in S$ is a homeomorphism from the space (S, \mathcal{S}) to itself.
2. If f is a homeomorphism from (S, \mathcal{S}) to (T, \mathcal{T}) then f^{-1} is a homeomorphism from (T, \mathcal{T}) to (S, \mathcal{S}) .
3. If f is a homeomorphism from (S, \mathcal{S}) to (T, \mathcal{T}) and g is a homeomorphism from (T, \mathcal{T}) to (U, \mathcal{U}) , then $g \circ f$ is a homeomorphism from (S, \mathcal{S}) to (U, \mathcal{U}) .

Continuity can also be defined *locally*, by restricting attention to the neighborhoods of a point.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces, and that $x \in S$. A function $f : S \rightarrow T$ is *continuous at x* if $f^{-1}(B)$ is a neighborhood of x in S whenever B is a neighborhood of $f(x)$ in T . If $A \subseteq S$, then f is *continuous on A* if f is continuous at each $x \in A$.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces, and that $f : S \rightarrow T$. Then f is continuous if and only if f is continuous at each $x \in S$.

Proof

Suppose that f is continuous. Let $x \in S$ and let B be a neighborhood of $f(x)$. Then there exists an open set V in T with $f(x) \in V \subseteq B$. But then $f^{-1}(V)$ is open in S , and $x \in f^{-1}(V) \subseteq f^{-1}(B)$, so $f^{-1}(B)$ is a neighborhood of x . Hence f is continuous at x .

Conversely, suppose that f is continuous at each $x \in S$, and suppose that $V \in \mathcal{T}$. If V contains no points in the range of f , then $f^{-1}(V) = \emptyset \in \mathcal{S}$. Otherwise, there exists $x \in S$ with $f(x) \in V$. But then V is a neighborhood of $f(x)$, so $U = f^{-1}(V)$ is a neighborhood of x . Let $y \in U$. Then $f(y) \in V$ also, so U is also a neighborhood of y . Hence $U \in \mathcal{S}$.

Properties that are defined for a topological space can be applied to a subset of the space, with the relative topology. But one has to be careful.

Suppose again that (S, \mathcal{S}) are topological spaces and that $f : S \rightarrow T$. Suppose also that $A \subseteq S$, and let \mathcal{A} denote the relative topology on A induced by \mathcal{S} , and let f_A denote the restriction of f to A . If f is continuous on A then f_A is continuous relative to the spaces (A, \mathcal{A}) and (T, \mathcal{T}) . The converse is not generally true.

Proof

Suppose that $V \in \mathcal{T}$. If $f(A) \cap V = \emptyset$ then $f_A^{-1}(V) = \emptyset \in \mathcal{A}$. Otherwise, suppose there exists $x \in A$ with $f(x) \in V$. Then V is a neighborhood of $f(x)$ in T so $f^{-1}(V)$ is a neighborhood of x in (S, \mathcal{S}) . Hence $f^{-1}(V) \cap A = f_A^{-1}(V)$ is a neighborhood of x in (A, \mathcal{A}) . Since f_A is continuous (relative to (A, \mathcal{A})) at each $x \in A$, f_A is continuous from the previous result.

For a simple counterexample, suppose that f is not continuous at a particular $x \in S$. The set $\{x\}$ has the trivial relative topology $\{\emptyset, \{x\}\}$ and so f restricted to $\{x\}$ is trivially continuous.

Product Spaces

Cartesian product sets are ubiquitous in mathematics, so a natural question is this: given topological spaces (S, \mathcal{S}) and (T, \mathcal{T}) , what is a natural topology for $S \times T$? The answer is very simple using the concept of a [base](#) above.

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces. The collection $\mathcal{B} = \{A \times B : A \in \mathcal{S}, B \in \mathcal{T}\}$ is a base for a topology on $S \times T$, called the *product topology* associated with the given spaces.

Proof

Trivially, $S \times T = \bigcup \mathcal{B}$. In fact $S \times T \in \mathcal{B}$. Next if $A \times B \in \mathcal{B}$ and $C \times D \in \mathcal{B}$, so that A, C are open in S and B, D are open in T , then

$$(A \times B) \cap (C \times D) = (A \cap C) \times (B \cap D) \in \mathcal{B} \quad (1.9.7)$$

Hence \mathcal{B} is a base for a topology on $S \times T$.

So basically, we want the product of open sets to be open in the product space. The product topology is the smallest topology that makes this happen. The definition above can be extended to very general product spaces, but to state the extension, let's recall how general product sets are constructed. Suppose that S_i is a set for each i in a nonempty index set I . Then the product set $\prod_{i \in I} S_i$ is the set of all functions $x : I \rightarrow \bigcup_{i \in I} S_i$ such that $x(i) \in S_i$ for $i \in I$.

Suppose that (S_i, \mathcal{S}_i) is a topological space for each i in a nonempty index set I . Then

$$\mathcal{B} = \left\{ \prod_{i \in I} A_i : A_i \in \mathcal{S}_i \text{ for all } i \in I \text{ and } A_i = S_i \text{ for all but finitely many } i \in I \right\} \quad (1.9.8)$$

is a base for a topology on $\prod_{i \in I} S_i$, known as the *product topology* associated with the given spaces.

Proof

The proof is just as before, except for the more complicated notation. Trivially $\prod_{i \in I} S_i = \bigcup \mathcal{B}$, and \mathcal{B} is closed under finite intersections.

Suppose again that S_i is a set for each i in a nonempty index set I . For $j \in I$, recall that *projection function* $p_j : \prod_{i \in I} S_i \rightarrow S_j$ is defined by $p_j(x) = x(j)$.

Suppose again that (S_i, \mathcal{S}_i) is a topological space for each $i \in I$, and give the product space $\prod_{i \in I} S_i$ the product topology. The projection function p_j is continuous for each $j \in I$.

Proof

If U is open in S_j then $p_j^{-1}(U) = \prod_{i \in I} A_i$ where $A_i = S_i$ for $i \in I$ with $i \neq j$, and $A_j = U$, so clearly this inverse image is open in the product space.

As a special case of all this, suppose that (S, \mathcal{S}) is a topological space, and that $S_i = S$ for all $i \in I$. Then the product space $\prod_{i \in I} S_i$ is the set of all functions from I to S , sometimes denoted S^I . In this case, the base for the product topology on S^I is

$$\mathcal{B} = \left\{ \prod_{i \in I} A_i : A_i \in \mathcal{S} \text{ for all } i \in I \text{ and } A_i = S \text{ for all but finitely many } i \in I \right\} \quad (1.9.9)$$

For $j \in I$, the projection function p_j just returns the value of a function $x : I \rightarrow S$ at j : $p_j(x) = x(j)$. This projection function is continuous. Note in particular that no topology is necessary on the domain I .

Examples and Special Cases

The Trivial Topology

Suppose that S is a nonempty set. Then $\{S, \emptyset\}$ is a topology on S , known as the *trivial topology*.

With the trivial topology, no two distinct points can be separated. So the topology cannot distinguish between points, in a sense, and all points in S are close to each other. Clearly, this topology is not very interesting, except as a place to start. Since there is only one nonempty open set (S itself), the space is connected, and every subset of S is compact. A sequence in S converges to every point in S .

Suppose that S has the trivial topology and that (T, \mathcal{T}) is another topological space.

1. Every function from T to S is continuous.
2. If (T, \mathcal{T}) is a Hausdorff space then the only continuous functions from S to T are constant functions.

Proof

1. Suppose $f : T \rightarrow S$. Then $f^{-1}(S) = T \in \mathcal{T}$ and $f^{-1}(\emptyset) = \emptyset \in \mathcal{T}$, so f is continuous.
2. Suppose that $f : S \rightarrow T$ is continuous and that u, v are distinct elements in the range of f . There exist disjoint open sets $U, V \in \mathcal{T}$ with $u \in U$ and $v \in V$. But $f^{-1}(U)$ and $f^{-1}(V)$ are nonempty and so must be S . If $x \in S$, $f(x) \in U$ and $f(x) \in V$, a contradiction.

The Discrete Topology

At the opposite extreme from the trivial topology, with the smallest collection of open sets, is the discrete topology, with the largest collection of open sets.

Suppose that S is a nonempty set. The power set $\mathcal{P}(S)$ (consisting of all subsets of S) is a topology, known as the *discrete topology*.

So in the discrete topology, every set is both open and closed. All points are separated, and in a sense, widely so. No point is close to another point. With the discrete topology, S is Hausdorff, disconnected, and the compact subsets are the finite subsets. A sequence in S converges to $x \in S$, if and only if all but finitely many terms of the sequence are x .

Suppose that S has the discrete topology and that (T, \mathcal{T}) is another topological space.

1. Every function from S to T is continuous.
2. If (T, \mathcal{T}) is connected, then the only continuous functions from T to S are constant functions.

Proof

1. Trivially, if $f : S \rightarrow T$, then $f^{-1}(U) \in \mathcal{P}(S)$ for $U \in \mathcal{T}$ so f is continuous.
2. Suppose that $f : T \rightarrow S$ is continuous and that x is in the range of f . Then $\{x\}$ is open and closed in S , so $f^{-1}\{x\}$ is open and closed in T . If T is connected, this means that $f^{-1}\{x\} = T$.

Euclidean Spaces

The standard topologies used in the Euclidean spaces are the topologies built from open sets that you familiar with.

For the set of real numbers \mathbb{R} , let $\mathcal{B} = \{(a, b) : a, b \in \mathbb{R}, a < b\}$, the collection of open intervals. Then \mathcal{B} is a base for a topology \mathcal{R} on \mathbb{R} , known as the *Euclidean topology*.

Proof

Clearly the conditions for \mathcal{B} to be a base given above are satisfied. First $\mathbb{R} = \bigcup \mathcal{B}$. Next, if $(a, b) \in \mathcal{B}$ and $(c, d) \in \mathcal{B}$ and $x \in (a, b) \cap (c, d)$, then $x \in (\max\{a, c\}, \min\{b, d\}) \subseteq (a, b) \cap (c, d)$.

The space $(\mathbb{R}, \mathcal{R})$ satisfies many properties that are motivations for definitions in topology in the first place. The convergence of a sequence in \mathbb{R} , in the topological sense given above, is the same as the definition of convergence in calculus. The same statement holds for the continuity of a function f from \mathbb{R} to \mathbb{R} .

Before listing other topological properties, we give a characterization of compact sets, known as the *Heine-Borel theorem*, named for Eduard Heine and Émile Borel. Recall that $A \subseteq \mathbb{R}$ is *bounded* if $A \subseteq [a, b]$ for some $a, b \in \mathbb{R}$ with $a < b$.

A subset $C \subseteq \mathbb{R}$ is compact if and only if C is closed and bounded.

So in particular, closed, bounded intervals of the form $[a, b]$ with $a, b \in \mathbb{R}$ and $a < b$ are compact.

The space $(\mathbb{R}, \mathcal{R})$ has the following properties:

1. Hausdorff.
2. Connected.
3. Locally compact.
4. Second countable.

Proof

1. Distinct points in \mathbb{R} can be separated by open intervals.
2. \mathbb{R} has no proper subset that is both open and closed.
3. If A is a neighborhood of $x \in \mathbb{R}$, then there exists $a, b \in \mathbb{R}$ with $a < b$ such that $x \in [a, b] \subseteq A$. The closed interval $[a, b]$ is compact.
4. The collection $\mathcal{Q} = \{(a, b) : a, b \in \mathbb{Q}, a < b\}$ is a countable base for \mathcal{R} , where as usual, \mathbb{Q} is the set of rational real numbers.

As noted in the proof, \mathbb{Q} , the set of rationals, is countable and is dense in \mathbb{R} . Another countable, dense subset is $\mathbb{D} = \{j/2^n : n \in \mathbb{N} \text{ and } j \in \mathbb{Z}\}$, the set of dyadic rationals (or *binary rationals*). For the higher-dimensional Euclidean spaces, we can use the product topology based on the topology of the real numbers.

For $n \in \{2, 3, \dots\}$, let $(\mathbb{R}^n, \mathcal{R}_n)$ be the n -fold product space corresponding to the space $(\mathbb{R}, \mathcal{R})$. Then \mathcal{R}_n is the *Euclidean topology* on \mathbb{R}^n .

A subset $A \subseteq \mathbb{R}^n$ is *bounded* if there exists $a, b \in \mathbb{R}$ with $a < b$ such that $A \subseteq [a, b]^n$, so that A fits inside of an n -dimensional “block”.

A subset $C \subseteq \mathbb{R}^n$ is compact if and only if C is closed and bounded.

The space $(\mathbb{R}^n, \mathcal{R}_n)$ has the following properties:

1. Hausdorff.
2. Connected.
3. Locally compact.
4. Second countable.

This page titled [1.9: Topological Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.10: Metric Spaces

Basic Theory

Most of the important topological spaces that occur in applications (like probability) have an additional structure that gives a *distance* between points in the space.

Definitions

A *metric space* consists of a nonempty set S and a function $d : S \times S \rightarrow [0, \infty)$ that satisfies the following axioms: For $x, y, z \in S$,

1. $d(x, y) = 0$ if and only if $x = y$.
2. $d(x, y) = d(y, x)$.
3. $d(x, z) \leq d(x, y) + d(y, z)$.

The function d is known as a *metric* or a *distance function*.

So as the name suggests, $d(x, y)$ is the *distance* between points $x, y \in S$. The axioms are intended to capture the essential properties of distance from geometry. Part (a) is the *positive property*; the distance is strictly positive if and only if the points are distinct. Part (b) is the *symmetric property*; the distance from x to y is the same as the distance from y to x . Part (c) is the *triangle inequality*; going from x to z cannot be longer than going from x to z by way of a third point y .

Note that if (S, d) is a metric space, and A is a nonempty subset of S , then the set A with d restricted to $A \times A$ is also a metric space (known as a *subspace*). The next definitions also come naturally from geometry:

Suppose that (S, d) is a metric space, and that $x \in S$ and $r \in (0, \infty)$.

1. $B(x, r) = \{y \in S : d(x, y) < r\}$ is the *open ball* with center x and radius r .
2. $C(x, r) = \{y \in S : d(x, y) \leq r\}$ is the *closed ball* with center x and radius r .

A metric on a space induces a topology on the space in a natural way.

Suppose that (S, d) is a metric space. By definition, a set $U \subseteq S$ is *open* if for every $x \in U$ there exists $r \in (0, \infty)$ such that $B(x, r) \subseteq U$. The collection \mathcal{S}_d of open subsets of S is a topology.

Proof

1. Trivially S is open and vacuously \emptyset is open.
2. Suppose that A_i is open for i in an arbitrary index set I , and let $A = \bigcup_{i \in I} A_i$. If $x \in A$ then $x \in A_i$ for some $i \in I$. Since A_i is open, there exists $r \in (0, \infty)$ with $B(x, r) \subseteq A_i$. But then $B(x, r) \subseteq A$ so A is open.
3. Suppose that A_i is open for i in a finite index set I , and let $A = \bigcap_{i \in I} A_i$. If $x \in A$ then $x \in A_i$ for every $i \in I$. Hence for each $i \in I$ there exist $r_i \in (0, \infty)$ such that $B(x, r_i) \subseteq A_i$. Let $r = \min\{r_i : i \in I\}$. Since I is finite, $r > 0$ and $B(x, r) \subseteq B(x, r_i) \subseteq A_i$ for each $i \in I$. Hence $B(x, r) \subseteq A$, so A is open.

As the names suggests, an open ball is in fact *open* and a closed ball is in fact *closed*.

Suppose again that (S, d) is a metric space, and that $x \in S$ and $r \in (0, \infty)$. Then

1. $B(x, r)$ is open.
2. $C(x, r)$ is closed.

Proof

1. Let $y \in B(x, r)$, and let $a = d(x, y)$, so that $a < r$. If $z \in B(y, r - a)$ then we have $d(x, y) = a$ and $d(y, z) < r - a$, so by the triangle inequality, $d(x, z) < a + (r - a) = r$. Hence $z \in B(x, r)$. Thus $B(y, r - a) \subseteq B(x, r)$. It follows that $B(x, r)$ is open.
2. We show that $U = [C(x, r)]^c$ is open. Suppose that $y \in U$, and let $a = d(x, y)$, so that $a > r$. Let $z \in B(y, a - r)$ and suppose that $z \in C(x, r)$, so that $d(z, x) \leq r$. By the triangle inequality again,

$$d(x, y) \leq d(x, z) + d(z, y) < r + (a - r) = a \quad (1.10.1)$$

a contradiction. Hence $z \in U$. So $B(y, a - r) \subseteq U$.

Recall that for a general topological space, a *neighborhood* of a point $x \in S$ is a set $A \subseteq S$ with the property that there exists an open set U with $x \in U \subseteq A$. It follows that in a metric space, $A \subseteq S$ is a neighborhood of x if and only if there exists $r > 0$ such that $B(x, r) \subseteq A$. In words, a neighborhood of a point must contain an open ball about that point.

It's easy to construct new metrics from ones that we already have. Here's one such result.

Suppose that S is a nonempty set, and that d, e are metrics on S , and $c \in (0, \infty)$. Then the following are also metrics on S :

1. cd
2. $d + e$

Proof

1. Recall that cd is the function defined by $(cd)(x, y) = cd(x, y)$ for $(x, y) \in S^2$. Since $c > 0$, it's easy to see that the axioms are satisfied.
2. Recall that $d + e$ is the function defined by $(d + e)(x, y) = d(x, y) + e(x, y)$ for $(x, y) \in S^2$. Again, it's easy to see that the axioms are satisfied.

Since a metric space produces a topological space, all of the definitions for general topological spaces apply to metric spaces as well. In particular, in a metric space, distinct points can always be separated.

A metric space (S, d) is a Hausdorff space.

Proof

Let x, y be distinct points in S . Then $r = d(x, y) > 0$. The sets $B(x, r/2)$ and $B(y, r/2)$ are open, and contain x and y , respectively. Suppose that $z \in B(x, r/2) \cap B(y, r/2)$. By the triangle inequality,

$$d(x, y) \leq d(x, z) + d(z, y) < \frac{r}{2} + \frac{r}{2} = r \quad (1.10.2)$$

a contradiction. Hence $B(x, r/2)$ and $B(y, r/2)$ are disjoint.

Metrizable Spaces

Again, every metric space is a topological space, but not conversely. A non-Hausdorff space, for example, cannot correspond to a metric space. We know there are such spaces; a set S with more than one point, and with the trivial topology $\mathcal{S} = \{S, \emptyset\}$ is non-Hausdorff.

Suppose that (S, \mathcal{S}) is a topological space. If there exists a metric d on S such that $\mathcal{S} = \mathcal{S}_d$, then (S, \mathcal{S}) is said to be *metrizable*.

It's easy to see that different metrics can induce the same topology. For example, if d is a metric and $c \in (0, \infty)$, then the metrics d and cd induce the same topology.

Let S be a nonempty set. Metrics d and e on S are *equivalent*, and we write $d \equiv e$, if $\mathcal{S}_d = \mathcal{S}_e$. The relation \equiv is an equivalence relation on the collection of metrics on S . That is, for metrics d, e, f on S ,

1. $d \equiv d$, the *reflexive property*.
2. If $d \equiv e$ then $e \equiv d$, the *symmetric property*.
3. If $d \equiv e$ and $e \equiv f$ then $d \equiv f$, the *transitive property*.

There is a simple condition that characterizes when the topology of one metric is finer than the topology of another metric, and then this in turn leads to a condition for equivalence of metrics.

Suppose again that S is a nonempty set and that d, e are metrics on S . Then \mathcal{S}_e is finer than \mathcal{S}_d if and only if every open ball relative to d contains an open ball relative to e .

Proof

Suppose that $\mathcal{S}_d \subseteq \mathcal{S}_e$ so that \mathcal{S}_e is finer than \mathcal{S}_d . If $x \in S$ and $a \in (0, \infty)$, then the open ball $B_d(x, a)$ centered at x of radius a for the metric d is in \mathcal{S}_d and hence in \mathcal{S}_e . Thus there exists $b \in (0, \infty)$ such that $B_e(x, b) \subseteq B_d(x, a)$. Conversely, suppose that the condition in the theorem holds and suppose that $U \in \mathcal{S}_d$. If $x \in U$ there exists $a \in (0, \infty)$ such that $B_d(x, a) \subseteq U$. Hence there exists $b \in (0, \infty)$ such that $B_e(x, b) \subseteq B_d(x, a) \subseteq U$. So $U \in \mathcal{S}_e$.

It follows that metrics d and e on S are equivalent if and only if every open ball relative to one of the metrics contains an open ball relative to the other metric.

So every metrizable topology on S corresponds to an equivalence class of metrics that produce that topology. Sometimes we want to know that a topological space is metrizable, because of the nice properties that it will have, but we don't really need to use a specific metric that generates the topology. At any rate, it's important to have conditions that are sufficient for a topological space to be metrizable. The most famous such result is the *Urysohn metrization theorem*, named for the Russian mathematician Pavel Uryshon:

Suppose that (S, \mathcal{S}) is a regular, second-countable, Hausdorff space. Then (S, \mathcal{S}) is metrizable.

Review of the terms

Recall that *regular* means that every closed set and point not in the set can be separated by disjoint open sets. As discussed earlier, *Hausdorff* means that any two distinct points can be separated by disjoint open sets. Finally, *second-countable* means that there is a countable base for the topology, that is, there is a countable collection of open sets with the property that every other open set is a union of sets in the collection.

Convergence

With a distance function, the convergence of a sequence can be characterized in a manner that is just like calculus. Recall that for a general topological space (S, \mathcal{S}) , if $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in S and $x \in S$, then $x_n \rightarrow x$ as $n \rightarrow \infty$ means that for every neighborhood U of x , there exists $m \in \mathbb{N}_+$ such that $x_n \in U$ for $n > m$.

Suppose that (S, d) is a metric space, and that $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in S and $x \in S$. Then $x_n \rightarrow x$ as $n \rightarrow \infty$ if and only if for every $\epsilon > 0$ there exists $m \in \mathbb{N}_+$ such that if $n > m$ then $d(x_n, x) < \epsilon$. Equivalently, $x_n \rightarrow x$ as $n \rightarrow \infty$ if and only if $d(x_n, x) \rightarrow 0$ as $n \rightarrow \infty$ (in the usual calculus sense).

Proof

Suppose that $x_n \rightarrow x$ as $n \rightarrow \infty$, and let $\epsilon > 0$. Then $B(x, \epsilon)$ is a neighborhood of x , so there exists $m \in \mathbb{N}_+$ such that $x_n \in B(x, \epsilon)$ for $n > m$, which is the condition in the theorem. Conversely, suppose that condition in the theorem holds, and let U be a neighborhood of x . Then there exists $\epsilon > 0$ such that $B(x, \epsilon) \subseteq U$. By assumption, there exists $m \in \mathbb{N}_+$ such that if $n > m$ then $x_n \in B(x, \epsilon) \subseteq U$.

So, no matter how tiny $\epsilon > 0$ may be, all but finitely many terms of the sequence are within ϵ distance of x . As one might hope, limits are unique.

Suppose again that (S, d) is a metric space. Suppose also that $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in S and that $x, y \in S$. If $x_n \rightarrow x$ as $n \rightarrow \infty$ and $x_n \rightarrow y$ as $n \rightarrow \infty$ then $x = y$.

Proof

This follows immediately since a metric space is a Hausdorff space, and the limit of a sequence in a Hausdorff space is unique. Here's a direct proof: Let $\epsilon > 0$. Then there exists $k \in \mathbb{N}_+$ such that $d(x_n, x) < \epsilon/2$ for $n > k$, and there exists $m \in \mathbb{N}_+$ such that $d(x_n, y) < \epsilon/2$ for $n > m$. Let $n > \max\{k, m\}$. By the triangle inequality,

$$d(x, y) \leq d(x, x_n) + d(x_n, y) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \quad (1.10.3)$$

So we have $d(x, y) < \epsilon$ for every $\epsilon > 0$ and hence $d(x, y) = 0$ and thus $x = y$.

Convergence of a sequence is a topological property, and so is preserved under equivalence of metrics.

Suppose that d, e are equivalent metrics on S , and that $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in S and $x \in S$. Then $x_n \rightarrow x$ as $n \rightarrow \infty$ relative to d if and only if $x_n \rightarrow x$ as $n \rightarrow \infty$ relative to e .

Closed subsets of a metric space have a simple characterization in terms of convergent sequences, and this characterization is more intuitive than the abstract axioms in a general topological space.

Suppose again that (S, d) is a metric space. Then $A \subseteq S$ is closed if and only if whenever a sequence of points in A converges, the limit is also in A .

Proof

Suppose that A is closed and that $(x_n : n \in \mathbb{N}_+)$ is a sequence of points in A with $x_n \rightarrow x \in S$ as $n \rightarrow \infty$. Suppose that $x \in A^c$. Since A^c is open, $x_n \in A^c$ for n sufficiently large, a contradiction. Hence $x \in A$. Conversely, suppose that A has the sequential closure property, but that A is not closed. Then A^c is not open. This means that there exists $x \in A^c$ with the property that every neighborhood of x has points in A . Specifically, for each $n \in \mathbb{N}_+$ there exists $x_n \in B(x, 1/n)$ with $x_n \in A$. But clearly $x_n \rightarrow x$ as $n \rightarrow \infty$, again a contradiction.

The following definition also shows up in standard calculus. The idea is to have a criterion for convergence of a sequence that does not require knowing *a-priori* the limit. But for metric spaces, this definition takes on added importance.

Suppose again that (S, d) is a metric space. A sequence of points $(x_n : n \in \mathbb{N}_+)$ in S is a *Cauchy sequence* if for every $\epsilon > 0$ there exist $k \in \mathbb{N}_+$ such that if $m, n \in \mathbb{N}_+$ with $m > k$ and $n > k$ then $d(x_m, x_n) < \epsilon$.

Cauchy sequences are named for the ubiquitous Augustin Cauchy. So for a Cauchy sequence, no matter how tiny $\epsilon > 0$ may be, all but finitely many terms of the sequence will be within ϵ distance of each other. A convergent sequence is always Cauchy.

Suppose again that (S, d) is a metric space. If a sequence of points $(x_n : n \in \mathbb{N}_+)$ in S converges, then the sequence is Cauchy.

Proof

By assumption, there exists $x \in S$ such that $x_n \rightarrow x$ as $n \rightarrow \infty$. Let $\epsilon > 0$. There exists $k \in \mathbb{N}_+$ such that if $n \in \mathbb{N}_+$ and $n > k$ then $d(x_n, x) < \epsilon/2$. Hence if $m, n \in \mathbb{N}_+$ with $m > k$ and $n > k$ then by the triangle inequality,

$$d(x_m, x_n) \leq d(x_m, x) + d(x, x_n) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \quad (1.10.4)$$

So the sequence is Cauchy.

Conversely, one might think that a Cauchy sequence should converge, but it's relatively trivial to create a situation where this is false. Suppose that (S, d) is a metric space, and that there is a point $x \in S$ that is the limit of a sequence of points in S that are all distinct from x . Then the space $T = S - \{x\}$ with the metric d restricted to $T \times T$ has a Cauchy sequence that does not converge. Essentially, we have created a “convergence hole”. So our next definition is very natural and very important.

Suppose again that (S, d) is metric space and that $A \subseteq S$. Then A is *complete* if every Cauchy sequence in A converges to a point in A .

Of course, completeness can be applied to the entire space S . Trivially, a complete set must be closed.

Suppose again that (S, d) is a metric space, and that $A \subseteq S$. If A is complete, then A is closed.

Proof

Suppose that $\mathbf{x} = (x_n : n \in \mathbb{N})$ is a sequence of points in A and that $x_n \rightarrow x \in S$ as $n \rightarrow \infty$. Then \mathbf{x} is a Cauchy sequence, and so by completeness, $x \in A$. Hence A is closed by (12).

Completeness is such a crucial property that it is often imposed as an assumption on metric spaces that occur in applications. Even though a Cauchy sequence may not converge, here is a partial result that will be useful later: if a Cauchy sequence has a convergent subsequence, then the sequence itself converges.

Suppose again the (S, d) is a metric space, and that $(x_n : n \in \mathbb{N}_+)$ is a Cauchy sequence in S . If there exists a subsequence $(x_{n_k} : k \in \mathbb{N}_+)$ such that $x_{n_k} \rightarrow x \in S$ as $k \rightarrow \infty$, then $x_n \rightarrow x$ as $n \rightarrow \infty$.

Proof

Recall that in the construction of a subsequence, the indices $(n_k : k \in \mathbb{N}_+)$ must be a strictly increasing sequence in \mathbb{N}_+ . In particular, $n_k \rightarrow \infty$ as $k \rightarrow \infty$. So let $\epsilon > 0$. From the hypotheses, there exists $j \in \mathbb{N}_+$ such that if $k > j$ then $d(x_{n_k}, x) < \epsilon/2$. There exists $N \in \mathbb{N}_+$ such that if $m > N$ and $p > N$ then $d(x_m, x_p) < \epsilon/2$. Now let $m > N$. Pick $k \in \mathbb{N}_+$ such that $k > j$ and $n_k > N$. By the triangle inequality,

$$d(x_m, x) \leq d(x_m, x_{n_k}) + d(x_{n_k}, x) \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \quad (1.10.5)$$

Continuity

In metric spaces, continuity of functions also has simple characterizations in terms of that are familiar from calculus. We start with local continuity. Recall that the general topological definition is that $f : S \rightarrow T$ is continuous at $x \in S$ if $f^{-1}(V)$ is a neighborhood of x in S for every open set V of $f(x)$ in T .

Suppose that (S, d) and (T, e) are metric spaces, and that $f : S \rightarrow T$. The continuity of f at $x \in S$ is equivalent to each of the following conditions:

1. If $(x_n : n \in \mathbb{N}_+)$ is a sequence in S with $x_n \rightarrow x$ as $n \rightarrow \infty$ then $f(x_n) \rightarrow f(x)$ as $n \rightarrow \infty$.
2. For every $\epsilon > 0$, there exists $\delta > 0$ such that if $y \in S$ and $d(x, y) < \delta$ then $e[f(y) - f(x)] < \epsilon$.

Proof

1. This condition is *sequential continuity* at x . Continuity at x implies sequential continuity at x for general topological spaces, and hence for metric spaces. Conversely, suppose that sequential continuity holds at $x \in S$, and let V be a neighborhood of $f(x)$ in T . If $U = f^{-1}(V)$ is not a neighborhood of x in S , then for every $n \in \mathbb{N}_+$, there exists $x_n \in B(x, 1/n)$ with $x_n \notin U$. But then clearly $x_n \rightarrow x$ as $n \rightarrow \infty$ but $f(x_n)$ does not converge to $f(x)$ as $n \rightarrow \infty$, a contradiction.
2. Suppose that f is continuous at x . For $\epsilon > 0$, $B_T[f(x), \epsilon]$ is a neighborhood of $f(x)$, and hence $U = f^{-1}(B_T[f(x), \epsilon])$ is a neighborhood of x . Hence there exists $\delta > 0$ such that $B_S(x, \delta) \subseteq U$. But this means that if $d(y, x) < \delta$ then $e[f(y), f(x)] < \epsilon$. Conversely suppose that the condition in (b) holds, and suppose that V is a neighborhood of $f(x)$. Then there exists $\epsilon > 0$ such that $B_T[f(x), \epsilon] \subseteq V$. By assumption, there exists $\delta > 0$ such that if $y \in B_S(x, \delta)$ then $f(y) \in B_T[f(x), \epsilon] \subseteq V$. This means that $f^{-1}(V)$ is a neighborhood of x .

More generally, recall that f continuous on $A \subseteq S$ means that f is continuous at each $x \in A$, and that f continuous means that f is continuous on S . So general continuity can be characterized in terms of sequential continuity and the ϵ - δ condition.

On a metric space, there are stronger versions of continuity.

Suppose again that (S, d) and (T, e) are metric spaces and that $f : S \rightarrow T$. Then f is *uniformly continuous* if for every $\epsilon > 0$ there exists $\delta > 0$ such that if $x, y \in S$ with $d(x, y) < \delta$ then $e[f(x), f(y)] < \epsilon$.

In the ϵ - δ formulation of ordinary point-wise continuity above, δ depends on the point x in addition to ϵ . With uniform continuity, there exists a δ depending only on ϵ that works *uniformly* in $x \in S$.

Suppose again that (S, d) and (T, e) are metric spaces, and that $f : S \rightarrow T$. If f is uniformly continuous then f is continuous.

Here is an even stronger version of continuity.

Suppose again that (S, d) and (T, e) are metric spaces, and that $f : S \rightarrow T$. Then f is *Hölder continuous* with exponent $\alpha \in (0, \infty)$ if there exists $C \in (0, \infty)$ such that $e[f(x), f(y)] \leq C[d(x, y)]^\alpha$ for all $x, y \in S$.

The definition is named for Otto Höder. The exponent α is more important than the constant C , which generally does not have a name. If $\alpha = 1$, f is said to be *Lipschitz continuous*, named for the German mathematician Rudolf Lipschitz.

Suppose again that (S, d) and (T, e) are metric spaces, and that $f : S \rightarrow T$. If f is Hölder continuous with exponent $\alpha > 0$ then f is uniformly continuous.

The case where $\alpha = 1$ and $C < 1$ is particularly important.

Suppose again that (S, d) and (T, e) are metric spaces. A function $f : S \rightarrow T$ is a *contraction* if there exists $C \in (0, 1)$ such that

$$e[f(x), f(y)] \leq Cd(x, y), \quad x, y \in S \quad (1.10.6)$$

So contractions shrink distance. By the result above, a contraction is uniformly continuous. Part of the importance of contraction maps is due to the famous *Banach fixed-point theorem*, named for Stefan Banach.

Suppose that (S, d) is a complete metric space and that $f : S \rightarrow S$ is a contraction. Then f has a unique *fixed point*. That is, there exists exactly one $x^* \in S$ with $f(x^*) = x^*$. Let $x_0 \in S$, and recursively define $x_n = f(x_{n-1})$ for $n \in \mathbb{N}_+$. Then $x_n \rightarrow x^*$ as $n \rightarrow \infty$.

Functions that preserve distance are particularly important. The term *isometry* means distance-preserving.

Suppose again that (S, d) and (T, e) are metric spaces, and that $f : S \rightarrow T$. Then f is an *isometry* if $e[f(x), f(y)] = d(x, y)$ for every $x, y \in S$.

Suppose again that (S, d) and (T, e) are metric spaces, and that $f : S \rightarrow T$. If f is an isometry, then f is one-to-one and Lipschitz continuous.

Proof

If $x, y \in S$ with $x \neq y$, then $e[f(x), f(y)] = d(x, y) > 0$, so $f(x) \neq f(y)$. Hence f is one-to-one. Directly from the definition, f is Hölder continuous with exponent $\alpha = 1$ and constant multiple $C = 1$.

In particular, an isometry f is uniformly continuous. If one metric space can be mapped isometrically *onto* another metric space, the spaces are essentially the same.

Metric spaces (S, d) and (T, e) are *isometric* if there exists an isometry f that maps S onto T . Isometry is an equivalence relation on metric spaces. That is, for metric spaces (S, d) , (T, e) , and (U, ρ) ,

1. (S, d) is isometric to (S, d) , the *reflexive property*.
2. If (S, d) is isometric to (T, e) then (T, e) is isometric to (S, d) , the *symmetric property*.
3. If (S, d) is isometric to (T, e) and (T, e) is isometric to (U, ρ) , then (S, d) is isometric to (U, ρ) , the *transitive property*.

Proof

1. The identity function $I : S \rightarrow S$ defined by $I(x) = x$ for $x \in S$ is an isometry from (S, d) onto (S, d) .
2. If f is an isometry from (S, d) onto (T, e) then f^{-1} is an isometry from (T, e) onto (S, d) .
3. If f is an isometry from (S, d) onto (T, e) and g is an isometry from (T, e) onto (U, ρ) , then $g \circ f$ is an isometry from (S, d) to (U, ρ) .

In particular, if metric spaces (S, d) and (T, e) are isometric, then as topological spaces, they are homeomorphic.

Compactness and Boundedness

In a metric space, various definitions related to a set being *bounded* are natural, and are related to the general concept of compactness.

Suppose again that (S, d) is a metric space, and that $A \subseteq S$. Then A is *bounded* if there exists $r \in (0, \infty)$ such that $d(x, y) \leq r$ for all $x, y \in A$. The *diameter* of A is

$$\text{diam}(A) = \inf\{r > 0 : d(x, y) < r \text{ for all } x, y \in A\} \quad (1.10.7)$$

Additional details

Recall that $\inf(\emptyset) = \infty$, so $\text{diam}(A) = \infty$ if A is unbounded. In the bounded case, note that if the distance between points in A is bounded by $r \in (0, \infty)$, then the distance is bounded by any $s \in [r, \infty)$. Hence the diameter definition makes sense.

So A is bounded if and only if $\text{diam}(A) < \infty$. Diameter is an increasing function relative to the subset partial order.

Suppose again that (S, d) is a metric space, and that $A \subseteq B \subseteq S$. Then $\text{diam}(A) \leq \text{diam}(B)$.

Our next definition is stronger, but first let's review some terminology that we used for general topological spaces: If S is a set, A a subset of S , and \mathcal{A} a collection of subsets of S , then \mathcal{A} is said to *cover* A if $A \subseteq \bigcup \mathcal{A}$. So with this terminology, we can talk about open covers, closed covers, finite covers, disjoint covers, and so on.

Suppose again that (S, d) is a metric space, and that $A \subseteq S$. Then A is *totally bounded* if for every $r > 0$ there is a finite cover of A with open balls of radius r .

Recall that for a general topological space, a set A is *compact* if every open cover of A has a finite subcover. So in a metric space, the term *precompact* is sometimes used instead of *totally bounded*: The set A is totally bounded if every cover of A with open balls of radius r has a finite subcover.

Suppose again that (S, d) is a metric space. If $A \subseteq S$ is totally bounded then A is bounded.

Proof

There exists a finite cover of A with open balls of radius 1. Let C denote the set of centers of the balls, and let $c = \max\{d(u, v) : u, v \in C\}$, the maximum distance between two centers. Since C is finite, $c < \infty$. Now let $x, y \in A$. Since the balls cover A , there exist $u, v \in C$ with $x \in B(u, 1)$ and $y \in B(v, 1)$. By the triangle inequality (what else?)

$$d(x, y) \leq d(x, u) + d(u, v) + d(v, y) \leq 2 + c \quad (1.10.8)$$

Hence A is bounded.

Since a metric space is a Hausdorff space, a compact subset of a metric space is closed. Compactness also has a simple characterization in terms of convergence of sequences.

Suppose again that (S, d) is a metric space. A subset $C \subseteq S$ is compact if and only if every sequence of points in C has a subsequence that converges to a point in C .

Proof

The condition in the theorem is known as *sequential compactness*, so we want to show that sequential compactness is equivalent to compactness. The proof is harder than most of the others in this section, but the proof presented here is the nicest I have found, and is due to Anton Schep.

Suppose that C is compact and that $\mathbf{x} = (x_n : n \in \mathbb{N}_+)$ is a sequence of points in C . Let $A = \{x_n : n \in \mathbb{N}_+\} \subseteq C$, the unordered set of distinct points in the sequence. If A is finite, then some element of $a \in A$ must occur infinitely many times in the sequence. In this case, we can construct a subsequence of \mathbf{x} all of whose terms are a , and so this subsequence trivially converges to $a \in C$. Suppose next that A is infinite. Since the space is Hausdorff, C is closed, and therefore $\text{cl}(A) \subseteq C$. Our next claim is that there exists $a \in \text{cl}(A)$ such that for every $r > 0$, the set $A \cap B(a, r)$ is infinite. If the claim is false, then for each $a \in \text{cl}(A)$ there exists $r_a > 0$ such that $A \cap B(a, r_a)$ is finite. It then follows that for each $a \in A$, there exists $\epsilon_a > 0$ such that $A \cap B(a, \epsilon_a) = \{a\}$. But then $\mathcal{U} = \{B(a, \epsilon_a) : a \in \text{cl}(A)\} \cup \{\text{cl}(A)^c\}$ is an open cover of C that has no finite subcover, a contradiction. So the claim is true and for some $a \in \text{cl}(A)$, the set $A \cap B(a, r)$ is infinite for each $r > 0$. We can construct a subsequence of \mathbf{x} that converges to $a \in C$.

Conversely, suppose that C is sequentially compact. If $\mathbf{x} = (x_n : n \in \mathbb{N}_+)$ is a Cauchy sequence in C , then by assumption, \mathbf{x} has a subsequence that converges to some $x \in C$. But then by (17) the sequence \mathbf{x} itself converges to x , so it follows that C is complete. We next show that C is totally bounded. Our goal is to show that C can be covered by a finite number of balls of an arbitrary radius $r > 0$. Pick $x_1 \in C$. If $C \subseteq B(x_1, r)$ then we are done. Otherwise, pick $x_2 \in C \setminus B(x_1, r)$. If $C \subseteq B(x_1, r) \cup B(x_2, r)$ then again we are done. Otherwise there exists $x_3 \in C \setminus [B(x_1, r) \cup B(x_2, r)]$. This process must terminate in a finite number of steps or otherwise we would have a sequence of points $(x_n : n \in \mathbb{N}_+)$ in C with the property that $d(x_n, x_m) \geq r$ for every $n, m \in \mathbb{N}_+$. Such a sequence does not have a convergent subsequence. Suppose now that \mathcal{U} is an open cover of C and let $c = \text{diam}(C)$. Then C can be covered by a finite number of closed balls of with centers in C and with radius $c/4$. It follows that at least one of these balls cannot be covered by a finite subcover from \mathcal{U} . Let C_1 denote the intersection of this ball with C . Then C_1 is closed and is sequentially compact with $\text{diam}(C_1) \leq c/4$. Repeating the argument, we generate a nested sequence of close sets (C_1, C_2, \dots) such that $\text{diam}(C_n) \leq c/2^n$, and with the property that C_n cannot be finitely covered by \mathcal{U} for each $n \in \mathbb{N}_+$. Pick $x_n \in C_n$ for each $n \in \mathbb{N}_+$. Then $\mathbf{x} = (x_n : n \in \mathbb{N}_+)$ is a Cauchy sequence in C and hence has a subsequence that converges to some $x \in C$. Then $x \in \bigcap_{n=1}^{\infty} C_n$ and since $\text{diam}(C_n) \rightarrow 0$ as $n \rightarrow \infty$ it follows that in fact, $\bigcap_{n=1}^{\infty} C_n = \{x\}$. Now, since \mathcal{U} covers C , there exists $U \in \mathcal{U}$ such that $x \in U$. Since U is open, there exists $r > 0$

such that $B(x, r) \subseteq U$. Now let $n \in \mathbb{N}_+$ be sufficiently large that $d(x, x_n) \leq r/2$ and $\text{diam}(C_n) < r/2$. Then $C_n \subseteq B(x, r) \subseteq U$, which contradicts the fact that C_n cannot be finitely covered by \mathcal{U} .

Hausdorff Measure and Dimension

Our last discussion is somewhat advanced, but is important for the study of certain random processes, particularly Brownian motion. The idea is to measure the “size” of a set in a metric space in a topological way, and then use this measure to define a type of “dimension”. We need a preliminary definition, using our convenient *cover* terminology. If (S, d) is a metric space, $A \subseteq S$, and $\delta \in (0, \infty)$, then a *countable δ cover* of A is a countable cover \mathcal{B} of A with the property that $\text{diam}(B) < \delta$ for each $B \in \mathcal{B}$.

Suppose again that (S, d) is a metric space and that $A \subseteq S$. For $\delta \in (0, \infty)$ and $k \in [0, \infty)$, define

$$H_\delta^k(A) = \inf \left\{ \sum_{B \in \mathcal{B}} [\text{diam}(B)]^k : \mathcal{B} \text{ is a countable } \delta \text{ cover of } A \right\} \quad (1.10.9)$$

The k -dimensional *Hausdorff measure* of A is

$$H^k(A) = \sup \{ H_\delta^k(A) : \delta > 0 \} = \lim_{\delta \downarrow 0} H_\delta^k(A) \quad (1.10.10)$$

Additional details

Note that if \mathcal{B} is a countable δ cover of A then it is also a countable ϵ cover of A for every $\epsilon > \delta$. This means that $H_\delta^k(A)$ is decreasing in $\delta \in (0, \infty)$ for fixed $k \in [0, \infty)$. Hence

$$\sup \{ H_\delta^k(A) : \delta > 0 \} = \lim_{\delta \downarrow 0} H_\delta^k(A) \quad (1.10.11)$$

Note that the k -dimensional Hausdorff measure is defined for every $k \in [0, \infty)$, not just nonnegative integers. Nonetheless, the integer dimensions are interesting. The 0-dimensional measure of A is the number of points in A . In Euclidean space, which we consider in (36), the measures of dimension 1, 2, and 3 are related to length, area, and volume, respectively.

Suppose again that (S, d) is a metric space and that $A \subseteq S$. The *Hausdorff dimension* of A is

$$\dim_H(A) = \inf \{ k \in [0, \infty) : H^k(A) = 0 \} \quad (1.10.12)$$

Of special interest, as before, is the case when $S = \mathbb{R}^n$ for some $n \in \mathbb{N}_+$ and d is the standard Euclidean distance, reviewed in (36). As you might guess, the Hausdorff dimension of a point is 0, the Hausdorff dimension of a “simple curve” is 1, the Hausdorff dimension of a “simple surface” is 2, and so on. But there are also sets with fractional Hausdorff dimension, and the stochastic process Brownian motion provides some fascinating examples. The graph of standard Brownian motion has Hausdorff dimension $3/2$ while the set of zeros has Hausdorff dimension $1/2$.

Examples and Special Cases

Normed Vector Spaces

A norm on a vector space generates a metric on the space in a very simple, natural way.

Suppose that $(S, +, \cdot)$ is a vector space, and that $\|\cdot\|$ is a norm on the space. Then d defined by $d(x, y) = \|y - x\|$ for $x, y \in S$ is a metric on S .

Proof

The metric axioms follow easily from the norm axioms.

1. The positive property for d follows since $\|x\| = 0$ if and only if $x = 0$.
2. The symmetric property for d follows since $\| -x \| = \|x\|$.
3. The triangle inequality for d follows from the triangle inequality for the norm: $\|x + y\| \leq \|x\| + \|y\|$.

On \mathbb{R}^n , we have a variety of norms, and hence a variety of metrics.

For $n \in \mathbb{N}_+$ and $k \in [1, \infty)$, the function d_k given below is a metric on \mathbb{R}^n :

$$d_k(\mathbf{x}, \mathbf{y}) = \left(\sum_{i=1}^n |x_i - y_i|^k \right)^{1/k}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n), \mathbf{y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n \quad (1.10.13)$$

Proof

This follows from the general result above, since $\|\cdot\|_k$ defined below is a norm on \mathbb{R}^n :

$$\|\mathbf{x}\|_k = \left(\sum_{i=1}^n |x_i|^k \right)^{1/k}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (1.10.14)$$

Of course the metric d_2 is *Euclidean distance*, named for Euclid of course. This is the most important one, in a practical sense because it's the usual one that we use in the real world, and in a mathematical sense because the associated norm corresponds to the standard inner product on \mathbb{R}^n given by

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i, \quad \mathbf{x} = (x_1, x_2, \dots, x_n), \mathbf{y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n \quad (1.10.15)$$

For $n \in \mathbb{N}_+$, the function d_∞ defined below is a metric on \mathbb{R}^n :

$$d_\infty(\mathbf{x}, \mathbf{y}) = \max\{|x_i - y_i| : i \in \{1, 2, \dots, n\}\}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (1.10.16)$$

Proof

This follows from the general result above, since $\|\cdot\|_\infty$ defined below is a norm on \mathbb{R}^n :

$$\|\mathbf{x}\|_\infty = \max\{|x_i| : i \in \{1, 2, \dots, n\}\}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (1.10.17)$$

To justify the notation, recall that $\|\mathbf{x}\|_k \rightarrow \|\mathbf{x}\|_\infty$ as $k \rightarrow \infty$ for $\mathbf{x} \in \mathbb{R}^n$, and hence $d_k(\mathbf{x}, \mathbf{y}) \rightarrow d_\infty(\mathbf{x}, \mathbf{y})$ as $k \rightarrow \infty$ for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.




Figure 1.10.1: From inside out, the boundaries of the unit balls centered at the origin in \mathbb{R}^2 for the metrics d_k with $k \in \{1, 3/4, 2, 3, \infty\}$.

Suppose now that S is a nonempty set. Recall that the collection \mathcal{V} of all functions $f : S \rightarrow \mathbb{R}$ is a vector space under the usual pointwise definition of addition and scalar multiplication. That is, if $f, g \in \mathcal{V}$ and $c \in \mathbb{R}$, then $f + g \in \mathcal{V}$ and $cf \in \mathcal{V}$ are defined by $(f + g)(x) = f(x) + g(x)$ and $(cf)(x) = cf(x)$ for $x \in S$. Recall further that the collection \mathcal{U} of *bounded* functions $f : S \rightarrow \mathbb{R}$ is a vector subspace of \mathcal{V} , and moreover, $\|\cdot\|$ defined by $\|f\| = \sup\{|f(x)| : x \in S\}$ is a norm on \mathcal{U} , known as the *supremum norm*. It follows that \mathcal{U} is a metric space with the metric d defined by

$$d(f, g) = \|f - g\| = \sup\{|f(x) - g(x)| : x \in S\} \quad (1.10.18)$$

Vector spaces of bounded, real-valued functions, with the supremum norm are very important in the study of probability and stochastic processes. Note that the supremum norm on \mathcal{U} generalizes the [maximum norm](#) on \mathbb{R}^n , since we can think of a point in \mathbb{R}^n as a function from $\{1, 2, \dots, n\}$ into \mathbb{R} . Later, as part of our discussion on integration with respect to a positive measure, we will see how to generalize the k norms on \mathbb{R}^n to spaces of functions.

Products of Metric Spaces

If we have a finite number of metric spaces, then we can combine the individual metrics, together with an norm on the vector space \mathbb{R}^n , to create a norm on the Cartesian product space.

Suppose $n \in \{2, 3, \dots\}$, and that (S_i, d_i) is a metric space for each $i \in \{1, 2, \dots, n\}$. Suppose also that $\|\cdot\|$ is a norm on \mathbb{R}^n . Then the function d given as follows is a metric on $S = S_1 \times S_2 \times \dots \times S_n$:

$$d(\mathbf{x}, \mathbf{y}) = \|(d_1(x_1, y_1), d_2(x_2, y_2), \dots, d_n(x_n, y_n))\|, \quad \mathbf{x} = (x_1, x_2, \dots, x_n), \mathbf{y} = (y_1, y_2, \dots, y_n) \in S \quad (1.10.19)$$

Proof

1. Note that $d(\mathbf{x}, \mathbf{y}) = 0$ if and only if $d_i(x_i, y_i) = 0$ for $i \in \{1, 2, \dots, n\}$ if and only if $x_i = y_i$ for $i \in \{1, 2, \dots, n\}$ if and only if $\mathbf{x} = \mathbf{y}$.
2. Since $d_i(x_i, y_i) = d_i(y_i, x_i)$ for $i \in \{1, 2, \dots, n\}$ we have $d(\mathbf{x}, \mathbf{y}) = d(\mathbf{y}, \mathbf{x})$.
3. The triangle inequality follows from the triangle inequality for each metric, and the triangle inequality for the norm.

Graphs

Recall that a *graph* (in the combinatorial sense) consists of a countable set S of *vertices* and a set $E \subseteq S \times S$ of *edges*. In this discussion, we assume that the graph is *undirected* in the sense that $(x, y) \in E$ if and only if $(y, x) \in E$, and has no loops so that $(x, x) \notin E$ for $x \in S$. Finally, recall that a *path* of length $n \in \mathbb{N}_+$ from $x \in S$ to $y \in S$ is a sequence $(x_0, x_1, \dots, x_n) \in S^{n+1}$ such that $x_0 = x$, $x_n = y$, and $(x_{i-1}, x_i) \in E$ for $i \in \{1, 2, \dots, n\}$. The graph is *connected* if there exists a path of finite length between any two distinct vertices in S . Such a graph has a natural metric:

Suppose that $G = (S, E)$ is a connected graph. Then d defined as follows is a metric on S : $d(x, x) = 0$ for $x \in S$, and $d(x, y)$ is the length of the shortest path from x to y for distinct $x, y \in S$.

Proof

1. The positive property follows from the definition: $d(x, y) = 0$ if and only if $x = y$.
2. The symmetric property follows since the graph is undirected: $d(x, y) = d(y, x)$ for all $x, y \in S$.
3. For the triangle inequality, suppose that $x, y, z \in S$, and that $m = d(x, y)$ and $n = d(y, z)$. Then there is a path of length m from x to y and a path of length n from y to z . Concatenating the paths produces a path of length $m + n$ from x to z . But $d(x, z)$ is the length of the shortest such path, so it follows that $d(x, z) \leq m + n$.

The Discrete Topology

Suppose that S is a nonempty set. Recall that the discrete topology on S is $\mathcal{P}(S)$, the power set of S , so that every subset of S is open (and closed). The discrete topology is [metrizable](#), and there are lots of metrics that generate this topology.

Suppose again that S is a nonempty set. A metric d on S with the property that there exists $c \in (0, \infty)$ such that $d(x, y) \geq c$ for distinct $x, y \in S$ generates the discrete topology.

Proof

Note that $B(x, c) = \{x\}$ for $x \in S$. Hence $\{x\}$ is open for each $x \in S$.

So any metric that is bounded from below (for distinct points) generates the discrete topology. It's easy to see that there are such metrics.

Suppose again that S is a nonempty set. The function d on $S \times S$ defined by $d(x, x) = 0$ for $x \in S$ and $d(x, y) = 1$ for distinct $x, y \in S$ is a metric on S , known as the *discrete metric*. This metric generates the discrete topology.

Proof

Clearly $d(x, y) = 0$ if and only if $x = y$, and $d(x, y) = d(y, x)$ for $x, y \in S$, so the positive and symmetric properties hold. For the triangle inequality, suppose $x, y, z \in S$. The inequality trivially holds if the points are not distinct. If the points are distinct, then $d(x, z) = 1$ and $d(x, y) + d(y, z) = 2$.

In probability applications, the discrete topology is often appropriate when S is countable. Note also that the discrete metric is the graph distance if S is made into the *complete graph*, so that (x, y) is an edge for every pair of distinct vertices $x, y \in S$.

This page titled [1.10: Metric Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.11: Measurable Spaces

In this section we discuss some topics from measure theory that are a bit more advanced than the topics in the early sections of this chapter. However, measure-theoretic ideas are essential for a deep understanding of probability, since probability is itself a measure. The most important of the definitions is the σ -algebra, a collection of subsets of a set with certain closure properties. Such collections play a fundamental role, even for applied probability, in encoding the state of information about a random experiment.

On the other hand, we won't be overly pedantic about measure-theoretic details in this text. Unless we say otherwise, we assume that all sets that appear are measurable (that is, members of the appropriate σ -algebras), and that all functions are measurable (relative to the appropriate σ -algebras).

Although this section is somewhat abstract, many of the proofs are straightforward. Be sure to try the proofs yourself before reading the ones in the text.

Algebras and σ -Algebras

Suppose that S is a set, playing the role of a universal set for a particular mathematical model. It is sometimes impossible to include *all* subsets of S in our model, particularly when S is uncountable. In a sense, the more sets that we include, the harder it is to have consistent theories. However, we almost always want the collection of admissible subsets to be *closed* under the basic set operations. This leads to some important definitions.

Algebras of Sets

Suppose that \mathcal{S} is a nonempty collection of subsets of S . Then \mathcal{S} is an *algebra* (or *field*) if it is closed under complement and union:

1. If $A \in \mathcal{S}$ then $A^c \in \mathcal{S}$.
2. If $A \in \mathcal{S}$ and $B \in \mathcal{S}$ then $A \cup B \in \mathcal{S}$.

If \mathcal{S} is an algebra of subsets of S then

1. $S \in \mathcal{S}$
2. $\emptyset \in \mathcal{S}$

Proof

1. Since \mathcal{S} is nonempty, there exists $A \in \mathcal{S}$. Hence $A^c \in \mathcal{S}$ so $S = A \cup A^c \in \mathcal{S}$.
2. $\emptyset = S^c \in \mathcal{S}$

Suppose that \mathcal{S} is an algebra of subsets of S and that $A_i \in \mathcal{S}$ for each i in a finite index set I .

1. $\bigcup_{i \in I} A_i \in \mathcal{S}$
2. $\bigcap_{i \in I} A_i \in \mathcal{S}$

Proof

1. This follows by induction on the number of elements in I .
2. This follows from (a) and DeMorgan's law. If $A_i \in \mathcal{S}$ for $i \in I$ then $A_i^c \in \mathcal{S}$ for $i \in I$. Therefore $\bigcup_{i \in I} A_i^c \in \mathcal{S}$ and hence $\bigcap_{i \in I} A_i = \left(\bigcup_{i \in I} A_i^c\right)^c \in \mathcal{S}$.

Thus it follows that an algebra of sets is closed under a finite number of set operations. That is, if we start with a finite number of sets in the algebra \mathcal{S} , and build a new set with a finite number of set operations (union, intersection, complement), then the new set is also in \mathcal{S} . However in many mathematical theories, probability in particular, this is not sufficient; we often need the collection of admissible subsets to be closed under a *countable* number of set operations.

σ -Algebras of Sets

Suppose that \mathcal{S} is a nonempty collection of subsets of S . Then \mathcal{S} is a σ -algebra (or σ -field) if the following axioms are satisfied:

1. If $A \in \mathcal{S}$ then $A^c \in \mathcal{S}$.
2. If $A_i \in \mathcal{S}$ for each i in a countable index set I , then $\bigcup_{i \in I} A_i \in \mathcal{S}$.

Clearly a σ -algebra of subsets is also an algebra of subsets, so the basic results for [algebras](#) above still hold. In particular, $S \in \mathcal{S}$ and $\emptyset \in \mathcal{S}$.

If $A_i \in \mathcal{S}$ for each i in a countable index set I , then $\bigcap_{i \in I} A_i \in \mathcal{S}$.

Proof

The proof is just like the one [above](#) for algebras. If $A_i \in \mathcal{S}$ for $i \in I$ then $A_i^c \in \mathcal{S}$ for $i \in I$. Therefore $\bigcup_{i \in I} A_i^c \in \mathcal{S}$ and hence $\bigcap_{i \in I} A_i = (\bigcup_{i \in I} A_i^c)^c \in \mathcal{S}$.

Thus a σ -algebra of subsets of S is closed under countable unions and intersections. This is the reason for the symbol σ in the name. As mentioned in the introductory paragraph, σ -algebras are of fundamental importance in mathematics generally and probability theory specifically, and thus deserve a special definition:

If S is a set and \mathcal{S} a σ -algebra of subsets of S , then the pair (S, \mathcal{S}) is called a *measurable space*.

The term *measurable space* will make more sense in the next chapter, when we discuss positive measures (and in particular, probability measures) on such spaces.

Suppose that S is a set and that \mathcal{S} is a finite algebra of subsets of S . Then \mathcal{S} is also a σ -algebra.

Proof

Any countable union of sets in \mathcal{S} reduces to a finite union.

However, there are algebras that are not σ -algebras. Here is the classic example:

Suppose that S is an infinite set. The collection of *finite and co-finite* subsets of S defined below is an algebra of subsets of S , but not a σ -algebra:

$$\mathcal{F} = \{A \subseteq S : A \text{ is finite or } A^c \text{ is finite}\} \quad (1.11.1)$$

Proof

$S \in \mathcal{F}$ since $S^c = \emptyset$ is finite. If $A \in \mathcal{F}$ then $A^c \in \mathcal{F}$ by the symmetry of the definition. Suppose that $A, B \in \mathcal{F}$. If A and B are both finite then $A \cup B$ is finite. If A^c or B^c is finite, then $(A \cup B)^c = A^c \cap B^c$ is finite. In either case, $A \cup B \in \mathcal{F}$. Thus \mathcal{F} is an algebra of subsets of S .

Since S is infinite, it contains a countably infinite subset $\{x_0, x_1, x_2, \dots\}$. Let $A_n = \{x_{2n}\}$ for $n \in \mathbb{N}$. Then A_n is finite, so $A_n \in \mathcal{F}$ for each $n \in \mathbb{N}$. Let $E = \bigcup_{n=0}^{\infty} A_n = \{x_0, x_2, x_4, \dots\}$. Then E is infinite by construction. Also $\{x_1, x_3, x_5, \dots\} \subseteq E^c$, so E^c is infinite as well. Hence $E \notin \mathcal{F}$ and so \mathcal{F} is not a σ -algebra.

General Constructions

Recall that $\mathcal{P}(S)$ denotes the collection of *all* subsets of S , called the *power set* of S . Trivially, $\mathcal{P}(S)$ is the largest σ -algebra of S . The power set is often the appropriate σ -algebra if S is countable, but as noted above, is sometimes too large to be useful if S is uncountable. At the other extreme, the smallest σ -algebra of S is given in the following result:

The collection $\{\emptyset, S\}$ is a σ -algebra.

Proof

Clearly $\{\emptyset, S\}$ is a finite algebra: S and \emptyset are complements of each other, and $S \cup \emptyset = S$. Hence $\{S, \emptyset\}$ is a σ -algebra by the [result above](#).

In many cases, we want to construct a σ -algebra that contains certain basic sets. The next two results show how to do this.

Suppose that \mathcal{S}_i is a σ -algebra of subsets of S for each i in a nonempty index set I . Then $\mathcal{S} = \bigcap_{i \in I} \mathcal{S}_i$ is also a σ -algebra of subsets of S .

Proof

The proof is completely straightforward. First, $S \in \mathcal{S}_i$ for each $i \in I$ so $S \in \mathcal{S}$. If $A \in \mathcal{S}$ then $A \in \mathcal{S}_i$ for each $i \in I$ and hence $A^c \in \mathcal{S}_i$ for each $i \in I$. Therefore $A^c \in \mathcal{S}$. Finally suppose that $A_j \in \mathcal{S}$ for each j in a countable index set J . Then $A_j \in \mathcal{S}_i$ for each $i \in I$ and $j \in J$ and therefore $\bigcup_{j \in J} A_j \in \mathcal{S}_i$ for each $i \in I$. It follows that $\bigcup_{j \in J} A_j \in \mathcal{S}$.

Note that no restrictions are placed on the index set I , other than it be nonempty, so in particular it may well be uncountable.

Suppose that S is a set and that \mathcal{B} is a collection of subsets of S . The σ -algebra generated by \mathcal{B} is

$$\sigma(\mathcal{B}) = \bigcap \{ \mathcal{S} : \mathcal{S} \text{ is a } \sigma\text{-algebra of subsets of } S \text{ and } \mathcal{B} \subseteq \mathcal{S} \} \quad (1.11.2)$$

If \mathcal{B} is countable then $\sigma(\mathcal{B})$ is said to be *countably generated*.

So the σ -algebra generated by \mathcal{B} is the intersection of all σ -algebras that contain \mathcal{B} , which by the previous result really is a σ -algebra. Note that the collection of σ -algebras in the intersection is not empty, since $\mathcal{P}(S)$ is in the collection. Think of the sets in \mathcal{B} as *basic sets* that we want to be measurable, but do not form a σ -algebra.

The σ -algebra $\sigma(\mathcal{B})$ is the smallest σ algebra containing \mathcal{B} .

1. $\mathcal{B} \subseteq \sigma(\mathcal{B})$
2. If \mathcal{S} is a σ -algebra of subsets of S and $\mathcal{B} \subseteq \mathcal{S}$ then $\sigma(\mathcal{B}) \subseteq \mathcal{S}$.

Proof

Both of these properties follows from the definition of $\sigma(\mathcal{B})$ as the intersection of all σ -algebras that contain \mathcal{B} .

Note that the conditions in the last theorem completely characterize $\sigma(\mathcal{B})$. If \mathcal{S}_1 and \mathcal{S}_2 satisfy the conditions, then by (a), $\mathcal{B} \subseteq \mathcal{S}_1$ and $\mathcal{B} \subseteq \mathcal{S}_2$. But then by (b), $\mathcal{S}_1 \subseteq \mathcal{S}_2$ and $\mathcal{S}_2 \subseteq \mathcal{S}_1$.

If A is a subset of S then $\sigma\{A\} = \{\emptyset, A, A^c, S\}$

Proof

Let $\mathcal{S} = \{\emptyset, A, A^c, S\}$. Clearly \mathcal{S} is an algebra: A and A^c are complements of each other, as are \emptyset and S . Also,

$$\begin{aligned} A \cup A^c &= A \cup S = A^c \cup S = S \cup S = \emptyset \cup S = S \\ A \cup \emptyset &= A \cup A = A \\ A^c \cup \emptyset &= A^c \cup A^c = A^c \\ \emptyset \cup \emptyset &= \emptyset \end{aligned}$$

Since \mathcal{S} is finite, it is a σ -algebra by (7). Next, $A \in \mathcal{S}$. Conversely, if \mathcal{T} is a σ -algebra and $A \in \mathcal{T}$ then of course $\emptyset, S, A^c \in \mathcal{T}$ so $\mathcal{S} \subseteq \mathcal{T}$. Hence $\mathcal{S} = \sigma\{A\}$

We can generalize the previous result. Recall that a collection of subsets $\mathcal{A} = \{A_i : i \in I\}$ is a *partition* of S if $A_i \cap A_j = \emptyset$ for $i, j \in I$ with $i \neq j$, and $\bigcup_{i \in I} A_i = S$.

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of S into nonempty subsets. Then $\sigma(\mathcal{A})$ is the collection of all unions of sets in \mathcal{A} . That is,

$$\sigma(\mathcal{A}) = \left\{ \bigcup_{j \in J} A_j : J \subseteq I \right\} \quad (1.11.3)$$

Proof

Let $\mathcal{S} = \left\{ \bigcup_{j \in J} A_j : J \subseteq I \right\}$. Note that $S \in \mathcal{S}$ since $S = \bigcup_{i \in I} A_i$. Next, suppose that $B \in \mathcal{S}$. Then $B = \bigcup_{j \in J} A_j$ for some $J \subseteq I$. But then $B^c = \bigcup_{j \in J^c} A_j$, so $B^c \in \mathcal{S}$. Next, suppose that $B_k \in \mathcal{S}$ for $k \in K$ where K is a countable index set. Then for each $k \in K$ there exists $J_k \subseteq I$ such that $B_k = \bigcup_{j \in J_k} A_j$. But then $\bigcup_{k \in K} B_k = \bigcup_{k \in K} \bigcup_{j \in J_k} A_j = \bigcup_{j \in J} A_j$ where $J = \bigcup_{k \in K} J_k$. Hence $\bigcup_{k \in K} B_k \in \mathcal{S}$. Therefore \mathcal{S} is a σ -algebra of subsets of S . Trivially, $\mathcal{A} \subseteq \mathcal{S}$. If \mathcal{T} is a σ -algebra of subsets of S and $\mathcal{A} \subseteq \mathcal{T}$, then clearly $\bigcup_{j \in J} A_j \in \mathcal{T}$ for every $J \subseteq I$. Hence $\mathcal{S} \subseteq \mathcal{T}$.

A σ -algebra of this form is said to be *generated by a countable partition*. Note that since $A_i \neq \emptyset$ for $i \in I$, the representation of a set in $\sigma(\mathcal{A})$ as a union of sets in \mathcal{A} is unique. That is, if $J, K \subseteq I$ and $J \neq K$ then $\bigcup_{j \in J} A_j \neq \bigcup_{k \in K} A_k$. In particular, if there are n nonempty sets in \mathcal{A} , so that $\#(I) = n$, then there are 2^n subsets of I and hence 2^n sets in $\sigma(\mathcal{A})$.

Suppose now that $\mathcal{A} = \{A_1, A_2, \dots, A_n\}$ is a collection of n subsets of S (not necessarily disjoint). To describe the σ -algebra generated by \mathcal{A} we need a bit more notation. For $x = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ (a bit string of length n), let $B_x = \bigcap_{i=1}^n A_i^{x_i}$ where $A_i^1 = A_i$ and $A_i^0 = A_i^c$.

In the setting above,

1. $\mathcal{B} = \{B_x : x \in \{0, 1\}^n\}$ partitions S .
2. $A_i = \bigcup \{B_x : x \in \{0, 1\}^n, x_i = 1\}$ for $i \in \{1, 2, \dots, n\}$.
3. $\sigma(\mathcal{A}) = \sigma(\mathcal{B}) = \{\bigcup_{x \in J} B_x : J \subseteq \{0, 1\}^n\}$.

Proof

1. Suppose that $x, y \in \{0, 1\}^n$ and that $x \neq y$. Without loss of generality we can suppose that for some $j \in \{1, 2, \dots, n\}$, $x_j = 0$ while $y_j = 1$. Then $B_x \subseteq A_j^c$ and $B_y \subseteq A_j$ so B_x and B_y are disjoint. Suppose that $s \in S$. Construct $x \in \{0, 1\}^n$ by $x_i = 1$ if $s \in A_i$ and $x_i = 0$ if $s \notin A_i$, for each $i \in \{1, 2, \dots, n\}$. Then by definition, $s \in B_x$. Hence \mathcal{B} partitions S .
2. Fix $i \in \{1, 2, \dots, n\}$. Again if $x \in \{0, 1\}^n$ and $x_i = 1$ then $B_x \subseteq A_i$. Hence $\bigcup \{B_x : x \in \{0, 1\}^n, x_i = 1\} \subseteq A_i$. Conversely, suppose $s \in A_i$. Define $y \in \{0, 1\}^n$ by $y_j = 1$ if $s \in A_j$ and $y_j = 0$ if $s \notin A_j$ for each $j \in \{1, 2, \dots, n\}$. Then $y_i = 1$ and $s \in B_y$. Hence $s \in \bigcup \{B_x : x \in \{0, 1\}^n, x_i = 1\}$.
3. Clearly, every σ -algebra of subsets of S that contains \mathcal{A} must also contain \mathcal{B} , and every σ -algebra of subsets of S that contains \mathcal{B} must also contain \mathcal{A} . It follows that $\sigma(\mathcal{A}) = \sigma(\mathcal{B})$. The characterization in terms of unions now follows from the [previous result](#).

Recall that there are 2^n bit strings of length n . The sets in \mathcal{A} are said to be in *general position* if the sets in \mathcal{B} are distinct (and hence there are 2^n of them) and are nonempty. In this case, there are 2^{2^n} sets in $\sigma(\mathcal{A})$.

Open the Venn diagram app. This app shows two subsets A and B of S in general position, and lists the 16 sets in $\sigma\{A, B\}$.

1. Select each of the 4 sets that partition S : $A \cap B$, $A \cap B^c$, $A^c \cap B$, $A^c \cap B^c$.
2. Select each of the other 12 sets in $\sigma\{A, B\}$ and note how each is a union of some of the sets in (a).

Sketch a Venn diagram with sets A_1, A_2, A_3 in general position. Identify the set B_x for each $x \in \{0, 1\}^3$.

If a σ -algebra is generated by a collection of basic sets, then each set in the σ -algebra is generated by a countable number of the basic sets.

Suppose that S is a set and \mathcal{B} a nonempty collection of subsets of S . Then

$$\sigma(\mathcal{B}) = \{A \subseteq S : A \in \sigma(\mathcal{C}) \text{ for some countable } \mathcal{C} \subseteq \mathcal{B}\} \quad (1.11.4)$$

Proof

Let \mathcal{S} denote the collection on the right. We first show that \mathcal{S} is a σ -algebra. First, pick $B \in \mathcal{B}$, which we can do since \mathcal{B} is nonempty. Then $S \in \sigma\{B\}$ so $S \in \mathcal{S}$. Let $A \in \mathcal{S}$ so that $A \in \sigma(\mathcal{C})$ for some countable $\mathcal{C} \subseteq \mathcal{B}$. Then $A^c \in \sigma(\mathcal{C})$ so $A^c \in \mathcal{S}$. Finally, suppose that $A_i \in \mathcal{S}$ for i in a countable index set I . Then for each $i \in I$, there exists a countable $\mathcal{C}_i \subseteq \mathcal{B}$ such that $A_i \in \sigma(\mathcal{C}_i)$. But then $\bigcup_{i \in I} \mathcal{C}_i$ is also countable and $\bigcup_{i \in I} A_i \in \sigma(\bigcup_{i \in I} \mathcal{C}_i)$. Hence $\bigcup_{i \in I} A_i \in \mathcal{S}$.

Next if $B \in \mathcal{B}$ then $B \in \sigma\{B\}$ so $B \in \mathcal{S}$. Hence $\sigma(\mathcal{B}) \subseteq \mathcal{S}$. Conversely, if $A \in \sigma(\mathcal{C})$ for some countable $\mathcal{C} \subseteq \mathcal{B}$ then trivially $A \in \sigma(\mathcal{B})$.

A σ -algebra on a set naturally leads to a σ -algebra on a subset.

Suppose that (S, \mathcal{S}) is a measurable space, and that $R \subseteq S$. Let $\mathcal{R} = \{A \cap R : A \in \mathcal{S}\}$. Then

1. \mathcal{R} is a σ -algebra of subsets of R .
2. If $R \in \mathcal{S}$ then $\mathcal{R} = \{B \in \mathcal{S} : B \subseteq R\}$.

Proof

1. First, $S \in \mathcal{S}$ and $S \cap R = R$ so $R \in \mathcal{R}$. Next suppose that $B \in \mathcal{R}$. Then there exists $A \in \mathcal{S}$ such that $B = A \cap R$. But then $A^c \in \mathcal{S}$ and $R \setminus B = R \cap B^c = R \cap A^c$, so $R \setminus B \in \mathcal{R}$. Finally, suppose that $B_i \in \mathcal{R}$ for i in a countable index set I . For each $i \in I$ there exists $A_i \in \mathcal{S}$ such that $B_i = A_i \cap R$. But then $\bigcup_{i \in I} A_i \in \mathcal{S}$ and $\bigcup_{i \in I} B_i = (\bigcup_{i \in I} A_i) \cap R$, so $\bigcup_{i \in I} B_i \in \mathcal{R}$.
2. Suppose that $R \in \mathcal{S}$. Then $A \cap R \in \mathcal{S}$ for every $A \in \mathcal{S}$, and of course, $A \cap R \subseteq R$. Conversely, if $B \in \mathcal{S}$ and $B \subseteq R$ then $B = B \cap R$ so $B \in \mathcal{R}$.

The σ -algebra \mathcal{R} is the σ -algebra on R induced by \mathcal{S} . The following construction is useful for counterexamples. Compare this example with the one for [finite and co-finite sets](#).

Let S be a nonempty set. The collection of *countable and co-countable* subsets of S is

$$\mathcal{C} = \{A \subseteq S : A \text{ is countable or } A^c \text{ is countable}\} \quad (1.11.5)$$

1. \mathcal{C} is a σ -algebra
2. $\mathcal{C} = \sigma\{\{x\} : x \in S\}$, the σ -algebra generated by the singleton sets.

Proof

1. First, $S \in \mathcal{C}$ since $S^c = \emptyset$ is countable. If $A \in \mathcal{C}$ then $A^c \in \mathcal{C}$ by the symmetry of the definition. Suppose that $A_i \in \mathcal{C}$ for each i in a countable index set I . If A_i is countable for each $i \in I$ then $\bigcup_{i \in I} A_i$ is countable. If A_j^c is countable for some $j \in I$ then $(\bigcup_{i \in I} A_i)^c = \bigcap_{i \in I} A_i^c \subseteq A_j^c$ is countable. In either case, $\bigcup_{i \in I} A_i \in \mathcal{C}$.
2. Let $\mathcal{D} = \sigma\{\{x\} : x \in S\}$. Clearly $\{x\} \in \mathcal{C}$ for $x \in S$. Hence $\mathcal{D} \subseteq \mathcal{C}$. Conversely, suppose that $A \in \mathcal{C}$. If A is countable, then $A = \bigcup_{x \in A} \{x\} \in \mathcal{D}$. If A^c is countable, then by an identical argument, $A^c \in \mathcal{D}$ and hence $A \in \mathcal{D}$.

Of course, if S is itself countable then $\mathcal{C} = \mathcal{P}(S)$. On the other hand, if S is uncountable, then there exists $A \subseteq S$ such that A and A^c are uncountable. Thus, $A \notin \mathcal{C}$, but $A = \bigcup_{x \in A} \{x\}$, and of course $\{x\} \in \mathcal{C}$. Thus, we have an example of a σ -algebra that is not closed under general unions.

Topology and Measure

One of the most important ways to generate a σ -algebra is by means of topology. Recall that a topological space consists of a set S and a topology \mathcal{S} , the collection of *open* subsets of S . Most spaces that occur in probability and stochastic processes are topological spaces, so it's crucial that the topological and measure-theoretic structures are compatible.

Suppose that (S, \mathcal{S}) is a topological space. Then $\sigma(\mathcal{S})$ is the *Borel σ -algebra* on S , and $(S, \sigma(\mathcal{S}))$ is a *Borel measurable space*.

So the Borel σ -algebra on S , named for Émile Borel is generated by the open subsets of S . Thus, a *topological space* (S, \mathcal{S}) naturally leads to a *measurable space* $(S, \sigma(\mathcal{S}))$. Since a closed set is simply the complement of an open set, the Borel σ -algebra contains the closed sets as well (and in fact is generated by the closed sets). Here are some other sets that are in the Borel σ -algebra:

Suppose again that (S, \mathcal{S}) is a topological space and that I is a countable index set.

1. If A_i is open for each $i \in I$ then $\bigcap_{i \in I} A_i \in \sigma(\mathcal{S})$. Such sets are called G_δ sets.
2. If A_i is closed for each $i \in I$ then $\bigcup_{i \in I} A_i \in \sigma(\mathcal{S})$. Such sets are called F_σ sets.
3. If (S, \mathcal{S}) is Hausdorff then $\{x\} \in \mathcal{S}$ for every $x \in S$.

Proof

1. This follows direction from the closure property for [intersections](#).
2. This follows from the [definition](#).
3. This follows since $\{x\}$ is closed for each $x \in S$ if the topology is Hausdorff.

In terms of part (c), recall that a topological space is *Hausdorff*, named for Felix Hausdorff, if the topology can distinguish individual points. Specifically, if $x, y \in S$ are distinct then there exist disjoint open sets U, V with $x \in U$ and $y \in V$. This is a

very basic property possessed by almost all topological spaces that occur in applications. A simple corollary of (c) is that if the topological space (S, \mathcal{S}) is Hausdorff then $A \in \sigma(\mathcal{S})$ for every countable $A \subseteq S$.

Let's note the extreme cases. If S has the *discrete topology* $\mathcal{P}(S)$, so that every set is open (and closed), then of course the Borel σ -algebra is also $\mathcal{P}(S)$. As noted above, this is often the appropriate σ -algebra if S is countable, but is often too large if S is uncountable. If S has the trivial topology $\{S, \emptyset\}$, then the Borel σ -algebra is also $\{S, \emptyset\}$, and so is also trivial.

Recall that a *base* for a topological space (S, \mathcal{T}) is a collection $\mathcal{B} \subseteq \mathcal{T}$ with the property that every set in \mathcal{T} is a union of a collection of sets in \mathcal{B} . In short, every open set is a union of some of the basic open sets.

Suppose that (S, \mathcal{S}) is a topological space with a countable base \mathcal{B} . Then $\sigma(\mathcal{B}) = \sigma(\mathcal{S})$.

Proof

Since $\mathcal{B} \subseteq \mathcal{S}$ it follows trivially that $\sigma(\mathcal{B}) \subseteq \sigma(\mathcal{S})$. Conversely, if $U \in \mathcal{S}$, there exists a collection of sets in \mathcal{B} whose union is U . Since \mathcal{B} is countable, $U \in \sigma(\mathcal{B})$.

The topological spaces that occur in probability and stochastic processes are usually assumed to have a countable base (along with other nice properties such as the Hausdorff property and locally compactness). The σ -algebra used for such a space is usually the Borel σ -algebra, which by the previous result, is countably generated.

Measurable Functions

Recall that a set usually comes with a σ -algebra of admissible subsets. A natural requirement on a function is that the inverse image of an admissible set in the range space be admissible in the domain space. Here is the formal definition.

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. A function $f : S \rightarrow T$ is *measurable* if $f^{-1}(A) \in \mathcal{S}$ for every $A \in \mathcal{T}$.

If the σ -algebra in the range space is generated by a collection of basic sets, then to check the measurability of a function, we need only consider inverse images of basic sets:

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces, and that $\mathcal{T} = \sigma(\mathcal{B})$ for a collection of subsets \mathcal{B} of T . Then $f : S \rightarrow T$ is measurable if and only if $f^{-1}(B) \in \mathcal{S}$ for every $B \in \mathcal{B}$.

Proof

First $\mathcal{B} \subseteq \mathcal{T}$, so if $f : S \rightarrow T$ is measurable then the condition in the theorem trivially holds. Conversely, suppose that the condition in the theorem holds, and let $\mathcal{U} = \{A \in \mathcal{T} : f^{-1}(A) \in \mathcal{S}\}$. Then $T \in \mathcal{U}$ since $f^{-1}(T) = S \in \mathcal{S}$. If $A \in \mathcal{U}$ then $f^{-1}(A^c) = [f^{-1}(A)]^c \in \mathcal{S}$, so $A^c \in \mathcal{U}$. If $A_i \in \mathcal{U}$ for i in a countable index set I , then $f^{-1}(\bigcup_{i \in I} A_i) = \bigcup_{i \in I} f^{-1}(A_i) \in \mathcal{S}$, and hence $\bigcup_{i \in I} A_i \in \mathcal{U}$. Thus \mathcal{U} is a σ -algebra of subsets of T . But $\mathcal{B} \subseteq \mathcal{U}$ by assumption, so $\mathcal{T} = \sigma(\mathcal{B}) \subseteq \mathcal{U}$. Of course $\mathcal{U} \subseteq \mathcal{T}$ by definition, so $\mathcal{U} = \mathcal{T}$ and hence f is measurable.

If you have reviewed the section on topology then you may have noticed a striking parallel between the definition of *continuity* for functions on topological spaces and the definition of *measurability* for functions on measurable spaces: A function from one topological space to another is *continuous* if the inverse image of an open set in the range space is open in the domain space. A function from one measurable space to another is *measurable* if the inverse image of a measurable set in the range space is measurable in the domain space. If we start with topological spaces, which we often do, and use the Borel σ -algebras to get measurable spaces, then we get the following (hardly surprising) connection.

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are topological spaces, and that we give S and T the Borel σ -algebras $\sigma(\mathcal{S})$ and $\sigma(\mathcal{T})$ respectively. If $f : S \rightarrow T$ is continuous, then f is measurable.

Proof

If $V \in \mathcal{T}$ then $f^{-1}(V) \in \mathcal{S} \subseteq \sigma(\mathcal{S})$. Hence f is measurable by the previous theorem.

Measurability is preserved under composition, the most important method for combining functions.

Suppose that (R, \mathcal{R}) , (S, \mathcal{S}) , and (T, \mathcal{T}) are measurable spaces. If $f : R \rightarrow S$ is measurable and $g : S \rightarrow T$ is measurable, then $g \circ f : R \rightarrow T$ is measurable.

Proof

If $A \in \mathcal{T}$ then $g^{-1}(A) \in \mathcal{S}$ since g is measurable, and hence $(g \circ f)^{-1}(A) = f^{-1}[g^{-1}(A)] \in \mathcal{R}$ since f is measurable.

If T is given the smallest possible σ -algebra or if S is given the largest one, then any function from S into T is measurable.

Every function $f : S \rightarrow T$ is measurable in each of the following cases:

1. $\mathcal{T} = \{\emptyset, T\}$ and \mathcal{S} is an arbitrary σ -algebra of subsets of S
2. $\mathcal{S} = \mathcal{P}(S)$ and \mathcal{T} is an arbitrary σ -algebra of subsets of T .

Proof

1. Suppose that $\mathcal{T} = \{\emptyset, T\}$ and that \mathcal{S} is an arbitrary σ -algebra on S . If $f : S \rightarrow T$, then $f^{-1}(T) = S \in \mathcal{S}$ and $f^{-1}(\emptyset) = \emptyset \in \mathcal{S}$ so f is measurable.
2. Suppose that $\mathcal{S} = \mathcal{P}(S)$ and that \mathcal{T} is an arbitrary σ -algebra on T . If $f : S \rightarrow T$, then trivially $f^{-1}(A) \in \mathcal{S}$ for every $A \in \mathcal{T}$ so f is measurable.

When there are several σ -algebras for the same set, then we use the phrase *with respect to* so that we can be precise. If a function is measurable with respect to a given σ -algebra on its domain, then it's measurable with respect to any larger σ -algebra on this space. If the function is measurable with respect to a σ -algebra on the range space then it's measurable with respect to any smaller σ -algebra on this space.

Suppose that S has σ -algebras \mathcal{R} and \mathcal{S} with $\mathcal{R} \subseteq \mathcal{S}$, and that T has σ -algebras \mathcal{T} and \mathcal{U} with $\mathcal{T} \subseteq \mathcal{U}$. If $f : S \rightarrow T$ is measurable with respect to \mathcal{R} and \mathcal{U} , then f is measurable with respect to \mathcal{S} and \mathcal{T} .

Proof

If $A \in \mathcal{T}$ then $A \in \mathcal{U}$. Hence $f^{-1}(A) \in \mathcal{R}$ so $f^{-1}(A) \in \mathcal{S}$.

The following construction is particularly important in probability theory:

Suppose that S is a set and (T, \mathcal{T}) is a measurable space. Suppose also that $f : S \rightarrow T$ and define $\sigma(f) = \{f^{-1}(A) : A \in \mathcal{T}\}$. Then

1. $\sigma(f)$ is a σ -algebra on S .
2. $\sigma(f)$ is the smallest σ -algebra on S that makes f measurable.

Proof

1. The key to the proof is that the inverse image preserves all set operations. First, $S \in \sigma(f)$ since $T \in \mathcal{T}$ and $f^{-1}(T) = S$. If $B \in \sigma(f)$ then $B = f^{-1}(A)$ for some $A \in \mathcal{T}$. But then $A^c \in \mathcal{T}$ and hence $B^c = f^{-1}(A^c) \in \sigma(f)$. Finally, suppose that $B_i \in \sigma(f)$ for i in a countable index set I . Then for each $i \in I$ there exists $A_i \in \mathcal{T}$ such that $B_i = f^{-1}(A_i)$. But then $\bigcup_{i \in I} A_i \in \mathcal{T}$ and $\bigcup_{i \in I} B_i = f^{-1}(\bigcup_{i \in I} A_i)$. Hence $\bigcup_{i \in I} B_i \in \sigma(f)$.
2. If \mathcal{S} is a σ -algebra on S and f is measurable with respect to \mathcal{S} and \mathcal{T} , then by definition $f^{-1}(A) \in \mathcal{S}$ for every $A \in \mathcal{T}$, so $\sigma(f) \subseteq \mathcal{S}$.

Appropriately enough, $\sigma(f)$ is called the σ -algebra generated by f . Often, S will have a given σ -algebra \mathcal{S} and f will be measurable with respect to \mathcal{S} and \mathcal{T} . In this case, $\sigma(f) \subseteq \mathcal{S}$. We can generalize to an arbitrary collection of functions on S .

Suppose S is a set and that (T_i, \mathcal{T}_i) is a measurable space for each i in a nonempty index set I . Suppose also that $f_i : S \rightarrow T_i$ for each $i \in I$. The σ -algebra generated by this collection of functions is

$$\sigma\{f_i : i \in I\} = \sigma\{\sigma(f_i) : i \in I\} = \sigma\{f_i^{-1}(A) : i \in I, A \in \mathcal{T}_i\} \quad (1.11.6)$$

Again, this is the smallest σ -algebra on S that makes f_i measurable for each $i \in I$.

Product Sets

Product sets arise naturally in the form of the higher-dimensional Euclidean spaces \mathbb{R}^n for $n \in \{2, 3, \dots\}$. In addition, product spaces are particularly important in probability, where they are used to describe the spaces associated with sequences of random variables. More general product spaces arise in the study of stochastic processes. We start with the product of two sets; the generalization to products of n sets and to general products is straightforward, although the notation gets more complicated.

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. The *product σ -algebra* on $S \times T$ is

$$\mathcal{S} \otimes \mathcal{T} = \sigma\{A \times B : A \in \mathcal{S}, B \in \mathcal{T}\} \quad (1.11.7)$$

So the definition is natural: the product σ -algebra is generated by products of measurable sets. Our next goal is to consider the measurability of functions defined on, or mapping into, product spaces. Of basic importance are the projection functions. If S and T are sets, let $p_1 : S \times T \rightarrow S$ and $p_2 : S \times T \rightarrow T$ be defined by $p_1(x, y) = x$ and $p_2(x, y) = y$ for $(x, y) \in S \times T$. Recall that p_1 is the *projection onto the first coordinate* and p_2 is the *projection onto the second coordinate*. The product σ algebra is the smallest σ -algebra that makes the projections measurable:

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. Then $\mathcal{S} \otimes \mathcal{T} = \sigma\{p_1, p_2\}$.

Proof

If $A \in \mathcal{S}$ then $p_1^{-1}(A) = A \times T \in \mathcal{S} \otimes \mathcal{T}$. Similarly, if $B \in \mathcal{T}$ then $p_2^{-1}(B) = S \times B \in \mathcal{S} \otimes \mathcal{T}$. Hence p_1 and p_2 are measurable, so $\sigma\{p_1, p_2\} \subseteq \mathcal{S} \otimes \mathcal{T}$. Conversely, if $A \in \mathcal{S}$ and $B \in \mathcal{T}$ then $A \times B = p_1^{-1}(A) \cap p_2^{-1}(B) \in \sigma\{p_1, p_2\}$. Since sets of this form generate the product σ -algebra, we have $\mathcal{S} \otimes \mathcal{T} \subseteq \sigma\{p_1, p_2\}$.

Projection functions make it easy to study functions mapping into a product space.

Suppose that (R, \mathcal{R}) , (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces, and that $S \times T$ is given the product σ -algebra $\mathcal{S} \otimes \mathcal{T}$. Suppose also that $f : R \rightarrow S \times T$, so that $f(x) = (f_1(x), f_2(x))$ for $x \in R$, where $f_1 : R \rightarrow S$ and $f_2 : R \rightarrow T$ are the *coordinate functions*. Then f is measurable if and only if f_1 and f_2 are measurable.

Proof

Note that $f_1 = p_1 \circ f$ and $f_2 = p_2 \circ f$. So if f is measurable then f_1 and f_2 are compositions of measurable functions, and hence are measurable. Conversely, suppose that f_1 and f_2 are measurable. If $A \in \mathcal{S}$ and $B \in \mathcal{T}$ then $f^{-1}(A \times B) = f_1^{-1}(A) \cap f_2^{-1}(B) \in \mathcal{R}$. Since products of measurable sets generate $\mathcal{S} \otimes \mathcal{T}$, it follows that f is measurable.

Our next goal is to consider cross sections of sets in a product space and cross sections of functions defined on a product space. It will help to introduce some new functions, which in a sense are complementary to the projection functions.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces, and that $S \times T$ is given the product σ -algebra $\mathcal{S} \otimes \mathcal{T}$.

1. For $x \in S$ the function $1_x : T \rightarrow S \times T$, defined by $1_x(y) = (x, y)$ for $y \in T$, is measurable.
2. For $y \in T$ the function $2_y : S \rightarrow S \times T$, defined by $2_y(x) = (x, y)$ for $x \in S$, is measurable.

Proof

To show that the functions are measurable, it suffices to consider inverse images of products of measurable sets, since such sets generate $\mathcal{S} \otimes \mathcal{T}$. Thus, let $A \in \mathcal{S}$ and $B \in \mathcal{T}$.

1. For $x \in S$ note that $1_x^{-1}(A \times B)$ is B if $x \in A$ and is \emptyset if $x \notin A$. In either case, $1_x^{-1}(A \times B) \in \mathcal{T}$.
2. Similarly, for $y \in T$ note that $2_y^{-1}(A \times B)$ is A if $y \in B$ and is \emptyset if $y \notin B$. In either case, $2_y^{-1}(A \times B) \in \mathcal{S}$.

Now our work is easy.

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces, and that $C \in \mathcal{S} \otimes \mathcal{T}$. Then

1. For $x \in S$, $\{y \in T : (x, y) \in C\} \in \mathcal{T}$.
2. For $y \in T$, $\{x \in S : (x, y) \in C\} \in \mathcal{S}$.

Proof

These result follow immediately from the measurability of the functions 1_x and 2_y :

1. For $x \in S$, $1_x^{-1}(C) = \{y \in T : (x, y) \in C\}$.
2. For $y \in T$, $2_y^{-1}(C) = \{x \in S : (x, y) \in C\}$.

The set in (a) is the *cross section of C in the first coordinate at x* , and the set in (b) is the *cross section of C in the second coordinate at y* . As a simple corollary to the theorem, note that if $A \subseteq S$, $B \subseteq T$ and $A \times B \in \mathcal{S} \otimes \mathcal{T}$ then $A \in \mathcal{S}$ and $B \in \mathcal{T}$. That is, the only measurable product sets are products of measurable sets. Here is the measurability result for cross-sectional functions:

Suppose again that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces, and that $S \times T$ is given the product σ -algebra $\mathcal{S} \otimes \mathcal{T}$. Suppose also that (U, \mathcal{U}) is another measurable space, and that $f : S \times T \rightarrow U$ is measurable. Then

1. The function $y \mapsto f(x, y)$ from T to U is measurable for each $x \in S$.
2. The function $x \mapsto f(x, y)$ from S to U is measurable for each $y \in T$.

Proof

Note that the function in (a) is just $f \circ 1_x$, and the function in (b) is just $f \circ 2_y$, both are compositions of measurable functions

The results for products of two spaces generalize in a completely straightforward way to a product of n spaces.

Suppose $n \in \mathbb{N}_+$ and that (S_i, \mathcal{S}_i) is a measurable space for each $i \in \{1, 2, \dots, n\}$. The *product σ -algebra* on the Cartesian product set $S_1 \times S_2 \times \dots \times S_n$ is

$$\mathcal{S}_1 \otimes \mathcal{S}_2 \otimes \dots \otimes \mathcal{S}_n = \sigma \{A_1 \times A_2 \times \dots \times A_n : A_i \in \mathcal{S}_i \text{ for all } i \in \{1, 2, \dots, n\}\} \quad (1.11.8)$$

So again, the product σ -algebra is generated by products of measurable sets. Results analogous to the theorems above hold. In the special case that $(S_i, \mathcal{S}_i) = (S, \mathcal{S})$ for $i \in \{1, 2, \dots, n\}$, the Cartesian product becomes S^n and the corresponding product σ -algebra is denoted \mathcal{S}^n . The notation is natural, but potentially confusing. Note that \mathcal{S}^n is not the Cartesian product of \mathcal{S} n times, but rather the σ -algebra generated by sets of the form $A_1 \times A_2 \times \dots \times A_n$ where $A_i \in \mathcal{S}$ for $i \in \{1, 2, \dots, n\}$.

We can also extend these ideas to a general product. To recall the definition, suppose that S_i is a set for each i in a nonempty index set I . The product set $\prod_{i \in I} S_i$ consists of all functions $x : I \rightarrow \bigcup_{i \in I} S_i$ such that $x(i) \in S_i$ for each $i \in I$. To make the notation look more like a simple Cartesian product, we often write x_i instead of $x(i)$ for the value of a function in the product set at $i \in I$. The next definition gives the appropriate σ -algebra for the product set.

Suppose that (S_i, \mathcal{S}_i) is a measurable space for each i in a nonempty index set I . The *product σ -algebra* on the product set $\prod_{i \in I} S_i$ is

$$\sigma \left\{ \prod_{i \in I} A_i : A_i \in \mathcal{S}_i \text{ for each } i \in I \text{ and } A_i = S_i \text{ for all but finitely many } i \in I \right\} \quad (1.11.9)$$

If you have reviewed the section on topology, the definition should look familiar. If the spaces were *topological spaces* instead of *measurable spaces*, with \mathcal{S}_i the topology of S_i for $i \in I$, then the set of products in the displayed expression above is a base for the product topology on $\prod_{i \in I} S_i$.

The definition can also be understood in terms of projections. Recall that the *projection* onto coordinate $j \in I$ is the function $p_j : \prod_{i \in I} S_i \rightarrow S_j$ given by $p_j(x) = x_j$. The product σ -algebra is the smallest σ -algebra on the product set that makes all of the projections measurable.

Suppose again that (S_i, \mathcal{S}_i) is a measurable space for each i in a nonempty index set I , and let \mathfrak{S} denote the product σ -algebra on the product set $S_I = \prod_{i \in I} S_i$. Then $\mathfrak{S} = \sigma\{p_i : i \in I\}$.

Proof

Let $j \in I$ and $A \in \mathcal{S}_j$. Then $p_j^{-1}(A) = \prod_{i \in I} A_i$ where $A_i = S_i$ for $i \neq j$ and $A_j = A$. This set is in \mathfrak{S} so p_j is measurable. Hence $\sigma\{p_i : i \in I\} \subseteq \mathfrak{S}$. For the other direction, consider a product set $\prod_{i \in I} A_i$ where $A_i = S_i$ except for $i \in J$, where

$J \subseteq I$ is finite. Then $\prod_{i \in I} A_i = \bigcap_{j \in J} p_j^{-1}(A_j)$. This set is in $\sigma\{p_i : i \in I\}$. Product sets of this form generate \mathfrak{S} so it follows that $\mathfrak{S} \subseteq \sigma\{p_i : i \in I\}$.

In the special case that (S, \mathcal{S}) is a fixed measurable space and $(S_i, \mathcal{S}_i) = (S, \mathcal{S})$ for all $i \in I$, the product set $\prod_{i \in I} S$ is just the collection of functions from I into S , often denoted S^I . The product σ -algebra is then denoted \mathcal{S}^I , a notation that is natural, but again potentially confusing. Here is the main measurability result for a function mapping into a product space.

Suppose that (R, \mathcal{R}) is a measurable space, and that (S_i, \mathcal{S}_i) is a measurable space for each i in a nonempty index set I . As before, let $\prod_{i \in I} S_i$ have the product σ -algebra. Suppose now that $f : R \rightarrow \prod_{i \in I} S_i$. For $i \in I$ let $f_i : R \rightarrow S_i$ denote the i th coordinate function of f , so that $f_i(x) = [f(x)]_i$ for $x \in R$. Then f is measurable if and only if f_i is measurable for each $i \in I$.

Proof

Suppose that f is measurable. For $i \in I$ note that $f_i = p_i \circ f$ is a composition of measurable functions, and hence is measurable. Conversely, suppose that f_i is measurable for each $i \in I$. To show that measurability of f we need only consider inverse images of sets that generate the product σ -algebra. Thus, suppose that $A_j \in \mathcal{S}_j$ for j in a finite subset $J \subseteq I$, and let $A_i = S_i$ for $i \in I - J$. Then $f^{-1}(\prod_{i \in I} A_i) = \bigcap_{j \in J} f_j^{-1}(A_j)$. This set is in \mathcal{R} since the intersection is over a finite index set.

Just as with the product of two sets, cross-sectional sets and functions are measurable with respect to the product measure. Again, it's best to work with some special functions.

Suppose that (S_i, \mathcal{S}_i) is a measurable space for each i in an index set I with at least two elements. For $j \in I$ and $u \in S_j$, define the function $j_u : \prod_{i \in I - \{j\}} S_i \rightarrow \prod_{i \in I} S_i$ by $j_u(x) = y$ where $y_i = x_i$ for $i \neq j$ and $y_j = u$. Then j_u is measurable with respect to the product σ -algebras.

Proof

Once again, it suffices to consider the inverse image of the sets that generate the product σ -algebra. So suppose $A_i \in \mathcal{S}_i$ for $i \in I$ with $A_i = S_i$ for all but finitely many $i \in I$. Then $j_u^{-1}(\prod_{i \in I} A_i) = \prod_{i \in I - \{j\}} A_i$ if $u \in A_j$, and the inverse image is \emptyset otherwise. In either case, $j_u^{-1}(\prod_{i \in I} A_i)$ is in the product σ -algebra on $\prod_{i \in I - \{j\}} S_i$.

In words, for $j \in I$ and $u \in S_j$, the function j_u takes a point in the product set $\prod_{i \in I - \{j\}} S_i$ and assigns u to coordinate j to give a point in $\prod_{i \in I} S_i$. If $A \subseteq \prod_{i \in I} S_i$, then $j_u^{-1}(A)$ is the cross section of A in coordinate j at u . So it follows immediately from the previous result that the cross sections of a measurable set are measurable. Cross sections of measurable functions are also measurable. Suppose that (T, \mathcal{T}) is another measurable space, and that $f : \prod_{i \in I} S_i \rightarrow T$ is measurable. The cross section of f in coordinate $j \in I$ at $u \in S_j$ is simply $f \circ j_u : S_{I - \{j\}} \rightarrow T$, a composition of measurable functions.

However, a non-measurable set can have measurable cross sections, even in a product of two spaces.

Suppose that S is an uncountable set with the σ -algebra \mathcal{C} of countable and co-countable sets as in (21). Consider $S \times S$ with the product σ -algebra $\mathcal{C} \otimes \mathcal{C}$. Let $D = \{(x, x) : x \in S\}$, the *diagonal* of $S \times S$. Then D has measurable cross sections, but D is not measurable.

Proof

For $x \in S$, the cross section of D in the first coordinate at x is $\{y \in S : (x, y) \in D\} = \{x\} \in \mathcal{C}$. Similarly, for $y \in S$, the cross section of D in the second coordinate at y is $\{x \in S : (x, y) \in D\} = \{y\} \in \mathcal{C}$. But D cannot be generated by a countable collection of sets of the form $A \times B$ with $A, B \in \mathcal{C}$, so $D \notin \mathcal{C} \otimes \mathcal{C}$, by the result above.

Special Cases

Most of the sets encountered in applied probability are either countable, or subsets of \mathbb{R}^n for some n , or more generally, subsets of a product of a countable number of sets of these types. In the study of stochastic processes, various spaces of functions play an important role. In this subsection, we will explore the most important special cases.

Discrete Spaces

If S is countable and $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S , then (S, \mathcal{S}) is a *discrete measurable space*.

Thus if (S, \mathcal{S}) is discrete, all subsets of S are measurable and every function from S to another measurable space is measurable. The power set is also the discrete topology on S , so \mathcal{S} is a Borel σ -algebra as well. As a topological space, (S, \mathcal{S}) is complete, locally compact, Hausdorff, and since S is countable, separable. Moreover, the discrete topology corresponds to the discrete metric d , defined by $d(x, x) = 0$ for $x \in S$ and $d(x, y) = 1$ for $x, y \in S$ with $x \neq y$.

Euclidean Spaces

Recall that for $n \in \mathbb{N}_+$, the Euclidean topology on \mathbb{R}^n is generated by the standard Euclidean metric d_n given by

$$d_n(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n), \mathbf{y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n \quad (1.11.10)$$

With this topology, \mathbb{R}^n is complete, connected, locally compact, Hausdorff, and separable.

For $n \in \mathbb{N}_+$, the n -dimensional *Euclidean measurable space* is $(\mathbb{R}^n, \mathcal{R}_n)$ where \mathcal{R}_n is the Borel σ -algebra corresponding to the standard Euclidean topology on \mathbb{R}^n .

The one-dimensional case is particularly important. In this case, the standard Euclidean metric d is given by $d(x, y) = |x - y|$ for $x, y \in \mathbb{R}$. The Borel σ -algebra \mathcal{R} can be generated by various collections of intervals.

Each of the following collections generates \mathcal{R} .

1. $\mathcal{B}_1 = \{I \subseteq \mathbb{R} : I \text{ is an interval}\}$
2. $\mathcal{B}_2 = \{(a, b] : a, b \in \mathbb{R}, a < b\}$
3. $\mathcal{B}_3 = \{(-\infty, b] : b \in \mathbb{R}\}$

Proof

The proof involves showing that each set in any one of the collections is in the σ -algebra of any other collection. Let $\mathcal{S}_i = \sigma(\mathcal{B}_i)$ for $i \in \{1, 2, 3\}$

1. Clearly $\mathcal{B}_2 \subseteq \mathcal{B}_1$ and $\mathcal{B}_3 \subseteq \mathcal{B}_1$ so $\mathcal{S}_2 \subseteq \mathcal{S}_1$ and $\mathcal{S}_3 \subseteq \mathcal{S}_1$.
2. If $a, b \in \mathbb{R}$ with $a \leq b$ then $[a, b] = \bigcap_{n=1}^{\infty} (a - \frac{1}{n}, b]$ and $(a, b) = \bigcup_{n=1}^{\infty} (a, b - \frac{1}{n}]$, so $[a, b], (a, b) \in \mathcal{S}_2$. Also $[a, b) = \bigcup_{n=1}^{\infty} [a, b - \frac{1}{n}]$ so $[a, b) \in \mathcal{S}_2$. Thus all bounded intervals are in \mathcal{S}_2 . Next, $[a, \infty) = \bigcup_{n=1}^{\infty} [a, a + n)$, $(-\infty, a] = \bigcap_{n=1}^{\infty} (a - n, a]$, and $(-\infty, a) = \bigcup_{n=1}^{\infty} (a - n, a)$, so each of these intervals is in \mathcal{S}_2 . Of course $\mathbb{R} \in \mathcal{S}_2$, so we now have that $I \in \mathcal{S}_2$ for every interval I . Thus $\mathcal{S}_1 \subseteq \mathcal{S}_2$, and so from (a), $\mathcal{S}_2 = \mathcal{S}_1$.
3. If $a, b \in \mathbb{R}$ with $a < b$ then $(a, b] = (-\infty, b] - (-\infty, a]$ so $(a, b] \in \mathcal{S}_3$. Hence $\mathcal{S}_2 \subseteq \mathcal{S}_3$. But then from (a) and (b) it follows that $\mathcal{S}_3 = \mathcal{S}_1$.

Since the Euclidean topology has a countable base, \mathcal{R} is countably generated. In fact each collection of intervals above, but with endpoints restricted to \mathbb{Q} , generates \mathcal{R} . Moreover, \mathcal{R} can also be constructed from σ -algebras that are generated by [countable partitions](#). First recall that for $n \in \mathbb{N}$, the set of dyadic rationals (or *binary rationals*) of rank n or less is $\mathbb{D}_n = \{j/2^n : j \in \mathbb{Z}\}$. Note that \mathbb{D}_n is countable and $\mathbb{D}_n \subseteq \mathbb{D}_{n+1}$ for $n \in \mathbb{N}$. Moreover, the set $\mathbb{D} = \bigcup_{n \in \mathbb{N}} \mathbb{D}_n$ of *all* dyadic rationals is dense in \mathbb{R} . The dyadic rationals are often useful in various applications because \mathbb{D}_n has the natural ordered enumeration $j \mapsto j/2^n$ for each $n \in \mathbb{N}$. Now let

$$\mathcal{D}_n = \{(j/2^n, (j+1)/2^n] : j \in \mathbb{Z}\}, \quad n \in \mathbb{N} \quad (1.11.11)$$

Then \mathcal{D}_n is a countable partition of \mathbb{R} into nonempty intervals of equal size $1/2^n$, so $\mathcal{E}_n = \sigma(\mathcal{D}_n)$ consists of unions of sets in \mathcal{D}_n as described [above](#). Every set \mathcal{D}_n is the union of two sets in \mathcal{D}_{n+1} so clearly $\mathcal{E}_n \subseteq \mathcal{E}_{n+1}$ for $n \in \mathbb{N}$. Finally, the Borel σ -algebra on \mathbb{R} is $\mathcal{R} = \sigma(\bigcup_{n=0}^{\infty} \mathcal{E}_n) = \sigma(\bigcup_{n=0}^{\infty} \mathcal{D}_n)$. This construction turns out to be useful in a number of settings.

For $n \in \{2, 3, \dots\}$, the Euclidean topology on \mathbb{R}^n is the n -fold product topology formed from the Euclidean topology on \mathbb{R} . So the Borel σ -algebra \mathcal{R}^n is also the n -fold power σ -algebra formed from \mathcal{R} . Finally, \mathcal{R}^n can be generated by n -fold products of sets in any of the three collections in the [previous theorem](#).

Space of Real Functions

Suppose that (S, \mathcal{S}) is a measurable space. From our general discussion of functions, recall that the usual arithmetic operations on functions from S into \mathbb{R} are defined pointwise.

If $f : S \rightarrow \mathbb{R}$ and $g : S \rightarrow \mathbb{R}$ are measurable and $a \in \mathbb{R}$, then each of the following functions from S into \mathbb{R} is also measurable:

1. $f + g$
2. $f - g$
3. fg
4. af

Proof

These results follow from the fact that the arithmetic operators are continuous, and hence measurable. That is, $(x, y) \mapsto x + y$, $(x, y) \mapsto x - y$, and $(x, y) \mapsto xy$ are continuous as functions from \mathbb{R}^2 into \mathbb{R} . Thus, if $f, g : S \rightarrow \mathbb{R}$ are measurable, then $(f, g) : S \rightarrow \mathbb{R}^2$ is measurable by the [result above](#). Then, $f + g$, $f - g$, fg are the compositions, respectively, of $+$, $-$, \cdot with (f, g) . Of course, (d) is a simple corollary of (c).

Similarly, if $f : S \rightarrow \mathbb{R} \setminus \{0\}$ is measurable, then so is $1/f$. Recall that the set of functions from S into \mathbb{R} is a vector space, under the pointwise definitions of addition and scalar multiplication. But once again, we usually want to restrict our attention to *measurable* functions. Thus, it's nice to know that the measurable functions from S into \mathbb{R} also form a vector space. This follows immediately from the *closure properties* (a) and (d) of the previous theorem. Of particular importance in probability and stochastic processes is the vector space of bounded, measurable functions $f : S \rightarrow \mathbb{R}$, with the supremum norm

$$\|f\| = \sup \{|f(x)| : x \in S\} \quad (1.11.12)$$

The *elementary functions* that we encounter in calculus and other areas of applied mathematics are functions from subsets of \mathbb{R} into \mathbb{R} . The elementary functions include algebraic functions (which in turn include the polynomial and rational functions), the usual transcendental functions (exponential, logarithm, trigonometric), and the usual functions constructed from these by composition, the arithmetic operations, and by piecing together. As we might hope, all of the elementary functions are measurable.

This page titled [1.11: Measurable Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

1.12: Special Set Structures

There are several other types of algebraic set structures that are weaker than σ -algebras. These are not particularly important in themselves, but are important for constructing σ -algebras and the measures on these σ -algebras. You may want to skip this section if you are not interested in questions of existence and uniqueness of positive measures.

Basic Theory

Definitions

Throughout this section, we assume that S is a set and \mathcal{S} is a nonempty collection of subsets of S . Here are the main definitions we will need.

\mathcal{S} is a π -system if \mathcal{S} is closed under finite intersections: if $A, B \in \mathcal{S}$ then $A \cap B \in \mathcal{S}$.

Closure under intersection is clearly a very simple property, but π systems turn out to be useful enough to deserve a name.

\mathcal{S} is a λ -system if it is closed under complements and countable *disjoint* unions.

1. If $A \in \mathcal{S}$ then $A^c \in \mathcal{S}$.
2. If $A_i \in \mathcal{S}$ for i in a countable index set I and $A_i \cap A_j = \emptyset$ for $i \neq j$ then $\bigcup_{i \in I} A_i \in \mathcal{S}$.

\mathcal{S} is a *semi-algebra* if it is closed under intersection and if complements can be written as finite, disjoint unions:

1. If $A, B \in \mathcal{S}$ then $A \cap B \in \mathcal{S}$.
2. If $A \in \mathcal{S}$ then there exists a finite, disjoint collection $\{B_i : i \in I\} \subseteq \mathcal{S}$ such that $A^c = \bigcup_{i \in I} B_i$.

For our final structure, recall that a sequence (A_1, A_2, \dots) of subsets of S is *increasing* if $A_n \subseteq A_{n+1}$ for all $n \in \mathbb{N}_+$. The sequence is *decreasing* if $A_{n+1} \subseteq A_n$ for all $n \in \mathbb{N}_+$. Of course, these are the standard meanings of *increasing* and *decreasing* relative to the ordinary order \leq on \mathbb{N}_+ and the subset partial order \subseteq on $\mathcal{P}(S)$.

\mathcal{S} is a *monotone class* if it is closed under increasing unions and decreasing intersections:

1. If (A_1, A_2, \dots) is an increasing sequence of sets in \mathcal{S} then $\bigcup_{n=1}^{\infty} A_n \in \mathcal{S}$.
2. If (A_1, A_2, \dots) is a decreasing sequence of sets in \mathcal{S} then $\bigcap_{n=1}^{\infty} A_n \in \mathcal{S}$.

If (A_1, A_2, \dots) is an increasing sequence of sets then we sometimes write $\bigcup_{n=1}^{\infty} A_n = \lim_{n \rightarrow \infty} A_n$. Similarly, if (A_1, A_2, \dots) is a decreasing sequence of sets we sometimes write $\bigcap_{n=1}^{\infty} A_n = \lim_{n \rightarrow \infty} A_n$. The reason for this notation will become clear in the section on Convergence in the chapter on Probability Spaces. With this notation, a monotone class \mathcal{S} is defined by the condition that if (A_1, A_2, \dots) is an increasing or decreasing sequence of sets in \mathcal{S} then $\lim_{n \rightarrow \infty} A_n \in \mathcal{S}$.

Basic Theorems

Our most important set structure, the σ -algebra, has all of the properties in the definitions above.

If \mathcal{S} is a σ -algebra then \mathcal{S} is a π -system, a λ -system, a semi-algebra, and a monotone class.

If \mathcal{S} is a λ -system then $S \in \mathcal{S}$ and $\emptyset \in \mathcal{S}$.

Proof

The proof is just like the one for an algebra. There exists $A \in \mathcal{S}$ since \mathcal{S} is non-empty. Hence $A^c \in \mathcal{S}$ and so $S = A \cup A^c \in \mathcal{S}$. Finally $\emptyset = S^c \in \mathcal{S}$.

Any type of algebraic structure on subsets of S that is defined purely in terms of *closure properties* will be preserved under intersection. That is, we will have results that are analogous to how σ -algebras are generated from more basic sets, with completely straightforward and analogous proofs. In the following two theorems, the term *system* could mean π -system, λ -system, or monotone class of subsets of S .

If \mathcal{S}_i is a system for each i in an index set I and $\bigcap_{i \in I} \mathcal{S}_i$ is nonempty, then $\bigcap_{i \in I} \mathcal{S}_i$ is a system of the same type.

The condition that $\bigcap_{i \in I} \mathcal{S}_i$ be nonempty is unnecessary for a λ -system, by the [result above](#). Now suppose that \mathcal{B} is a nonempty collection of subsets of S , thought of as basic sets of some sort. Then the system generated by \mathcal{B} is the intersection of all systems that contain \mathcal{B} .

The system \mathcal{S} generated by \mathcal{B} is the smallest system containing \mathcal{B} , and is characterized by the following properties:

1. $\mathcal{B} \subseteq \mathcal{S}$.
2. If \mathcal{T} is a system and $\mathcal{B} \subseteq \mathcal{T}$ then $\mathcal{S} \subseteq \mathcal{T}$.

Note however, that the previous two results do not apply to semi-algebras, because the semi-algebra is not defined purely in terms of closure properties (the condition on A^c is not a closure property).

If \mathcal{S} is a monotone class and an algebra, then \mathcal{S} is a σ -algebra.

Proof

All that is needed is to prove closure under countable unions. Thus, suppose that $A_i \in \mathcal{S}$ for $i \in \mathbb{N}_+$. Then $B_n = \bigcup_{i=1}^n A_i \in \mathcal{S}$ since \mathcal{S} is an algebra. The sequence (B_1, B_2, \dots) is increasing, so $\bigcup_{n=1}^{\infty} B_n \in \mathcal{S}$, since \mathcal{S} is a monotone class. But $\bigcup_{n=1}^{\infty} B_n = \bigcup_{i=1}^{\infty} A_i$.

By definition, a semi-algebra is a π -system. More importantly, a semi-algebra can be used to construct an algebra.

Suppose that \mathcal{S} is a semi-algebra of subsets of S . Then the collection \mathcal{S}^* of finite, disjoint unions of sets in \mathcal{S} is an algebra.

Proof

Suppose that $A, B \in \mathcal{S}^*$. Then there exist finite, disjoint collections $\{A_i : i \in I\} \subseteq \mathcal{S}$ and $\{B_j : j \in J\} \subseteq \mathcal{S}$ such that $A = \bigcup_{i \in I} A_i$ and $B = \bigcup_{j \in J} B_j$. Hence

$$A \cap B = \bigcup_{(i,j) \in I \times J} (A_i \cap B_j) \quad (1.12.1)$$

But $\{A_i \cap B_j : (i,j) \in I \times J\}$ is a finite, disjoint collection of sets in \mathcal{S} , so $A \cap B \in \mathcal{S}^*$. Suppose $A \in \mathcal{S}^*$, so that there exists a finite, disjoint collection $\{A_i : i \in I\}$ such that $A = \bigcup_{i \in I} A_i$. Then $A^c = \bigcap_{i \in I} A_i^c$. But $A_i^c \in \mathcal{S}^*$ by definition of semi-algebra, and we just showed that \mathcal{S}^* is closed under finite intersections, so $A^c \in \mathcal{S}^*$.

We will say that our nonempty collection \mathcal{S} is *closed under proper set difference* if $A, B \in \mathcal{S}$ and $A \subseteq B$ implies $B \setminus A \in \mathcal{S}$. The following theorem gives the basic relationship between λ -systems and monotone classes.

Suppose that \mathcal{S} is a nonempty collection of subsets of S .

1. If \mathcal{S} is a λ -system then \mathcal{S} is a monotone class and is closed under proper set difference.
2. If \mathcal{S} is a monotone class, is closed under proper set difference, and contains S , then \mathcal{S} is a λ -system.

Proof

1. Suppose that \mathcal{S} is a λ -system. Suppose that $A, B \in \mathcal{S}$ and $A \subseteq B$. Then $B^c \in \mathcal{S}$, and A and B^c are disjoint, so $A \cup B^c \in \mathcal{S}$. But then $(A \cup B^c)^c = B \cap A^c = B \setminus A \in \mathcal{S}$. Hence \mathcal{S} is closed under proper set difference. Next suppose that (A_1, A_2, \dots) is an increasing sequence of sets in \mathcal{S} . Let $B_1 = A_1$ and $B_n = A_n \setminus A_{n-1}$ for $n \in \{2, 3, \dots\}$. Then $B_i \in \mathcal{S}$ for each $i \in \mathbb{N}_+$. But the sequence (B_1, B_2, \dots) is disjoint and has the same union as (A_1, A_2, \dots) . Hence $\bigcup_{i=1}^{\infty} A_i = \bigcup_{i=1}^{\infty} B_i \in \mathcal{S}$. Finally, suppose that (A_1, A_2, \dots) is a decreasing sequence of sets in \mathcal{S} . Then $A_i^c \in \mathcal{S}$ for each $i \in \mathbb{N}_+$ and (A_1^c, A_2^c, \dots) is increasing. Hence $\bigcup_{i=1}^{\infty} A_i^c \in \mathcal{S}$ and therefore $(\bigcup_{i=1}^{\infty} A_i^c)^c = \bigcap_{i=1}^{\infty} A_i \in \mathcal{S}$.
2. Suppose that \mathcal{S} is a monotone class, is closed under proper set difference, and $S \in \mathcal{S}$. If $A \in \mathcal{S}$ then trivially $A \subseteq S$ so $A^c = S \setminus A \in \mathcal{S}$. Next, suppose that $A, B \in \mathcal{S}$ are disjoint. Then $A^c \in \mathcal{S}$ and $B \subseteq A^c$, so $A^c \setminus B = A^c \cap B^c \in \mathcal{S}$. Hence $A \cup B = (A^c \cap B^c)^c \in \mathcal{S}$. Finally, suppose that (A_1, A_2, \dots) is a disjoint sequence of sets in \mathcal{S} . We just showed that \mathcal{S} is closed under finite, disjoint unions, so $B_n = \bigcup_{i=1}^n A_i \in \mathcal{S}$. But the sequence (B_1, B_2, \dots) is increasing, and hence $\bigcup_{n=1}^{\infty} B_n = \bigcup_{i=1}^{\infty} A_i \in \mathcal{S}$.

The following theorem is known as the *monotone class theorem*, and is due to the mathematician Paul Halmos.

Suppose that \mathcal{A} is an algebra, \mathcal{M} is a monotone class, and $\mathcal{A} \subseteq \mathcal{M}$. Then $\sigma(\mathcal{A}) \subseteq \mathcal{M}$.

Proof

First let $m(\mathcal{A})$ denote the monotone class generated by \mathcal{A} , as defined [above](#). The outline of the proof is to show that $m(\mathcal{A})$ is an algebra, so that by [\(9\)](#), $m(\mathcal{A})$ is a σ -algebra. It then follows that $\sigma(\mathcal{A}) \subseteq m(\mathcal{A}) \subseteq \mathcal{M}$. To show that $m(\mathcal{A})$ is an algebra, we first show that it is closed under complements and then under simple union.

Since $m(\mathcal{A})$ is a monotone class, the collection $m^*(\mathcal{A}) = \{A \subseteq S : A^c \in m(\mathcal{A})\}$ is also a monotone class. Moreover, $\mathcal{A} \subseteq m^*(\mathcal{A})$ so it follows that $m(\mathcal{A}) \subseteq m^*(\mathcal{A})$. Hence if $A \in m(\mathcal{A})$ then $A \in m^*(\mathcal{A})$ so $A^c \in m(\mathcal{A})$. Thus $m(\mathcal{A})$ is closed under complements.

Let $\mathcal{M}_1 = \{A \subseteq S : A \cup B \in m(\mathcal{A}) \text{ for all } B \in \mathcal{A}\}$. Then \mathcal{M}_1 is a monotone class and $\mathcal{A} \subseteq \mathcal{M}_1$ so $m(\mathcal{A}) \subseteq \mathcal{M}_1$. Next let $\mathcal{M}_2 = \{A \subseteq S : A \cup B \in m(\mathcal{A}) \text{ for all } B \in m(\mathcal{A})\}$. Then \mathcal{M}_2 is also a monotone class. Let $A \in \mathcal{A}$. If $B \in m(\mathcal{A})$ then $B \in \mathcal{M}_1$ and hence $A \cup B \in m(\mathcal{A})$. Hence $A \in \mathcal{M}_2$. Thus we have $\mathcal{A} \subseteq \mathcal{M}_2$, so $m(\mathcal{A}) \subseteq \mathcal{M}_2$. Finally, let $A, B \in m(\mathcal{A})$. Then $A \in \mathcal{M}_2$ so $A \cup B \in m(\mathcal{A})$ and therefore $m(\mathcal{A})$ is closed under simple union.

As noted in [\(5\)](#), a σ -algebra is both a π -system and a λ -system. The converse is also true, and is one of the main reasons for studying these structures.

If \mathcal{S} is a π -system and a λ -system then \mathcal{S} is a σ -algebra.

Proof

$S \in \mathcal{S}$, and if $A \in \mathcal{S}$ then $A^c \in \mathcal{S}$ by definition of a λ -system. Thus, all that is left is to show closure under countable unions. Thus, suppose that (A_1, A_2, \dots) is a sequence of sets in \mathcal{S} . Then $A_i^c \in \mathcal{S}$ for each $i \in \mathbb{N}_+$. Since \mathcal{S} is also a π -system, it follows that for each $n \in \mathbb{N}_+$, $B_n = A_n \cap A_1^c \cap \dots \cap A_{n-1}^c \in \mathcal{S}$ (by convention $B_1 = A_1$). But the sequence (B_1, B_2, \dots) is disjoint and has the same union as (A_1, A_2, \dots) . Hence $\bigcup_{i=1}^{\infty} A_i = \bigcup_{i=1}^{\infty} B_i \in \mathcal{S}$.

The importance of π -systems and λ -systems stems in part from Dynkin's π - λ theorem given next. It's named for the mathematician Eugene Dynkin.

Suppose that \mathcal{A} is a π -system of subsets of S , \mathcal{B} is a λ -system of subsets of S , and $\mathcal{A} \subseteq \mathcal{B}$. Then $\sigma(\mathcal{A}) \subseteq \mathcal{B}$.

Proof

Let \mathcal{L} denote the λ -system generated by \mathcal{A} . Then of course $\mathcal{A} \subseteq \mathcal{L} \subseteq \mathcal{B}$. For $A \in \mathcal{L}$, let

$$\mathcal{L}_A = \{B \subseteq S : B \cap A \in \mathcal{L}\} \quad (1.12.2)$$

We will show that \mathcal{L}_A is a λ -system. Note that $S \cap A = A \in \mathcal{L}$ and therefore $S \in \mathcal{L}_A$. Next, suppose that $B_1, B_2 \in \mathcal{L}_A$ and that $B_1 \subseteq B_2$. Then $B_1 \cap A \in \mathcal{L}$ and $B_2 \cap A \in \mathcal{L}$ and $B_1 \cap A \subseteq B_2 \cap A$. Hence $(B_2 \setminus B_1) \cap A = (B_2 \cap A) \setminus (B_1 \cap A) \in \mathcal{L}$. Hence $B_2 \setminus B_1 \in \mathcal{L}_A$. Finally, suppose that $\{B_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{L}_A . Then $B_i \cap A \in \mathcal{L}$ for each $i \in I$, and $\{B_i \cap A : i \in I\}$ is also a disjoint collection. Therefore, $\bigcup_{i \in I} (B_i \cap A) = (\bigcup_{i \in I} B_i) \cap A \in \mathcal{L}$. Hence $\bigcup_{i \in I} B_i \in \mathcal{L}_A$.

Next fix $A \in \mathcal{A}$. If $B \in \mathcal{A}$ then $A \cap B \in \mathcal{A}$, so $A \cap B \in \mathcal{L}$ and hence $B \in \mathcal{L}_A$. But \mathcal{L} is the smallest λ -system containing \mathcal{A} so we have shown that $\mathcal{L} \subseteq \mathcal{L}_A$ for every $A \in \mathcal{A}$. Now fix $B \in \mathcal{L}$. If $A \in \mathcal{A}$ then $B \in \mathcal{L}_A$ so $A \cap B \in \mathcal{L}$ and therefore $A \in \mathcal{L}_B$. Again, \mathcal{L} is the smallest λ -system containing \mathcal{A} so we have now shown that $\mathcal{L} \subseteq \mathcal{L}_B$ for every $B \in \mathcal{L}$. Finally, let $B, C \in \mathcal{L}$. Then $C \in \mathcal{L}_B$ and hence $B \cap C \in \mathcal{L}$. It now follows that \mathcal{L} is a π -system, as well as a λ -system, and therefore by the [theorem above](#), \mathcal{L} is a sigma-algebra. But $\mathcal{A} \subseteq \mathcal{L}$ and hence $\sigma(\mathcal{A}) \subseteq \mathcal{L}$.

Examples and Special Cases

Suppose that S is a set and \mathcal{A} is a finite partition of S . Then $\mathcal{S} = \{\emptyset\} \cup \mathcal{A}$ is a semi-algebra of subsets of S .

Proof

If $A, B \in \mathcal{A}$ then $A \cap B = \emptyset \in \mathcal{S}$. If $A \in \mathcal{S}$ then $A^c = \bigcup \{B \in \mathcal{A} : B \neq A\}$

Euclidean Spaces

The following example is particularly important because it will be used to construct positive measures on \mathbb{R} . Let

$$\mathcal{B} = \{(a, b] : a, b \in \mathbb{R}, a < b\} \cup \{(-\infty, b] : b \in \mathbb{R}\} \cup \{(a, \infty) : a \in \mathbb{R}\} \quad (1.12.3)$$

\mathcal{B} is a semi-algebra of subsets of \mathbb{R} .

Proof

Note that the intersection of two intervals of the type in \mathcal{B} is another interval of this type. The complement of an interval of this type is either another interval of this type or the union of two disjoint intervals of this type.

It follows from the [theorem above](#) that the collection \mathcal{A} of finite disjoint unions of intervals in \mathcal{B} is an algebra. Recall also that $\sigma(\mathcal{B}) = \sigma(\mathcal{A})$ is the Borel σ -algebra of \mathbb{R} , named for Émile Borel. We can generalize all of this to \mathbb{R}^n for $n \in \mathbb{N}_+$

The collection $\mathcal{B}_n = \{\prod_{i=1}^n A_i : A_i \in \mathcal{B} \text{ for each } i \in \{1, 2, \dots, n\}\}$ is a semi-algebra of subsets of \mathbb{R}^n .

Recall also that $\sigma(\mathcal{B}_n)$ is the σ -algebra of Borel sets of \mathbb{R}^n .

Product Spaces

The examples in this discussion are important for constructing positive measures on product spaces.

Suppose that \mathcal{S} is a semi-algebra of subsets of a set S and that \mathcal{T} is a semi-algebra of subsets of a set T . Then

$$\mathcal{U} = \{A \times B : A \in \mathcal{S}, B \in \mathcal{T}\} \quad (1.12.4)$$

is a semi-algebra of subsets of $S \times T$.

Proof

1. Suppose that $A \times B, C \times D \in \mathcal{U}$, so that $A, C \in \mathcal{S}$ and $B, D \in \mathcal{T}$. Recall that $(A \times B) \cap (C \times D) = (A \cap C) \times (B \cap D)$. But $A \cap C \in \mathcal{S}$ and $B \cap D \in \mathcal{T}$ so $(A \times B) \cap (C \times D) \in \mathcal{U}$.
2. Suppose that $A \times B \in \mathcal{U}$ so that $A \in \mathcal{S}$ and $B \in \mathcal{T}$. Then

$$(A \times B)^c = (A^c \times B) \cup (A \times B^c) \cup (A^c \times B^c) \quad (1.12.5)$$

There exists a finite, disjoint collection $\{A_i : i \in I\}$ of sets in \mathcal{S} and a finite, disjoint collection $\{B_j : j \in J\}$ of sets in \mathcal{T} such that $A^c = \bigcup_{i \in I} A_i$ and $B^c = \bigcup_{j \in J} B_j$. Hence

$$(A \times B)^c = \left[\bigcup_{i \in I} (A_i \times B) \right] \cup \left[\bigcup_{j \in J} (A \times B_j) \right] \cup \left[\bigcup_{i \in I} \bigcup_{j \in J} (A_i \times B_j) \right] \quad (1.12.6)$$

All of the product sets in this union are in \mathcal{U} and the product sets are disjoint.

This result extends in a completely straightforward way to a product of a finite number of sets.

Suppose that $n \in \mathbb{N}_+$ and that \mathcal{S}_i is a semi-algebra of subsets of a set S_i for $i \in \{1, 2, \dots, n\}$. Then

$$\mathcal{U} = \left\{ \prod_{i=1}^n A_i : A_i \in \mathcal{S}_i \text{ for all } i \in \{1, 2, \dots, n\} \right\} \quad (1.12.7)$$

is a semi-algebra of subsets of $\prod_{i=1}^n S_i$.

Note that the semi-algebra of products of intervals in \mathbb{R}^n described [above](#) is a special case of this result. For the product of an infinite sequence of sets, the result is bit more tricky.

Suppose that \mathcal{S}_i is a semi-algebra of subsets of a set S_i for $i \in \mathbb{N}_+$. Then

$$\mathcal{U} = \left\{ \prod_{i=1}^{\infty} A_i : A_i \in \mathcal{S}_i \text{ for all } i \in \mathbb{N}_+ \text{ and } A_i = S_i \text{ for all but finitely many } i \in \mathbb{N}_+ \right\} \quad (1.12.8)$$

is a semi-algebra of subsets of $\prod_{i=1}^n S_i$.

Proof

The proof is very much like the previous ones.

1. Suppose that $A = \prod_{i=1}^{\infty} A_i \in \mathcal{U}$ and $B = \prod_{i=1}^{\infty} B_i \in \mathcal{U}$, so that $A_i, B_i \in \mathcal{S}_i$ for $i \in \mathbb{N}_+$ and $A_i = S_i$ for all but finitely many $i \in \mathbb{N}_+$ and $B_i = S_i$ for all but finitely many $i \in \mathbb{N}_+$. Then $A \cap B = \prod_{i=1}^{\infty} (A_i \cap B_i)$. Also, $A_i \cap B_i \in \mathcal{S}_i$ for $i \in \mathbb{N}_+$ and $A_i \cap B_i = S_i$ for all but finitely many $i \in \mathbb{N}_+$. So $A \cap B \in \mathcal{U}$.
2. Suppose that $A = \prod_{i=1}^{\infty} A_i \in \mathcal{U}$, where $A_i \in \mathcal{S}_i$ for $i \in \mathbb{N}_+$ and $A_i = S_i$ for $i > n$, for some $n \in \mathbb{N}_+$. Then $A^c = \bigcup_{j=1}^n B_j$ where

$$B_j = A_1 \times \cdots \times A_{j-1} \times A_j^c \times S_{j+1} \times S_{j+2} \times \cdots, \quad j \in \{1, 2, \dots, n\} \quad (1.12.9)$$

Note that the product sets in this union are disjoint. But for each $j \in \{1, 2, \dots, n\}$ there exists a finite disjoint collection $\{C_{j,k} : k \in K_j\}$ such that $A_j^c = \bigcup_{k \in K_j} C_{j,k}$. Substituting and distributing then gives A^c as a finite, disjoint union of sets in \mathcal{U} .

Note that this result would not be true with $\mathcal{U} = \{\prod_{i=1}^{\infty} A_i : A_i \in \mathcal{S}_i \text{ for all } i \in \mathbb{N}_+\}$. In general, the complement of a set in \mathcal{U} cannot be written as a finite disjoint union of sets in \mathcal{U} .

This page titled [1.12: Special Set Structures](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

2: Probability Spaces

The basic topics in this chapter are fundamental to probability theory, and should be accessible to new students of probability. We start with the paradigm of the random experiment and its mathematical model, the probability space. The main objects in this model are sample spaces, events, random variables, and probability measures. We also study several concepts of fundamental importance: conditional probability and independence.

The advanced topics can be skipped if you are a new student of probability, or can be studied later, as the need arises. These topics include the convergence of random variables, the measure-theoretic foundations of probability theory, and the existence and construction of probability measures and random processes.

[2.1: Random Experiments](#)

[2.2: Events and Random Variables](#)

[2.3: Probability Measures](#)

[2.4: Conditional Probability](#)

[2.5: Independence](#)

[2.6: Convergence](#)

[2.7: Measure Spaces](#)

[2.8: Existence and Uniqueness](#)

[2.9: Probability Spaces Revisited](#)

[2.10: Stochastic Processes](#)

[2.11: Filtrations and Stopping Times](#)

This page titled [2: Probability Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.1: Random Experiments

Experiments

Probability theory is based on the paradigm of a *random experiment*; that is, an experiment whose outcome cannot be predicted with certainty, before the experiment is run. In *classical* or *frequency-based* probability theory, we also assume that the experiment can be repeated indefinitely under essentially the same conditions. The repetitions can be *in time* (as when we toss a single coin over and over again) or *in space* (as when we toss a bunch of similar coins all at once). The repeatability assumption is important because the classical theory is concerned with the long-term behavior as the experiment is replicated. By contrast, *subjective* or *belief-based* probability theory is concerned with measures of belief about what will happen when we run the experiment. In this view, repeatability is a less crucial assumption. In any event, a complete description of a random experiment requires a careful definition of precisely what information about the experiment is being recorded, that is, a careful definition of what constitutes an *outcome*.

The term *parameter* refers to a non-random quantity in a model that, once chosen, remains constant. Many probability models of random experiments have one or more parameters that can be adjusted to fit the physical experiment being modeled.

The subjects of probability and statistics have an inverse relationship of sorts. In probability, we start with a completely specified mathematical model of a random experiment. Our goal is perform various computations that help us understand the random experiment, help us predict what will happen when we run the experiment. In statistics, by contrast, we start with an incompletely specified mathematical model (one or more parameters may be unknown, for example). We run the experiment to collect data, and then use the data to draw inferences about the unknown factors in the mathematical model.

Compound Experiments

Suppose that we have n experiments (E_1, E_2, \dots, E_n) . We can form a new, compound experiment by performing the n experiments in sequence, E_1 first, and then E_2 and so on, independently of one another. The term *independent* means, intuitively, that the outcome of one experiment has no influence over any of the other experiments. We will make the term mathematically precise later.

In particular, suppose that we have a basic experiment. A fixed number (or even an infinite number) of independent replications of the basic experiment is a new, compound experiment. Many experiments turn out to be compound experiments and moreover, as noted above, (classical) probability theory itself is based on the idea of replicating an experiment.

In particular, suppose that we have a simple experiment with two outcomes. Independent replications of this experiment are referred to as Bernoulli trials, named for Jacob Bernoulli. This is one of the simplest, but most important models in probability. More generally, suppose that we have a simple experiment with k possible outcomes. Independent replications of this experiment are referred to as multinomial trials.

Sometimes an experiment occurs in well-defined stages, but in a *dependent* way, in the sense that the outcome of a given stage is influenced by the outcomes of the previous stages.

Sampling Experiments

In most statistical studies, we start with a population of objects of interest. The objects may be people, memory chips, or acres of corn, for example. Usually there are one or more numerical measurements of interest to us—the height and weight of a person, the lifetime of a memory chip, the amount of rain, amount of fertilizer, and yield of an acre of corn.

Although our interest is in the entire population of objects, this set is usually too large and too amorphous to study. Instead, we collect a random sample of objects from the population and record the measurements of interest of for each object in the sample.

There are two basic types of sampling. If we sample *with replacement*, each item is replaced in the population before the next draw; thus, a single object may occur several times in the sample. If we sample *without replacement*, objects are not replaced in the population. The chapter on Finite Sampling Models explores a number of models based on sampling from a finite population.

Sampling with replacement can be thought of as a compound experiment, based on independent replications of the simple experiment of drawing a single object from the population and recording the measurements of interest. Conversely, a compound experiment that consists of n independent replications of a simple experiment can usually be thought of as a sampling experiment.

On the other hand, sampling without replacement is an experiment that consists of dependent stages, because the population changes with each draw.

Examples and Applications

Probability theory is often illustrated using simple devices from games of chance: coins, dice, card, spinners, urns with balls, and so forth. Examples based on such devices are pedagogically valuable because of their simplicity and conceptual clarity. On the other hand, it would be a terrible shame if you were to think that probability is *only* about gambling and games of chance. Rather, try to see problems involving coins, dice, etc. as metaphors for more complex and realistic problems.

Coins and Dice

In terms of probability, the important fact about a coin is simply that when tossed it lands on one side or the other. Coins in Western societies, dating to antiquity, usually have the head of a prominent person engraved on one side and something of lesser importance on the other. In non-Western societies, coins often did not have a head on either side, but did have distinct engravings on the two sides, one typically more important than the other. Nonetheless, *heads* and *tails* are the ubiquitous terms used in probability theory to distinguish the *front* or *obverse* side of the coin from the *back* or *reverse* side of the coin.



Figure 2.1.1: Obverse and reverse sides of a Roman coin, about 241 CE, from [Wikipedia](#)

Consider the *coin experiment* of tossing a coin n times and recording the score (1 for heads or 0 for tails) for each toss.

1. Identify a parameter of the experiment.
2. Interpret the experiment as a compound experiment.
3. Interpret the experiment as a sampling experiment.
4. Interpret the experiment as n Bernoulli trials.

Answer

1. The number of coins n is the parameter.
2. The experiment consists of n independent replications of the simple experiment of tossing the coin one time.
3. The experiment can be thought of as selecting a sample of size n with replacement from the population $\{0, 1\}$.
4. There are two outcomes on each toss and the tosses are independent.

In the simulation of the coin experiment, set $n = 5$. Run the simulation 100 times and observe the outcomes.

Dice are randomizing devices that, like coins, date to antiquity and come in a variety of sizes and shapes. Typically, the faces of a die have numbers or other symbols engraved on them. Again, the important fact is that when a die is thrown, a unique face is chosen (usually the upward face, but sometimes the downward one). For more on dice, see the introductory section in the chapter on Games of Chance.

Consider the *dice experiment* of throwing a k -sided die (with faces numbered 1 to k), n times and recording the scores for each throw.

1. Identify the parameters of the experiment.
2. Interpret the experiment as a compound experiment.
3. Interpret the experiment as a sampling experiment.
4. Identify the experiment as n multinomial trials.

Answer

1. The parameters are the number of dice n and the number of faces k .
2. The experiment consists of n independent replications of the simple experiment of throwing one die.
3. The experiment can be thought of as selecting a sample of size n with replacement from the population $\{1, 2, \dots, k\}$.
4. The same k outcomes occur for each die the throws are independent.

In reality, most dice are *Platonic solids* (named for Plato of course) with 4, 6, 8, 12, or 20 sides. The six-sided die is the *standard die*.



Figure 2.1.2: Blue Platonic dice

In the simulation of the dice experiment, set $n = 5$. Run the simulation 100 times and observe the outcomes.

In the *die-coin experiment*, a standard die is thrown and then a coin is tossed the number of times shown on the die. The sequence of coin scores is recorded (1 for heads and 0 for tails). Interpret the experiment as a compound experiment.

Answer

The first stage consists rolling the die and the second stage consists of tossing the coin. The stages are dependent because the number of tosses depends on the outcome of the die throw.

Note that this experiment can be obtained by *randomizing* the parameter n in the basic coin experiment in (1).

Run the simulation of the die-coin experiment 100 times and observe the outcomes.

In the *coin-die experiment*, a coin is tossed. If the coin lands heads, a red die is thrown and if the coin lands tails, a green die is thrown. The coin score (1 for heads and 0 for tails) and the die score are recorded. Interpret the experiment as a compound experiment.

Answer

The first stage consists of tossing the coin and the second stage consists of rolling the die. The stages are dependent because different dice (that may behave differently) are thrown, depending on the outcome of the coin toss.

Run the simulation of the coin-die experiment 100 times and observe the outcomes.

Cards

Playing cards, like coins and dice, date to antiquity. From the point of view of probability, the important fact is that a playing card encodes a number of properties or attributes on the front of the card that are hidden on the back of the card. (Later in this chapter, these properties will become random variables.) In particular, a standard *card deck* can be modeled by the Cartesian product set

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (2.1.1)$$

where the first coordinate encodes the *denomination* or *kind* (ace, 2–10, jack, queen, king) and where the second coordinate encodes the *suit* (clubs, diamonds, hearts, spades). Sometimes we represent a card as a *string* rather than an ordered pair (for example $q\heartsuit$ rather than (q, \heartsuit) for the queen of hearts). Some other properties, derived from the two main ones, are *color* (diamonds and hearts are red, clubs and spades are black), *face* (jacks, queens, and kings have faces, the other cards do not), and *suit order* (from least to highest rank: $(\clubsuit, \diamondsuit, \heartsuit, \spadesuit)$).

Consider the card experiment that consists of dealing n cards from a standard deck (without replacement).

1. Identify a parameter of the experiment.
2. Interpret the experiment as a compound experiment.
3. Interpret the experiment as a sampling experiment.

Answer

1. The parameter is n , the number of cards dealt.
2. At each stage, we draw a card from a deck, but the deck changes from one draw to the next, so the stages are dependent.
3. The experiment is to select a sample of size n from the population D , without replacement.

In the simulation of the card experiment, set $n = 5$. Run the simulation 100 times and observe the outcomes.

The special case $n = 5$ is the *poker experiment* and the special case $n = 13$ is the *bridge experiment*.

Open each of the following to see depictions of card playing in some famous paintings.

1. Cheat with the Ace of Clubs by Georges de La Tour
2. The Cardsharps by Michelangelo Carravagio
3. The Card Players by Paul Cézanne
4. His Station and Four Aces by CM Coolidge
5. Waterloo by CM Coolidge

Urn Models

Urn models are often used in probability as simple metaphors for sampling from a finite population.

An urn contains m distinct balls, labeled from 1 to m . The experiment consists of selecting n balls from the urn, without replacement, and recording the sequence of ball numbers.

1. Identify the parameters of the experiment.
2. Interpret the experiment as a compound experiment.
3. Interpret the experiment as a sampling experiment.

Answer

1. The parameters are the number of balls m and the sample size n .
2. At each stage, we draw a ball from the urn, but the contents of the urn change from one draw to the next, so the stages are dependent
3. The experiment is to select a sample of size n from the balls in the urn (the population), without replacement.

Consider the basic urn model of the previous exercise. Suppose that r of the m balls are red and the remaining $m - r$ balls are green. Identify an additional parameter of the model. This experiment is a metaphor for sampling from a general *dichotomous* population

Answer

The parameters are the population size m , the sample size n , and the number of red balls r .

In the simulation of the urn experiment, set $m = 100$, $r = 40$, and $n = 25$. Run the experiment 100 times and observe the results.

An urn initially contains m balls; r are red and $m - r$ are green. A ball is selected from the urn and removed, and then replaced with k balls of the same color. The process is repeated. This is known as *Pólya's urn model*, named after George Pólya.

1. Identify the parameters of the experiment.
2. Interpret the case $k = 0$ as a sampling experiment.
3. Interpret the case $k = 1$ as a sampling experiment.

Answer

1. The parameters are the population size m , the initial number of red balls r , and the number new balls added k .
2. When $k = 0$, each ball drawn is removed and no new balls are added, so the experiment is to select a sample of size n from the urn, without replacement.
3. When $k = 1$, each ball drawn is replaced with another ball of the same color. So at least in terms of the colors of the balls, the experiment is equivalent to selecting a sample of size n from the urn, with replacement.

Open the image of the painting *Allegory of Fortune* by Dosso Dossi. Presumably the young man has chosen lottery tickets from an urn.

Buffon's Coin Experiment

Buffon's coin experiment consists of tossing a coin with radius $r \leq \frac{1}{2}$ on a floor covered with square tiles of side length 1. The coordinates of the center of the coin are recorded, relative to axes through the center of the square, parallel to the sides. The experiment is named for comte de Buffon.

1. Identify a parameter of the experiment
2. Interpret the experiment as a compound experiment.
3. Interpret the experiment as sampling experiment.

Answer

1. The parameter is the coin radius r .
2. The experiment can be thought of as selecting the coordinates of the coin center independently of one another.
3. The experiment is equivalent to selecting a sample of size 2 from the population $[-\frac{1}{2}, \frac{1}{2}]$, with replacement.

In the simulation of Buffon's coin experiment, set $r = 0.1$. Run the experiment 100 times and observe the outcomes.

Reliability

In the usual model of *structural reliability*, a system consists of n components, each of which is either *working* or *failed*. The states of the components are uncertain, and hence define a random experiment. The system as a whole is also either working or failed, depending on the states of the components and how the components are connected. For example, a *series system* works if and only if each component works, while a *parallel system* works if and only if at least one component works. More generally, a *k out of n system* works if at least k components work.

Consider the k out of n reliability model.

1. Identify two parameters.
2. What value of k gives a series system?
3. What value of k gives a parallel system?

Answer

1. The parameters are k and n
2. $k = n$ gives a series system.
3. $k = 1$ gives a parallel system.

The reliability model above is a *static* model. It can be extended to a *dynamic* model by assuming that each component is initially working, but has a random time until failure. The system as a whole would also have a random time until failure that would depend on the component failure times and the structure of the system.

Genetics

In ordinary sexual reproduction, the genetic material of a child is a random combination of the genetic material of the parents. Thus, the birth of a child is a random experiment with respect to outcomes such as eye color, hair type, and many other physical traits. We are often particularly interested in the random transmission of traits and the random transmission of genetic disorders.

For example, let's consider an overly simplified model of an *inherited trait* that has two possible states (*phenotypes*), say a pea plant whose pods are either green or yellow. The term *allele* refers to alternate forms of a particular gene, so we are assuming that there is a gene that determines pod color, with two alleles: g for green and y for yellow. A pea plant has two alleles for the trait (one from each parent), so the possible *genotypes* are

- gg , alleles for green pods from each parent.
- gy , an allele for green pods from one parent and an allele for yellow pods from the other (we usually cannot observe which parent contributed which allele).
- yy , alleles for yellow pods from each parent.

The genotypes gg and yy are called *homozygous* because the two alleles are the same, while the genotype gy is called *heterozygous* because the two alleles are different. Typically, one of the alleles of the inherited trait is *dominant* and the other *recessive*. Thus, for example, if g is the dominant allele for pod color, then a plant with genotype gg or gy has green pods, while a plant with genotype

yy has yellow pods. Genes are passed from parent to child in a random manner, so each new plant is a random experiment with respect to pod color.

Pod color in peas was actually one of the first examples of an inherited trait studied by Gregor Mendel, who is considered the father of modern genetics. Mendel also studied the color of the flowers (yellow or purple), the length of the stems (short or long), and the texture of the seeds (round or wrinkled).

For another example, the *ABO blood type* in humans is controlled by three alleles: *a*, *b*, and *o*. Thus, the possible genotypes are *aa*, *ab*, *ao*, *bb*, *bo* and *oo*. The alleles *a* and *b* are *co-dominant* and *o* is *recessive*. Thus there are four possible blood types (phenotypes):

- Type *A*: genotype *aa* or *ao*
- Type *B*: genotype *bb* or *bo*
- Type *AB*: genotype *ab*
- type *O*: genotype *oo*

Of course, blood may be typed in much more extensive ways than the simple *ABO* typing. The *RH* factor (positive or negative) is the most well-known example.

For our third example, consider a *sex-linked* hereditary disorder in humans. This is a disorder due to a defect on the *X chromosome* (one of the two chromosomes that determine gender). Suppose that *h* denotes the healthy allele and *d* the defective allele for the gene linked to the disorder. Women have two X chromosomes, and typically *d* is *recessive*. Thus, a woman with genotype *hh* is completely normal with respect to the condition; a woman with genotype *hd* does not have the disorder, but is a *carrier*, since she can pass the defective allele to her children; and a woman with genotype *dd* has the disorder. A man has only one X chromosome (his other sex chromosome, the *Y chromosome*, typically plays no role in the disorder). A man with genotype *h* is normal and a man with genotype *d* has the disorder. Examples of sex-linked hereditary disorders are *dichromatism*, the most common form of color-blindness, and *hemophilia*, a bleeding disorder. Again, genes are passed from parent to child in a random manner, so the birth of a child is a random experiment in terms of the disorder.

Point Processes

There are a number of important processes that generate “random points in time”. Often the random points are referred to as *arrivals*. Here are some specific examples:

- times that a piece of radioactive material emits elementary particles
- times that customers arrive at a store
- times that requests arrive at a web server
- failure times of a device

To formalize an experiment, we might record the number of arrivals during a specified interval of time or we might record the times of successive arrivals.

There are other processes that produce “random points in space”. For example,

- flaws in a piece of sheet metal
- errors in a string of symbols (in a computer program, for example)
- raisins in a cake
- misprints on a page
- stars in a region of space

Again, to formalize an experiment, we might record the number of points in a given region of space.

Statistical Experiments

In 1879, Albert Michelson constructed an experiment for measuring the speed of light with an interferometer. The velocity of light data set contains the results of 100 repetitions of Michelson's experiment. Explore the data set and explain, in a general way, the variability of the data.

Answer

The variability is due to measurement and other experimental errors beyond the control of Michelson.

In 1998, two students at the University of Alabama in Huntsville designed the following experiment: purchase a bag of M&Ms (of a specified advertised size) and record the counts for red, green, blue, orange, and yellow candies, and the net weight (in grams). Explore the M&M data set and explain, in a general way, the variability of the data.

Answer

The variability in weight is due to measurement error on the part of the students and to manufacturing errors on the part of the company. The variability in color counts is less clear and may be due to purposeful randomness on the part of the company.

In 1999, two researchers at Belmont University designed the following experiment: capture a cicada in the Middle Tennessee area, and record the body weight (in grams), the wing length, wing width, and body length (in millimeters), the gender, and the species type. The cicada data set contains the results of 104 repetitions of this experiment. Explore the cicada data and explain, in a general way, the variability of the data.

Answer

The variability in body measurements is due to differences in the three species, to all sorts of environmental factors, and to measurement errors by the researchers.

On June 6, 1761, James Short made 53 measurements of the parallax of the sun, based on the transit of Venus. Explore the Short data set and explain, in a general way, the variability of the data.

Answer

The variability is due to measurement and other experimental errors beyond the control of Short.

In 1954, two massive field trials were conducted in an attempt to determine the effectiveness of the new vaccine developed by Jonas Salk for the prevention of polio. In both trials, a *treatment group* of children were given the vaccine while a *control group* of children were not. The incidence of polio in each group was measured. Explore the polio field trial data set and explain, in a general way, the underlying random experiment.

Answer

The basic random experiment is to observe whether a given child, in the treatment group or control group, comes down with polio in a specified period of time. Presumably, a lower incidence of polio in the treatment group compared with the control group would be evidence that the vaccine was effective.

Each year from 1969 to 1972 a lottery was held in the US to determine who would be drafted for military service. Essentially, the lottery was a *ball and urn model* and became famous because many believed that the process was not sufficiently random. Explore the Vietnam draft lottery data set and speculate on how one might judge the degree of randomness.

Answer

This is a difficult problem, but presumably in a sufficiently random lottery, one would not expect to see dates in the same month clustered too closely together. Observing such clustering, then, would be evidence that the lottery was not random.

Deterministic Versus Probabilistic Models

One could argue that some of the examples discussed above are inherently *deterministic*. In tossing a coin, for example, if we know the initial conditions (involving position, velocity, rotation, etc.), the forces acting on the coin (gravity, air resistance, etc.), and the makeup of the coin (shape, mass density, center of mass, etc.), then the laws of physics should allow us to predict precisely how the coin will land. This is true in a technical, theoretical sense, but false in a very real sense. Coins, dice, and many more complicated and important systems are *chaotic* in the sense that the outcomes of interest depend in a very sensitive way on the initial conditions and other parameters. In such situations, it might well be impossible to ever know the initial conditions and forces accurately enough to use deterministic methods.

In the coin experiment, for example, even if we strip away most of the real world complexity, we are still left with an essentially random experiment. Joseph Keller in his article "[The Probability of Heads](#)" deterministically analyzed the toss of a coin under a number of ideal assumptions:

1. The coin is a perfect circle and has negligible thickness

2. The center of gravity of the coin is the geometric center.
3. The coin is initially heads up and is given an initial upward velocity u and angular velocity ω .
4. In flight, the coin rotates about a horizontal axis along a diameter of the coin.
5. In flight, the coin is governed only by the force of gravity. All other possible forces (air resistance or wind, for example) are neglected.
6. The coin does not bounce or roll after landing (as might be the case if it lands in sand or mud).

Of course, few of these ideal assumptions are valid for real coins tossed by humans. Let $t = u/g$ where g is the acceleration of gravity (in appropriate units). Note that the t just has units of time (in seconds) and hence is independent of how distance is measured. The scaled parameter t actually represents the time required for the coin to reach its maximum height.

Keller showed that the regions of the parameter space (t, ω) where the coin lands either heads up or tails up are separated by the curves

$$\omega = \left(2n \pm \frac{1}{2}\right) \frac{\pi}{2t}, \quad n \in \mathbb{N} \quad (2.1.2)$$

The parameter n is the total number of revolutions in the toss. A plot of some of these curves is given below. The largest region, in the lower left corner, corresponds to the event that the coin does not complete even one rotation, and so of course lands heads up, just as it started. The next region corresponds to one rotation, with the coin landing tails up. In general, the regions alternate between heads and tails.




Figure 2.1.3: Regions of heads and tails

The important point, of course, is that for even moderate values of t and ω , the curves are very close together, so that a small change in the initial conditions can easily shift the outcome from heads up to tails up or conversely. As noted in Keller's article, the probabilist and statistician Persi Diaconis determined experimentally that typical values of the initial conditions for a real coin toss are $t = \frac{1}{4}$ seconds and $\omega = 76\pi \approx 238.6$ radians per second. These values correspond to $n = 19$ revolutions in the toss. Of course, this parameter point is far beyond the region shown in our graph, in a region where the curves are exquisitely close together.

This page titled [2.1: Random Experiments](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.2: Events and Random Variables

The purpose of this section is to study two basic types of objects that form part of the model of a random experiment. If you are a new student of probability, just ignore the measure-theoretic terminology and skip the technical details.

Sample Spaces

The Set of Outcomes

Recall that in a random experiment, the outcome cannot be predicted with certainty, before the experiment is run. On the other hand:

We assume that we can identify a fixed set S that includes all possible outcomes of a random experiment. This set plays the role of the universal set when modeling the experiment.

For simple experiments, S may be precisely the set of possible outcomes. More often, for complex experiments, S is a mathematically convenient set that includes the possible outcomes and perhaps other elements as well. For example, if the experiment is to throw a standard die and record the score that occurs, we would let $S = \{1, 2, 3, 4, 5, 6\}$, the set of possible outcomes. On the other hand, if the experiment is to capture a cicada and measure its body weight (in milligrams), we might conveniently take $S = [0, \infty)$, even though most elements of this set are impossible (we hope!). The problem is that we may not know exactly the outcomes that are possible. Can a light bulb burn without failure for one thousand hours? For one thousand days? for one thousand years?

Often the outcome of a random experiment consists of one or more real measurements, and thus, the S consists of all possible measurement sequences, a subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$. More generally, suppose that we have n experiments and that S_i is the set of outcomes for experiment $i \in \{1, 2, \dots, n\}$. Then the Cartesian product $S_1 \times S_2 \times \dots \times S_n$ is the natural set of outcomes for the compound experiment that consists of performing the n experiments in sequence. In particular, if we have a basic experiment with S as the set of outcomes, then S^n is the natural set of outcomes for the compound experiment that consists of n replications of the basic experiment. Similarly, if we have an infinite sequence of experiments and S_i is the set of outcomes for experiment $i \in \mathbb{N}_+$, then $S_1 \times S_2 \times \dots$ is the natural set of outcomes for the compound experiment that consists of performing the given experiments in sequence. In particular, the set of outcomes for the compound experiment that consists of indefinite replications of a basic experiment is $S^\infty = S \times S \times \dots$. This is an essential special case, because (classical) probability theory is based on the idea of replicating a given experiment.

Events

Consider again a random experiment with S as the set of outcomes. Certain subsets of S are referred to as *events*. Suppose that $A \subseteq S$ is a given event, and that the experiment is run, resulting in outcome $s \in S$.

1. If $s \in A$ then we say that A occurs.
2. If $s \notin A$ then we say that A does not occur.

Intuitively, you should think of an event as a meaningful *statement* about the experiment: every such statement translates into an event, namely the set of outcomes for which the statement is true. In particular, S itself is an event; by definition it *always* occurs. At the other extreme, the empty set \emptyset is also an event; by definition it *never* occurs.

For a note on terminology, recall that a *mathematical space* consists of a set together with other mathematical structures defined on the set. An example you may be familiar with is a *vector space*, which consists of a set (the vectors) together with the operations of addition and scalar multiplication. In probability theory, many authors use the term *sample space* for the set of outcomes of a random experiment, but here is the more careful definition:

The *sample space* of an experiment is (S, \mathcal{S}) where S is the set of outcomes and \mathcal{S} is the collection of events.

Details

Sometimes not every subset of S can be allowed as an event, but the collection of events \mathcal{S} is required to be a σ -algebra, so that the sample space (S, \mathcal{S}) is a measurable space. The axioms of a σ -algebra ensure that new sets that are constructed in a

reasonable way from given events, using the set operations, are themselves valid events. Most of the sample spaces that occur in elementary probability fall into two general categories.

1. *Discrete*: S is countable and $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S . In this case, the sample space (S, \mathcal{S}) is *discrete*.
2. *Euclidean*: S is a measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$ and \mathcal{S} is the collection of measurable subsets of S .

In (b), the measurable subsets of \mathbb{R}^n include all of the sets encountered in calculus and in standard applications of probability theory, and many more besides. Nonetheless, for technical reasons, certain very weird subsets must be excluded. Typically S is a set defined by a finite number of inequalities involving elementary functions.

The Algebra of Events

The standard algebra of sets leads to a grammar for discussing random experiments and allows us to construct new events from given events. In the following results, suppose that S is the set of outcomes of a random experiment, and that A and B are events.

$A \subseteq B$ if and only if the occurrence of A *implies* the occurrence of B .

Proof

Recall that \subseteq is the subset relation. So by definition, $A \subseteq B$ means that $s \in A$ implies $s \in B$.

$A \cup B$ is the event that occurs if and only if A occurs *or* B occurs.

Proof

Recall that $A \cup B$ is the *union* of A and B . So by definition, $s \in A \cup B$ if and only if $s \in A$ or $s \in B$.

$A \cap B$ is the event that occurs if and only if A occurs *and* B occurs.

Proof

Recall that $A \cap B$ is the *intersection* of A and B . So by definition, $s \in A \cap B$ if and only if $s \in A$ and $s \in B$.

A and B are disjoint if and only if they are *mutually exclusive*; they cannot both occur on the same run of the experiment.

Proof

By definition, A and B *disjoint* means that $A \cap B = \emptyset$.

A^c is the event that occurs if and only if A does *not* occur.

Proof

Recall that A^c is the *complement* of A , so $s \in A^c$ if and only if $s \notin A$.

$A \setminus B$ is the event that occurs if and only if A occurs *and* B does *not* occur.

Proof

Recall that $A \setminus B = A \cap B^c$. Hence $s \in A \setminus B$ if and only if $s \in A$ and $s \notin B$.

$(A \cap B^c) \cup (B \cap A^c)$ is the event that occurs if and only if *one but not both* of the given events occurs.

Proof

The events in the union are disjoint. So for s is in the given event if and only if either $s \in A$ and $s \notin B$, or $s \in B$ and $s \notin A$.

Recall that the event in (10) is the *symmetric difference* of A and B , and is sometimes denoted $A \Delta B$. This event corresponds to *exclusive or*, as opposed to the ordinary union $A \cup B$ which corresponds to *inclusive or*.

$(A \cap B) \cup (A^c \cap B^c)$ is the event that occurs if and only if *both or neither* of the given events occurs.

Proof

The events in the union are disjoint. Thus s is in the given event if and only if either $s \in A$ and $s \in B$, or $s \notin A$ and $s \notin B$.

In the Venn diagram app, observe the diagram of each of the 16 events that can be constructed from A and B .

Suppose now that $\mathcal{A} = \{A_i : i \in I\}$ is a collection of events for the random experiment, where I is a countable index set.

$\bigcup_{i \in I} A_i$ is the event that occurs if and only if *at least one* event in the collection occurs.

Proof

Note that $s \in \bigcup_{i \in I} A_i$ if and only if $s \in A_i$ for some $i \in I$.

$\bigcap_{i \in I} A_i$ is the event that occurs if and only if *every* event in the collection occurs:

Proof

Note that $s \in \bigcap_{i \in I} A_i$ if and only if $s \in A_i$ for every $i \in I$.

\mathcal{A} is a pairwise disjoint collection if and only if the events are *mutually exclusive*; at most one of the events could occur on a given run of the experiment.

Proof

By definition, $A_i \cap A_j = \emptyset$ for distinct $i, j \in I$.

Suppose now that (A_1, A_2, \dots) is an infinite sequence of events.

$\bigcup_{n=1}^{\infty} \bigcup_{i=n}^{\infty} A_i$ is the event that occurs if and only if *infinitely many* of the given events occur. This event is sometimes called the *limit superior* of (A_1, A_2, \dots) .

Proof

Note that s is in the given event if and only if for every $n \in \mathbb{N}_+$ there exists $i \in \mathbb{N}_+$ with $i \geq n$ such that $s \in A_i$. In turn this means that $s \in A_i$ for infinitely many $i \in I$.

$\bigcup_{n=1}^{\infty} \bigcap_{i=n}^{\infty} A_i$ is the event that occurs if and only if *all but finitely many* of the given events occur. This event is sometimes called the *limit inferior* of (A_1, A_2, \dots) .

Proof

Note that s is in the given event if and only if there exists $n \in \mathbb{N}_+$ such that $s \in A_i$ for every $i \in \mathbb{N}_+$ with $i \geq n$. In turn, this means that $s \in A_i$ for all but finitely many $i \in I$.

Limit superiors and inferiors are discussed in more detail in the section on convergence.

Random Variables

Intuitively, a *random variable* is a measurement of interest in the context of the experiment. Simple examples include the number of heads when a coin is tossed several times, the sum of the scores when a pair of dice are thrown, the lifetime of a device subject to random stress, the weight of a person chosen from a population. Many more examples are given below in the [exercises below](#). Mathematically, a random variable is a function defined on the set of outcomes.

A function X from S into a set T is a *random variable* for the experiment with values in T .

Details

The set T will also come with a σ -algebra \mathcal{T} of admissible subsets, so that (T, \mathcal{T}) is a measurable space, just like (S, \mathcal{S}) . The function X is required to be measurable, an assumption which ensures that meaningful statements involving X define events. In the discussion below, all subsets of T are assumed to be in \mathcal{T} .

Probability has its own notation, very different from other branches of mathematics. As a case in point, random variables, even though they are functions, are usually denoted by capital letters near the end of the alphabet. The use of a letter near the end of the alphabet is intended to emphasize the idea that the object is a *variable* in the context of the experiment. The use of a capital letter is intended to emphasize the fact that it is not an ordinary *algebraic* variable to which we can assign a specific value, but rather a

random variable whose value is indeterminate until we run the experiment. Specifically, when we run the experiment an outcome $s \in S$ occurs, and random variable X takes the value $X(s) \in T$.




Figure 2.2.1: A random variable as a function defined on the set of outcomes.

If $B \subseteq T$, we use the notation $\{X \in B\}$ for the inverse image $\{s \in S : X(s) \in B\}$, rather than $X^{-1}(B)$. Again, the notation is more natural since we think of X as a variable in the experiment. Think of $\{X \in B\}$ as a *statement* about X , which then translates into the event $\{s \in S : X(s) \in B\}$




Figure 2.2.2: The event $\{X \in B\}$ corresponding to $B \subseteq T$

Again, every statement about a random variable X with values in T translates into an inverse image of the form $\{X \in B\}$ for some $B \subseteq T$. So, for example, if $x \in T$ then $\{X = x\} = \{X \in \{x\}\} = \{s \in S : X(s) = x\}$. If X is a real-valued random variable and $a, b \in \mathbb{R}$ with $a < b$ then $\{a \leq X \leq b\} = \{X \in [a, b]\} = \{s \in S : a \leq X(s) \leq b\}$.

Suppose that X is a random variable taking values in T , and that $A, B \subseteq T$. Then

1. $\{X \in A \cup B\} = \{X \in A\} \cup \{X \in B\}$
2. $\{X \in A \cap B\} = \{X \in A\} \cap \{X \in B\}$
3. $\{X \in A \setminus B\} = \{X \in A\} \setminus \{X \in B\}$
4. $A \subseteq B \implies \{X \in A\} \subseteq \{X \in B\}$
5. If A and B are disjoint, then so are $\{X \in A\}$ and $\{X \in B\}$.

Proof

This is a restatement of the fact that inverse images of a function preserve the set operations; only the notation changes (and is simpler).

1. $s \in \{X \in A \cup B\}$ if and only if $X(s) \in A \cup B$ if and only if $X(s) \in A$ or $X(s) \in B$ if and only if $s \in \{X \in A\}$ or $s \in \{X \in B\}$ if and only if $s \in \{X \in A\} \cup \{X \in B\}$.
2. The proof is exactly the same as (a), with *and* replacing *or*.
3. The proof is also exactly the same as (a), with *but not* replacing *or*.
4. If $s \in \{X \in A\}$ then $X(s) \in A$ so $X(s) \in B$ and hence $s \in \{X \in B\}$.
5. This follows from part (b).

As with a general function, the result in part (a) holds for the union of a countable collection of subsets, and the result in part (b) holds for the intersection of a countable collection of subsets. No new ideas are involved; only the notation is more complicated.

Often, a random variable takes values in a subset T of \mathbb{R}^k for some $k \in \mathbb{N}_+$. We might express such a random variable as $\mathbf{X} = (X_1, X_2, \dots, X_k)$ where X_i is a real-valued random variable for each $i \in \{1, 2, \dots, k\}$. In this case, we usually refer to \mathbf{X} as a *random vector*, to emphasize its higher-dimensional character. A random variable can have an even more complicated structure. For example, if the experiment is to select n objects from a population and record various real measurements for each object, then the outcome of the experiment is a vector of vectors: $\mathbf{X} = (X_1, X_2, \dots, X_n)$ where X_i is the vector of measurements for the i th object. There are other possibilities; a random variable could be an infinite sequence, or could be set-valued. Specific examples are given in the computational exercises below. However, the important point is simply that a random variable is a function defined on the set of outcomes S .

The outcome of the experiment itself can be thought of as a random variable. Specifically, let $T = S$ and let X denote the identity function on S so that $X(s) = s$ for $s \in S$. Then trivially X is a random variable, and the events that can be defined in terms of X

are simply the original events of the experiment. That is, if A is an event then $\{X \in A\} = A$. Conversely, every random variable effectively defines a new random experiment.

In the general setting above, a random variable X defines a new random experiment with T as the new set of outcomes and subsets of T as the new collection of events.

Details

Technically, the σ -algebra \mathcal{T} would be the new collection of events.

In fact, often a random experiment is modeled by specifying the random variables of interest, in the language of the experiment. Then, a mathematical definition of the random variables specifies the sample space. A function (or transformation) of a random variable defines a new random variable.

Suppose that X is a random variable for the experiment with values in T and that g is a function from T into another set U . Then $Y = g(X)$ is a random variable with values in U .

Details

Technically, T and U both come with σ -algebras of admissible subsets \mathcal{T} and \mathcal{U} , respectively. The function g , just like the function X , is required to be measurable. This assumption ensures that $Y = g(X)$ is a measurable function from S into U , and hence is a valid random variable.

Note that, as functions, $g(X) = g \circ X$, the composition of g with X . But again, thinking of X and Y as *variables* in the context of the experiment, the notation $Y = g(X)$ is much more natural.

Indicator Variables

For an event A , the indicator function of A is called the *indicator variable* of A .

The value of this random variables tells us whether or not A has occurred:

$$\mathbf{1}_A = \begin{cases} 1, & A \text{ occurs} \\ 0, & A \text{ does not occur} \end{cases} \quad (2.2.1)$$

That is, as a function on S ,

$$\mathbf{1}_A(s) = \begin{cases} 1, & s \in A \\ 0, & s \notin A \end{cases} \quad (2.2.2)$$

If X is a random variable that takes values 0 and 1, then X is the indicator variable of the event $\{X = 1\}$.

Proof

Note that for $s \in S$, $X(s) = 1$ if $s \in \{X = 1\}$ and $X(s) = 0$ otherwise.

Recall also that the set algebra of events translates into the arithmetic algebra of indicator variables.

Suppose that A and B are events.

1. $\mathbf{1}_{A \cap B} = \mathbf{1}_A \mathbf{1}_B = \min \{\mathbf{1}_A, \mathbf{1}_B\}$
2. $\mathbf{1}_{A \cup B} = 1 - (1 - \mathbf{1}_A)(1 - \mathbf{1}_B) = \max \{\mathbf{1}_A, \mathbf{1}_B\}$
3. $\mathbf{1}_{B \setminus A} = \mathbf{1}_B (1 - \mathbf{1}_A)$
4. $\mathbf{1}_{A^c} = 1 - \mathbf{1}_A$
5. $A \subseteq B$ if and only if $\mathbf{1}_A \leq \mathbf{1}_B$

The results in part (a) extends to arbitrary intersections and the results in part (b) extends to arbitrary unions. If the event A has a complicated description, sometimes we use $\mathbf{1}(A)$ for the indicator variable rather than $\mathbf{1}_A$.

Examples and Applications

Recall that probability theory is often illustrated using simple devices from games of chance: coins, dice, cards, spinners, urns with balls, and so forth. Examples based on such devices are pedagogically valuable because of their simplicity and conceptual clarity. On the other hand, remember that probability is not only about gambling and games of chance. Rather, try to see problems involving coins, dice, etc. as metaphors for more complex and realistic problems.

Coins and Dice

The basic *coin experiment* consists of tossing a coin n times and recording the sequence of scores (X_1, X_2, \dots, X_n) (where 1 denotes heads and 0 denotes tails). This experiment is a generic example of n Bernoulli trials, named for Jacob Bernoulli.

Consider the coin experiment with $n = 4$, and Let Y denote the number of heads.

1. Give the set of outcomes S in list form.
2. Give the event $\{Y = k\}$ in list form for each $k \in \{0, 1, 2, 3, 4\}$.

Answer

To simplify the notation, we represent outcomes a *bit strings* rather than ordered sequences.

$$1. S = \{1111, 1110, 1101, 1011, 0111, 1100, 1010, 1001, 0110, 0101, 0011, 1000, 0100, 0010, 0001, 0000\}$$

$$\begin{aligned} 2. \quad \{Y = 0\} &= \{0000\} & (2.2.3) \\ \{Y = 1\} &= \{1000, 0100, 0010, 0001\} & (2.2.4) \\ \{Y = 2\} &= \{1100, 1010, 1001, 0110, 0101, 0011\} & (2.2.5) \\ \{Y = 3\} &= \{1110, 1101, 1011, 0111\} & (2.2.6) \\ \{Y = 4\} &= \{1111\} & (2.2.7) \end{aligned}$$

In the simulation of the coin experiment, set $n = 4$. Run the experiment 100 times and count the number of times that the event $\{Y = 2\}$ occurs.

Now consider the general coin experiment with the coin tossed n times, and let Y denote the number of heads.

1. Give the set of outcomes S in Cartesian product form, and give the cardinality of S .
2. Express Y as a function on S .
3. Find $\#\{Y = k\}$ (as a subset of S) for $k \in \{0, 1, \dots, n\}$

Answer

1. $S = \{0, 1\}^n$ and $\#(S) = 2^n$.
2. $Y(x_1, x_2, \dots, x_n) = x_1 + x_2 + \dots + x_n$. The set of possible values is $\{0, 1, \dots, n\}$
3. $\#\{Y = k\} = \binom{n}{k}$

The basic *dice experiment* consists of throwing n distinct k -sided dice (with faces numbered from 1 to k) and recording the sequence of scores (X_1, X_2, \dots, X_n) . This experiment is a generic example of n multinomial trials. The special case $k = 6$ corresponds to *standard dice*.

Consider the dice experiment with $n = 2$ standard dice. Let S denote the set of outcomes, A the event that the first die score is 1, and B the event that the sum of the scores is 7. Give each of the following events in the form indicated:

1. S in Cartesian product form
2. A in list form
3. B in list form
4. $A \cup B$ in list form
5. $A \cap B$ in list form
6. $A^c \cap B^c$ in predicate form

Answer

1. $S = \{1, 2, 3, 4, 5, 6\}^2$
2. $A = \{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6)\}$

3. $B = \{(1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}$
4. $A \cup B = \{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}$
5. $A \cap B = \{(1, 6)\}$
6. $A^c \cap B^c = \{(x, y) \in S : x + y \neq 7 \text{ and } x \neq 1\}$

In the simulation of the dice experiment, set $n = 2$. Run the experiment 100 times and count the number of times each event in the previous exercise occurs.

Consider the dice experiment with $n = 2$ standard dice, and let S denote the set of outcomes, Y the sum of the scores, U the minimum score, and V the maximum score.

1. Express Y as a function on S and give the set of possible values in list form.
2. Express U as a function on S and give the set of possible values in list form.
3. Express V as a function on the S and give the set of possible values in list form.
4. Give the set of possible values of (U, V) in predicate form

Answer

Note that $S = \{1, 2, 3, 4, 5, 6\}^2$. The following functions are defined on S .

1. $Y(x_1, x_2) = x_1 + x_2$. The set of values is $\{2, 3, \dots, 12\}$
2. $U(x_1, x_2) = \min\{x_1, x_2\}$. The set of values is $\{1, 2, \dots, 6\}$
3. $V(x_1, x_2) = \max\{x_1, x_2\}$. The set of values is $\{1, 2, \dots, 6\}$
4. $\{(u, v) \in \{1, 2, 3, 4, 5, 6\}^2 : u \leq v\}$

Consider again the dice experiment with $n = 2$ standard dice, and let S denote the set of outcomes, Y the sum of the scores, U the minimum score, and V the maximum score. Give each of the following as subsets of S , in list form.

1. $\{X_1 < 3, X_2 > 4\}$
2. $\{Y = 7\}$
3. $\{U = 2\}$
4. $\{V = 4\}$
5. $\{U = V\}$

Answer

1. $\{(1, 5), (2, 5), (1, 6), (2, 6)\}$
2. $\{(1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}$
3. $\{(2, 2), (2, 3), (3, 2), (2, 4), (4, 2), (2, 5), (5, 2), (2, 6), (6, 2)\}$
4. $\{(4, 1), (1, 4), (2, 4), (4, 2), (4, 3), (3, 4), (4, 4)\}$
5. $\{(1, 1), (2, 2), (3, 3), (4, 4), (5, 5), (6, 6)\}$

In the dice experiment, set $n = 2$. Run the experiment 100 times. Count the number of times each event in the previous exercise occurred.

In the general dice experiment with n distinct k -sided dice, let Y denote the sum of the scores, U the minimum score, and V the maximum score.

1. Give the set of outcomes S and find $\#(S)$.
2. Express Y as a function on S , and give the set of possible values in list form.
3. Express U as a function on S , and give the set of possible values in list form.
4. Express V as a function on S , and give the set of possible values in list form.
5. Give the set of possible values of (U, V) in predicate form.

Answer

1. $S = \{1, 2, \dots, k\}^n$ and $\#(S) = k^n$
2. $Y(x_1, x_2, \dots, x_n) = x_1 + x_2 + \dots + x_n$. The set of possible values is $\{n, n+1, \dots, nk\}$
3. $U(x_1, x_2, \dots, x_n) = \min\{x_1, x_2, \dots, x_n\}$. The set of possible values is $\{1, 2, \dots, k\}$

4. $V(x_1, x_2, \dots, x_n) = \max\{x_1, x_2, \dots, x_n\}$. The set of possible values is $\{1, 2, \dots, k\}$
5. $\{(u, v) \in \{1, 2, \dots, k\}^2 : u \leq v\}$

The set of outcomes of a random experiment depends of course on what information is recorded. The following exercise is an illustration.

An experiment consists of throwing a pair of standard dice repeatedly until the sum of the two scores is either 5 or 7. Let A denote the event that the sum is 5 rather than 7 on the final throw. Experiments of this type arise in the casino game craps.

1. Suppose that the pair of scores on each throw is recorded. Define the set of outcomes of the experiment and describe A as a subset of this set.
2. Suppose that the pair of scores on the final throw is recorded. Define the set of outcomes of the experiment and describe A as a subset of this set.

Answer

Let $D_5 = \{(1, 4), (2, 3), (3, 2), (4, 1)\}$ $D_7 = \{(1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}$ $D = D_5 \cup D_7$, and $C = D^c$

1. $S = D \cup (C \times D) \cup (C^2 \times D) \cup \dots$, $A = D_5 \cup (C \times D_5) \cup (C^2 \times D_5) \cup \dots$
2. $S = D$, $A = D_5$

Suppose that 3 standard dice are rolled and the sequence of scores (X_1, X_2, X_3) is recorded. A person pays \$1 to play. If some of the dice come up 6, then the player receives her \$1 back, plus \$1 for each 6. Otherwise she loses her \$1. Let W denote the person's net winnings. This is the game of chuck-a-luck and is treated in more detail in the chapter on Games of Chance.

1. Give the set of outcomes S in Cartesian product form.
2. Express W as a function on S and give the set of possible values in list form.

Answer

1. $S = \{1, 2, 3, 4, 5, 6\}^3$
2. $W(x_1, x_2, x_3) = \mathbf{1}(x_1 = 6) + \mathbf{1}(x_2 = 6) + \mathbf{1}(x_3 = 6) - \mathbf{1}(x_1 \neq 6, x_2 \neq 6, x_3 \neq 6)$. The set of possible values is $\{-1, 1, 2, 3\}$

Play the chuck-a-luck experiment a few times and see how you do.

In the *die-coin experiment*, a standard die is rolled and then a coin is tossed the number of times shown on the die. The sequence of coin scores \mathbf{X} is recorded (0 for tails, 1 for heads). Let N denote the die score and Y the number of heads.

1. Give the set of outcomes S in terms of Cartesian powers and find $\#(S)$.
2. Express N as a function on S and give the set of possible values in list form.
3. Express Y as a function on S and give the set of possible values in list form.
4. Give the event A that all tosses result in heads in list form.

Answer

1. $S = \bigcup_{n=1}^6 \{0, 1\}^n$, $\#(S) = 126$
2. $N(x_1, x_2, \dots, x_n) = n$ for $(x_1, x_2, \dots, x_n) \in S$. The set of values is $\{1, 2, 3, 4, 5, 6\}$
3. $Y(x_1, x_2, \dots, x_n) = \sum_{i=1}^n x_i$ for $(x_1, x_2, \dots, x_n) \in S$. The set of possible values is $\{0, 1, 2, 3, 4, 5, 6\}$
4. $A = \{1, 11, 111, 1111, 11111, 111111\}$

Run the simulation of the die-coin experiment 10 times. For each run, give the values of the random variables \mathbf{X} , N , and Y of the previous exercise. Count the number of times the event A occurs.

In the *coin-die experiment*, we have a coin and two distinct dice, say one red and one green. First the coin is tossed, and then if the result is heads the red die is thrown, while if the result is tails the green die is thrown. The coin score X and the score of the chosen die Y are recorded. Suppose now that the red die is a standard 6-sided die, and the green die a 4-sided die.

1. Give the set of outcomes S in list form.
2. Express X as a function on S .

- Express Y as a function on S .
- Give the event $\{Y \geq 3\}$ as a subset of S in list form.

Answer

- $\{(0, 1), (0, 2), (0, 3), (0, 4), (1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6)\}$
- $X(i, j) = i$ for $(i, j) \in S$
- $Y(i, j) = j$ for $(i, j) \in S$
- $\{(0, 3), (0, 4), (1, 3), (1, 4), (1, 5), (1, 6)\}$

Run the coin-die experiment 100 times, with various types of dice.

Sampling Models

Recall that many random experiments can be thought of as *sampling experiments*. For the general finite sampling model, we start with a population D with m (distinct) objects. We select a sample of n objects from the population. If the sampling is done in a random way, then we have a random experiment with the sample as the basic outcome. Thus, the set of outcomes of the experiment is literally the *set of samples*; this is the historical origin of the term *sample space*. There are four common types of sampling from a finite population, based on the criteria of *order* and *replacement*. Recall the following facts from the section on combinatorial structures:

Samples of size n chosen from a population with m elements.

- If the sampling is with replacement and with regard to order, then the set of samples is the Cartesian power D^n . The number of samples is m^n .
- If the sampling is without replacement and with regard to order, then the set of samples is the set of all *permutations* of size n from D . The number of samples is $m^{(n)} = m(m-1) \cdots [m-(n-1)]$.
- If the sampling is without replacement and without regard to order, then the set of samples is the set of all *combinations* (or *subsets*) of size n from D . The number of samples is $\binom{m}{n}$.
- If the sampling is with replacement and without regard to order, then the set of samples is the set of all *multisets* of size n from D . The number of samples is $\binom{m+n-1}{n}$.

If we sample with replacement, the sample size n can be any positive integer. If we sample without replacement, the sample size cannot exceed the population size, so we must have $n \in \{1, 2, \dots, m\}$.

The basic [coin and dice experiments](#) are examples of sampling *with* replacement. If we toss a coin n times and record the sequence of scores (where as usual, 0 denotes tails and 1 denotes heads), then we generate an ordered sample of size n with replacement from the population $\{0, 1\}$. If we throw n (distinct) standard dice and record the sequence of scores, then we generate an ordered sample of size n with replacement from the population $\{1, 2, 3, 4, 5, 6\}$.

Suppose that the sampling is without replacement (the most common case). If we record the *ordered* sample $\mathbf{X} = (X_1, X_2, \dots, X_n)$, then the *unordered* sample $\mathbf{W} = \{X_1, X_2, \dots, X_n\}$ is a random variable (that is, a function of \mathbf{X}). On the other hand, if we just record the unordered sample \mathbf{W} in the first place, then we cannot recover the ordered sample. Note also that the number of ordered samples of size n is simply $n!$ times the number of unordered samples of size n . No such simple relationship exists when the sampling is with replacement. This will turn out to be an important point when we study probability models based on random samples, in the next section.

Consider a sample of size $n = 3$ chosen without replacement from the population $\{a, b, c, d, e\}$.

- Give T , the set of unordered samples in list form.
- Give in list form the set of all ordered samples that correspond to the unordered sample $\{b, c, e\}$.
- Note that for every unordered sample, there are 6 ordered samples.
- Give the cardinality of S , the set of ordered samples.

Answer

- $T = \{\{a, b, c\}, \{a, b, d\}, \{a, b, e\}, \{a, c, d\}, \{a, c, e\}, \{a, d, e\}, \{b, c, d\}, \{b, c, e\}, \{b, d, e\}, \{c, d, e\}\}$
- $\{(b, c, e), (b, e, c), (c, b, e), (c, e, b), (e, b, c), (e, c, b)\}$
- 60

Traditionally in probability theory, an *urn* containing *balls* is often used as a metaphor for a finite population.

Suppose that an urn contains 50 (distinct) balls. A sample of 10 balls is selected from the urn. Find the number of samples in each of the following cases:

1. Ordered samples with replacement
2. Ordered samples without replacement
3. Unordered samples without replacement
4. Unordered samples with replacement

Answer

1. 97 656 250 000 000 000
2. 37 276 043 023 296 000
3. 10 272 278 170
4. 62 828 356 305

Suppose again that we have a population D with m (distinct) objects, but suppose now that each object is one of two types—either type 1 or type 0. Such populations are said to be *dichotomous*. Here are some specific examples:

- The population consists of persons, each either *male* or *female*.
- The population consists of voters, each either *democrat* or *republican*.
- The population consists of devices, each either *good* or *defective*.
- The population consists of balls, each either *red* or *green*.

Suppose that the population D has r type 1 objects and hence $m - r$ type 0 objects. Of course, we must have $r \in \{0, 1, \dots, m\}$. Now suppose that we select a sample of size n without replacement from the population. Note that this model has three parameters: the population size m , the number of type 1 objects in the population r , and the sample size n .

Let Y denote the number of type 1 objects in the sample. Then

1. $\#\{Y = k\} = \binom{n}{k} r^{(k)} (m - r)^{(n-k)}$ for each $k \in \{0, 1, \dots, n\}$, if the event is considered as a subset of S , the set of ordered samples.
2. $\#\{Y = k\} = \binom{r}{k} \binom{m-r}{n-k}$ for each $k \in \{0, 1, \dots, n\}$, if the event is considered as a subset of T , the set of unordered samples.
3. The expression in (a) is $n!$ times the expression in (b).

Proof

1. $\binom{n}{k}$ is the number of ways to pick the coordinates (in the ordered sample) where the type 1 objects will go, $r^{(k)}$ is the number of ways to select a permutation of k type 1 objects, and $(m - r)^{(n-k)}$ is the number of ways to select a permutation of $n - k$ type 0 objects. The result follows from the multiplication principle.
2. $\binom{r}{k}$ is the number of ways to select a combination of k type 1 objects and $\binom{m-r}{n-k}$ is the number of ways to select a combination of $n - k$ type 0 objects. The result again follows from the multiplication principle.
3. This result can be shown algebraically, but a combinatorial argument is better. For every combination of size n there are $n!$ permutations of those objects.

A batch of 50 components consists of 40 good components and 10 defective components. A sample of 5 components is selected, without replacement. Let Y denote the number of defectives in the sample.

1. Let S denote the set of ordered samples. Find $\#(S)$.
2. Let T denote the set of unordered samples. Find $\#(T)$.
3. As a subset of T , find $\#\{Y = k\}$ for each $k \in \{0, 1, 2, 3, 4, 5\}$.

Answer

1. 254 251 200
2. 2 118 760
3. $\#\{Y = 0\} = 658 008$ $\#\{Y = 1\} = 913 900$ $\#\{Y = 2\} = 444 600$ $\#\{Y = 3\} = 93 600$ $\#\{Y = 4\} = 8 400$,
 $\#\{Y = 5\} = 252$

Run the simulation of the ball and urn experiment 100 times for the parameter values in the last exercise: $m = 50$, $r = 10$, $n = 5$. Note the values of the random variable Y .

Cards

Recall that a standard *card deck* can be modeled by the Cartesian product set

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (2.2.8)$$

where the first coordinate encodes the *denomination* or *kind* (ace, 2–10, jack, queen, king) and where the second coordinate encodes the *suit* (clubs, diamonds, hearts, spades). Sometimes we represent a card as a *string* rather than an ordered pair (for example $q\heartsuit$ rather than (q, \heartsuit) for the queen of hearts).

Most card games involve **sampling** without replacement from the deck D , which plays the role of the population. Thus, the basic *card experiment* consists of dealing n cards from a standard deck without replacement; in this special context, the sample of cards is often referred to as a *hand*. Just as in the general sampling model, if we record the *ordered* hand $\mathbf{X} = (X_1, X_2, \dots, X_n)$, then the *unordered* hand $\mathbf{W} = \{X_1, X_2, \dots, X_n\}$ is a random variable (that is, a function of \mathbf{X}). On the other hand, if we just record the unordered hand \mathbf{W} in the first place, then we cannot recover the ordered hand. Finally, recall that $n = 5$ is the *poker experiment* and $n = 13$ is the *bridge experiment*. The game of poker is treated in more detail in the chapter on Games of Chance.

Suppose that a single card is dealt from a standard deck. Let Q denote the event that the card is a queen and H the event that the card is a heart. Give each of the following events in list form:

1. Q
2. H
3. $Q \cup H$
4. $Q \cap H$
5. $Q \setminus H$

Answer

1. $Q = \{q\clubsuit, q\diamondsuit, q\heartsuit, q\spadesuit\}$
2. $H = \{1\heartsuit, 2\heartsuit, \dots, 10\heartsuit, j\heartsuit, q\heartsuit, k\heartsuit\}$
3. $Q \cup H = \{1\heartsuit, 2\heartsuit, \dots, 10\heartsuit, j\heartsuit, q\heartsuit, k\heartsuit, q\clubsuit, q\diamondsuit, q\spadesuit\}$
4. $Q \cap H = \{q\heartsuit\}$
5. $Q \setminus H = \{q\clubsuit, q\diamondsuit, q\spadesuit\}$

In the card experiment, set $n = 1$. Run the experiment 100 times and count the number of times each event in the previous exercise occurs.

Suppose that two cards are dealt from a standard deck and the sequence of cards recorded. Let S denote the set of outcomes, and let Q_i denote the event that the i th card is a queen and H_i the event that the i th card is a heart for $i \in \{1, 2\}$. Find the number of outcomes in each of the following events:

1. S
2. H_1
3. H_2
4. $H_1 \cap H_2$
5. $Q_1 \cap H_1$
6. $Q_1 \cap H_2$
7. $H_1 \cup H_2$

Answer

1. 2652
2. 663
3. 663
4. 156
5. 51

6. 51
7. 1170

Consider the general card experiment in which n cards are dealt from a standard deck, and the ordered hand \mathbf{X} is recorded.

1. Give cardinality of S , the set of values of the ordered hand \mathbf{X} .
2. Give the cardinality of T , the set of values of the unordered hand \mathbf{W} .
3. How many ordered hands correspond to a given unordered hand?
4. Explicitly compute the numbers in (a) and (b) when $n = 5$ (poker).
5. Explicitly compute the numbers in (a) and (b) when $n = 13$ (bridge).

Answer

3. $\#(S) = 52^{(n)}$
4. $\#(T) = \binom{52}{n}$
5. $n!$
6. 311 875 2002 598 960
7. 3 954 242 643 911 239 680 000 501 359 600

Consider the bridge experiment of dealing 13 cards from a deck and recording the unordered hand. In the most common *point counting system*, an ace is worth 4 points, a king 3 points, a queen 2 points, and a jack 1 point. The other cards are worth 0 points. Let S denote the set of outcomes of the experiment and V the point value of the hand.

1. Find the set of possible values of V .
2. Find the cardinality of the event $\{V = 0\}$ as a subset of S .

Answer

1. $\{0, 1, \dots, 37\}$
2. $\#\{V = 0\} = 2\,310\,789\,600$

In the card experiment, set $n = 13$ and run the experiment 100 times. For each run, compute the value of each of the random variable V in the previous exercise.

Consider the poker experiment of dealing 5 cards from a deck. Find the cardinality of each of the events below, as a subset of the set of unordered hands.

1. A : the event that the hand is a *full house* (3 cards of one kind and 2 of another kind).
2. B : the event that the hand has *4 of a kind* (4 cards of one kind and 1 of another kind).
3. C : the event that all cards in the hand are in the same suit (the hand is a *flush* or a *straight flush*).

Answer

1. $\#(A) = 3744$
2. $\#(B) = 624$
3. $\#(C) = 5148$

Run the poker experiment 1000 times. Note the number of times that the events A , B , and C in the previous exercise occurred.

Consider the bridge experiment of dealing 13 cards from a standard deck. Let S denote the set of unordered hands, Y the number of hearts in the hand, and Z the number of queens in the hand.

1. Find the cardinality of the event $\{Y = y\}$ as a subset of S for each $y \in \{0, 1, \dots, 13\}$.
2. Find the cardinality of the event $\{Z = z\}$ as a subset of S for each $z \in \{0, 1, 2, 3, 4\}$.

Answer

1. $\#(Y = y) = \binom{13}{y} \binom{39}{13-y}$ for $y \in \{0, 1, \dots, 13\}$
2. $\#(Z = z) = \binom{4}{z} \binom{48}{4-z}$ for $z \in \{0, 1, 2, 3, 4\}$

Geometric Models

In the experiments that we have considered so far, the sample spaces have all been discrete (so that the set of outcomes is finite or countably infinite). In this subsection, we consider Euclidean sample spaces where the set of outcomes S is *continuous* in a sense that we will make clear later. The experiments we consider are sometimes referred to as *geometric models* because they involve selecting a point *at random* from a Euclidean set.

We first consider *Buffon's coin experiment*, which consists of tossing a coin with radius $r \leq \frac{1}{2}$ randomly on a floor covered with square tiles of side length 1. The coordinates (X, Y) of the center of the coin are recorded relative to axes through the center of the square in which the coin lands. Buffon's experiments are studied in more detail in the chapter on Geometric Models and are named for Comte de Buffon




Figure 2.2.3: Buffon's coin experiment

In Buffon's coin experiment, let S denote the set of outcomes, A the event that the coin does not touch the sides of the square, and let Z denote the distance from the center of the coin to the center of the square.

1. Describe S as a Cartesian product.
2. Describe A as a subset of S .
3. Describe A^c as a subset of S .
4. Express Z as a function on S .
5. Express the event $\{X < Y\}$ as a subset of S .
6. Express the event $\{Z \leq \frac{1}{2}\}$ as a subset of S .

Answer

1. $S = [-\frac{1}{2}, \frac{1}{2}]^2$
2. $A = [r - \frac{1}{2}, \frac{1}{2} - r]^2$
3. $A^c = \{(x, y) \in S : x < r - \frac{1}{2} \text{ or } x > \frac{1}{2} - r \text{ or } y < r - \frac{1}{2} \text{ or } y > \frac{1}{2} - r\}$
4. $Z(x, y) = \sqrt{x^2 + y^2}$ for $(x, y) \in S$
5. $\{X < Y\} = \{(x, y) \in S : x < y\}$
6. $\{Z < \frac{1}{2}\} = \{(x, y) \in S : x^2 + y^2 < \frac{1}{4}\}$

Run Buffon's coin experiment 100 times with $r = 0.2$. For each run, note whether event A occurs and compute the value of random variable Z .

A point (X, Y) is chosen at random in the circular region of radius 1 in \mathbb{R}^2 centered at the origin. Let S denote the set of outcomes. Let A denote the event that the point is in the inscribed square region centered at the origin, with sides parallel to the coordinate axes. Let B denote the event that the point is in the inscribed square with vertices $(\pm 1, 0), (0, \pm 1)$.

1. Describe S mathematically and sketch the set.
2. Describe A mathematically and sketch the set.
3. Describe B mathematically and sketch the set.
4. Sketch $A \cup B$
5. Sketch $A \cap B$
6. Sketch $A \cap B^c$

Answer

1. $S = \{(x, y) : x^2 + y^2 \leq 1\}$

2. $A = \left\{ (x, y) : -\frac{1}{\sqrt{2}} \leq x \leq \frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}} \leq y \leq \frac{1}{\sqrt{2}} \right\}$
3. $B = \{(x, y) \in S : -1 \leq |x + y| \leq 1, -1 \leq |y - x| \leq 1\}$

Reliability

In the simple model of *structural reliability*, a system is composed of n components, each of which is either *working* or *failed*. The state of component i is an indicator random variable X_i , where 1 means working and 0 means failure. Thus, $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a vector of indicator random variables that specifies the states of all of the components, and therefore the set of outcomes of the experiment is $S = \{0, 1\}^n$. The system as a whole is also either working or failed, depending only on the states of the components and how the components are connected together. Thus, the state of the system is also an indicator random variable and is a function of \mathbf{X} . The state of the system (working or failed) as a function of the states of the components is the *structure function*.

A *series system* is working if and only if each component is working. The state of the system is

$$U = X_1 X_2 \cdots X_n = \min \{X_1, X_2, \dots, X_n\} \quad (2.2.9)$$

A *parallel system* is working if and only if at least one component is working. The state of the system is

$$V = 1 - (1 - X_1)(1 - X_2) \cdots (1 - X_n) = \max \{X_1, X_2, \dots, X_n\} \quad (2.2.10)$$

More generally, a k out of n system is working if and only if at least k of the n components are working. Note that a parallel system is a 1 out of n system and a series system is an n out of n system. A k out of $2k$ system is a *majority rules system*.

The state of the k out of n system is $U_{n,k} = \mathbf{1}(\sum_{i=1}^n X_i \geq k)$. The structure function can also be expressed as a polynomial in the variables.

Explicitly give the state of the k out of 3 system, as a polynomial function of the component states (X_1, X_2, X_3) , for each $k \in \{1, 2, 3\}$.

Answer

1. $U_{3,1} = X_1 + X_2 + X_3 - X_1 X_2 - X_1 X_3 - X_2 X_3 + X_1 X_2 X_3$
2. $U_{3,2} = X_1 X_2 + X_1 X_3 + X_2 X_3 - 2 X_1 X_2 X_3$
3. $U_{3,3} = X_1 X_2 X_3$

In some cases, the system can be represented as a *graph* or *network*. The *edges* represent the components and the *vertices* the connections between the components. The system functions if and only if there is a working path between two designated vertices, which we will denote by a and b .

Find the state of the *Wheatstone bridge network* shown below, as a function of the component states. The network is named for Charles Wheatstone.

Answer




Figure 2.2.4: The Wheatstone bridge network

Not every function $u : \{0, 1\}^n \rightarrow \{0, 1\}$ makes sense as a structure function. Explain why the following properties might be desirable:

1. $u(0, 0, \dots, 0) = 0$ and $u(1, 1, \dots, 1) = 1$
2. u is an increasing function, where $\{0, 1\}$ is given the ordinary order and $\{0, 1\}^n$ the corresponding product order.
3. For each $i \in \{1, 2, \dots, n\}$, there exist \mathbf{x} and \mathbf{y} in $\{0, 1\}^n$ all of whose coordinates agree, except $x_i = 0$ and $y_i = 1$, and $u(\mathbf{x}) = 0$ while $u(\mathbf{y}) = 1$.

Answer

1. This means that if all components have failed then the system has failed, and if all components are working then the system is working.
2. This means that if a particular component is changed from failed to working, then the system may also go from failed to working, but not from working to failed. That is, the system can only improve.
3. This means that every component is relevant to the system, that is, there exists a configuration in which changing component i from failed to working changes the system from failed to working.

The model just discussed is a *static model*. We can extend it to a *dynamic model* by assuming that component i is initially working, but has a random *time to failure* T_i , taking values in $[0, \infty)$, for each $i \in \{1, 2, \dots, n\}$. Thus, the basic outcome of the experiment is the random vector of failure times (T_1, T_2, \dots, T_n) , and so the set of outcomes is $[0, \infty)^n$.

Consider the dynamic reliability model for a system with structure function u (valid in the sense of the previous exercise).

1. The state of component i at time $t \geq 0$ is $X_i(t) = \mathbf{1}(T_i > t)$.
2. The state of the system at time t is $X(t) = s[X_1(t), X_2(t), \dots, X_n(t)]$.
3. The time to failure of the system is $T = \min\{t \geq 0 : X(t) = 0\}$.

Suppose that we have two devices and that we record (X, Y) , where X is the failure time of device 1 and Y is the failure time of device 2. Both variables take values in the interval $[0, \infty)$, where the units are in hundreds of hours. Sketch each of the following events:

1. The set of outcomes S
2. $\{X < Y\}$
3. $\{X + Y > 2\}$

Answer

1. $S = [0, \infty)^2$, the first quadrant of the coordinate plane.
2. $\{X < Y\} = \{(x, y) \in S : x < y\}$. This is the region below the diagonal line $x = y$.
3. $\{X + Y > 2\} = \{(x, y) \in S : x + y > 2\}$. This is the region above (or to the right) of the line $x + y = 2$.

Genetics

Please refer to the discussion of genetics in the section on random experiments if you need to review some of the definitions in this section.

Recall first that the *ABO blood type* in humans is determined by three alleles: a , b , and o . Furthermore, o is recessive and a and b are co-dominant.

Suppose that a person is chosen at random and his genotype is recorded. Give each of the following in list form.

1. The set of outcomes S
2. The event that the person is type A
3. The event that the person is type B
4. The event that the person is type AB
5. The event that the person is type O

Answer

1. $S = \{aa, ab, ao, bb, bo, oo\}$
2. $A = \{aa, ao\}$
3. $B = \{bb, bo\}$
4. $AB = \{ab\}$
5. $O = \{oo\}$

Suppose next that pod color in certain type of pea plant is determined by a gene with two alleles: g for green and y for yellow, and that g is dominant.

Suppose that n (distinct) pea plants are collected and the sequence of pod color genotypes is recorded.

1. Give the set of outcomes S in Cartesian product form and find $\#(S)$.
2. Let N denote the number of plants with green pods. Find $\#(N = k)$ (as a subset of S) for each $k \in \{0, 1, \dots, n\}$.

Answer

1. $S = \{gg, gy, yy\}^n$, $\#(S) = 3^n$
2. $\binom{n}{k} 2^k$

Next consider a sex-linked hereditary disorder in humans (such as colorblindness or hemophilia). Let h denote the healthy allele and d the defective allele for the gene linked to the disorder. Recall that d is recessive for women.

Suppose that n women are sampled and the sequence of genotypes is recorded.

1. Give the set of outcomes S in Cartesian product form and find $\#(S)$.
2. Let N denote the number of women who are completely healthy (genotype hh). Find $\#(N = k)$ (as a subset of S) for each $k \in \{0, 1, \dots, n\}$.

Answer

1. $S = \{hh, hd, dd\}^n$, $\#(S) = 3^n$
2. $\binom{n}{k} 2^{n-k}$

Radioactive Emissions

The emission of elementary particles from a sample of radioactive material occurs in a random way. Suppose that the time of emission of the i th particle is a random variable T_i taking values in $(0, \infty)$. If we measure these *arrival times*, then basic outcome vector is (T_1, T_2, \dots) and so the set of outcomes is $S = \{(t_1, t_2, \dots) : 0 < t_1 < t_2 < \dots\}$.

Run the simulation of the gamma experiment in single-step mode for different values of the parameters. Observe the arrival times.

Now let N_t denote the number of emissions in the interval $(0, t]$. Then

1. $N_t = \max \{n \in \mathbb{N}_+ : T_n \leq t\}$.
2. $N_t \geq n$ if and only if $T_n \leq t$.

Run the simulation of the Poisson experiment in single-step mode for different parameter values. Observe the arrivals in the specified time interval.

Statistical Experiments

In the basic cicada experiment, a cicada in the Middle Tennessee area is captured and the following measurements recorded: body weight (in grams), wing length, wing width, and body length (in millimeters), species type, and gender. The cicada data set gives the results of 104 repetitions of this experiment.

1. Define the set of outcomes S for the basic experiment.
2. Let F be the event that a cicada is female. Describe F as a subset of S . Determine whether F occurs for each cicada in the data set.
3. Let V denote the ratio of wing length to wing width. Compute V for each cicada.
4. Give the set of outcomes for the compound experiment that consists of 104 repetitions of the basic experiment.

Answer

For gender, let 0 denote female and 1 male, for species, let 1 denote tredecula, 2 tredecim, and 3 tredecassini.

1. $S = (0, \infty)^4 \times \{0, 1\} \times \{1, 2, 3\}$
2. $F = \{(x_1, x_2, x_3, x_4, y, z) \in S : y = 0\}$
5. S^{104}

In the basic M&M experiment, a bag of M&Ms (of a specified size) is purchased and the following measurements recorded: the number of red, green, blue, yellow, orange, and brown candies, and the net weight (in grams). The M&M data set gives the

results of 30 repetitions of this experiment.

1. Define the set of outcomes S for the basic experiment.
2. Let A be the event that a bag contains at least 57 candies. Describe A as a subset of S .
3. Determine whether A occurs for each bag in the data set.
4. Let N denote the total number of candies. Compute N for each bag in the data set.
5. Give the set of outcomes for the compound experiment that consists of 30 repetitions of the basic experiment.

Answer

1. $S = \mathbb{N}^6 \times (0, \infty)$
2. $A = \{(n_1, n_2, n_3, n_4, n_5, n_6, w) \in S : n_1 + n_2 + \cdots + n_6 > 57\}$
5. S^{30}

This page titled [2.2: Events and Random Variables](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.3: Probability Measures

This section contains the final and most important ingredient in the basic model of a random experiment. If you are a new student of probability, skip the technical details.

Definitions and Interpretations

Suppose that we have a random experiment with sample space (S, \mathcal{S}) , so that S is the set of outcomes of the experiment and \mathcal{S} is the collection of events. When we run the experiment, a given event A either *occurs* or *does not occur*, depending on whether the outcome of the experiment is in A or not. Intuitively, the *probability* of an event is a measure of how likely the event is to occur when we run the experiment. Mathematically, probability is a function on the collection of events that satisfies certain axioms.

Definition

A *probability measure* (or *probability distribution*) \mathbb{P} on the sample space (S, \mathcal{S}) is a real-valued function defined on the collection of events \mathcal{S} that satisfies the following axioms:

1. $\mathbb{P}(A) \geq 0$ for every event A .
2. $\mathbb{P}(S) = 1$.
3. If $\{A_i : i \in I\}$ is a countable, pairwise disjoint collection of events then

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mathbb{P}(A_i) \quad (2.3.1)$$

Details

Recall that the collection of events \mathcal{S} is required to be a σ -algebra, which guarantees that the union of the events in (c) is itself an event. A probability measure is a special case of a positive measure.

Axiom (c) is known as *countable additivity*, and states that the probability of a union of a finite or countably infinite collection of disjoint events is the sum of the corresponding probabilities. The axioms are known as the *Kolmogorov axioms*, in honor of Andrei Kolmogorov who was the first to formalize probability theory in an axiomatic way. More informally, we say that \mathbb{P} is a probability measure (or distribution) on S , the collection of events \mathcal{S} usually being understood.

Axioms (a) and (b) are really just a matter of convention; we choose to measure the probability of an event with a number between 0 and 1 (as opposed, say, to a number between -5 and 7). Axiom (c) however, is fundamental and inescapable. It is required for probability for precisely the same reason that it is required for other measures of the “size” of a set, such as cardinality for finite sets, length for subsets of \mathbb{R} , area for subsets of \mathbb{R}^2 , and volume for subsets of \mathbb{R}^3 . In all these cases, the size of a set that is composed of countably many disjoint pieces is the sum of the sizes of the pieces.




Figure 2.3.1: The union of 4 disjoint events

On the other hand, *uncountable additivity* (the extension of axiom (c) to an uncountable index set I) is unreasonable for probability, just as it is for other measures. For example, an interval of positive length in \mathbb{R} is a union of uncountably many points, each of which has length 0.

We now have defined the three essential ingredients for the model a random experiment:

A *probability space* $(S, \mathcal{S}, \mathbb{P})$ consists of

1. A set of outcomes S
2. A collection of events \mathcal{S}
3. A probability measure \mathbb{P} on the sample space (S, \mathcal{S})

Details

Again, the collection of events \mathcal{S} is a σ -algebra, so that the sample space (S, \mathcal{S}) is a measurable space. The probability space $(S, \mathcal{S}, \mathbb{P})$ is a special case of a positive measure space.

The Law of Large Numbers

Intuitively, the probability of an event is supposed to measure the long-term relative frequency of the event—in fact, this concept was taken as the definition of probability by Richard Von Mises. Here are the relevant definitions:

Suppose that the experiment is repeated indefinitely, and that A is an event. For $n \in \mathbb{N}_+$,

1. Let $N_n(A)$ denote the number of times that A occurred. This is the *frequency* of A in the first n runs.
2. Let $P_n(A) = N_n(A)/n$. This is the *relative frequency* or *empirical probability* of A in the first n runs.

Note that repeating the original experiment indefinitely creates a new, compound experiment, and that $N_n(A)$ and $P_n(A)$ are random variables for the new experiment. In particular, the values of these variables are uncertain until the experiment is run n times. The basic idea is that if we have chosen the correct probability measure for the experiment, then in some sense we expect that the relative frequency of an event should converge to the probability of the event. That is,

$$P_n(A) \rightarrow \mathbb{P}(A) \text{ as } n \rightarrow \infty, \quad A \in \mathcal{S} \quad (2.3.2)$$

regardless of the uncertainty of the relative frequencies on the left. The precise statement of this is the law of large numbers or *law of averages*, one of the fundamental theorems in probability. To emphasize the point, note that in general there will be lots of *possible* probability measures for an experiment, in the sense of the axioms. However, only the probability measure that models the experiment correctly will satisfy the law of large numbers.

Given the data from n runs of the experiment, the empirical probability function P_n is a probability measure on S .

Proof

If we run the experiment n times, we generate n points in S (although of course, some of these points may be the same). The function $A \mapsto N_n(A)$ for $A \subseteq S$ is simply counting measure corresponding to the n points. Clearly $P_n(A) \geq 0$ for an event A and $P_n(S) = n/n = 1$. Countable additivity holds by the addition rule for counting measure.

The Distribution of a Random Variable

Suppose now that X is a random variable for the experiment, taking values in a set T . Recall that mathematically, X is a function from S into T , and $\{X \in B\}$ denotes the event $\{s \in S : X(s) \in B\}$ for $B \subseteq T$. Intuitively, X is a variable of interest for the experiment, and every meaningful statement about X defines an event.

The function $B \mapsto \mathbb{P}(X \in B)$ for $B \subseteq T$ defines a probability measure on T .

Proof




Figure 2.3.2: A set $B \in \mathcal{T}$ corresponds to the event $\{X \in B\} \in \mathcal{S}$

The probability measure in (5) is called the *probability distribution* of X , so we have all of the ingredients for a new probability space.

A random variable X with values in T defines a new probability space:

1. T is the set of outcomes.
2. Subsets of T are the events.
3. The probability distribution of X is the probability measure on T .

This probability space corresponds to the new random experiment in which the outcome is X . Moreover, recall that the outcome of the experiment itself can be thought of as a random variable. Specifically, if we let $T = S$ we let X be the identity function on S , so that $X(s) = s$ for $s \in S$. Then X is a random variable with values in S and $\mathbb{P}(X \in A) = \mathbb{P}(A)$ for each event A . Thus, every probability measure can be thought of as the distribution of a random variable.

Constructions

Measures

How can we construct probability measures? As noted briefly above, there are other measures of the “size” of sets; in many cases, these can be converted into probability measures. First, a positive measure μ on the sample space (S, \mathcal{S}) is a real-valued function defined on \mathcal{S} that satisfies axioms (a) and (c) in (1), and then (S, \mathcal{S}, μ) is a *measure space*. In general, $\mu(A)$ is allowed to be infinite. However, if $\mu(S)$ is positive and finite (so that μ is a *finite positive measure*), then μ can easily be re-scaled into a probability measure.

If μ is a positive measure on S with $0 < \mu(S) < \infty$ then \mathbb{P} defined below is a probability measure.

$$\mathbb{P}(A) = \frac{\mu(A)}{\mu(S)}, \quad A \in \mathcal{S} \quad (2.3.3)$$

Proof

1. $\mathbb{P}(A) \geq 0$ since $\mu(A) \geq 0$ and $0 < \mu(S) < \infty$.
2. $\mathbb{P}(S) = \mu(S)/\mu(S) = 1$
3. If $\{A_i : i \in I\}$ is a countable collection of disjoint events then

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) = \frac{1}{\mu(S)} \mu\left(\bigcup_{i \in I} A_i\right) = \frac{1}{\mu(S)} \sum_{i \in I} \mu(A_i) = \sum_{i \in I} \frac{\mu(A_i)}{\mu(S)} = \sum_{i \in I} \mathbb{P}(A_i) \quad (2.3.4)$$

In this context, $\mu(S)$ is called the *normalizing constant*. In the next two subsections, we consider some very important special cases.

Discrete Distributions

In this discussion, we assume that the sample space (S, \mathcal{S}) is *discrete*. Recall that this means that the set of outcomes S is countable and that $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S , so that every subset is an event. The standard measure on a discrete space is counting measure $\#$, so that $\#(A)$ is the number of elements in A for $A \subseteq S$. When S is finite, the probability measure corresponding to counting measure as constructed in above is particularly important in combinatorial and sampling experiments.

Suppose that S is a finite, nonempty set. The *discrete uniform distribution* on S is given by

$$\mathbb{P}(A) = \frac{\#(A)}{\#(S)}, \quad A \subseteq S \quad (2.3.5)$$

The underlying model is referred to as the *classical probability model*, because historically the very first problems in probability (involving coins and dice) fit this model.

In the general discrete case, if \mathbb{P} is a probability measure on S , then since S is countable, it follows from countable additivity that \mathbb{P} is completely determined by its values on the singleton events. Specifically, if we define $f(x) = \mathbb{P}(\{x\})$ for $x \in S$, then $\mathbb{P}(A) = \sum_{x \in A} f(x)$ for every $A \subseteq S$. By axiom (a), $f(x) \geq 0$ for $x \in S$ and by axiom (b), $\sum_{x \in S} f(x) = 1$. Conversely, we can give a general construction for defining a probability measure on a discrete space.

Suppose that $g: S \rightarrow [0, \infty)$. Then μ defined by $\mu(A) = \sum_{x \in A} g(x)$ for $A \subseteq S$ is a positive measure on S . If $0 < \mu(S) < \infty$ then \mathbb{P} defined as follows is a probability measure on S .

$$\mathbb{P}(A) = \frac{\mu(A)}{\mu(S)} = \frac{\sum_{x \in A} g(x)}{\sum_{x \in S} g(x)}, \quad A \subseteq S \quad (2.3.6)$$

Proof

Trivially $\mu(A) \geq 0$ for $A \subseteq S$ since g is nonnegative. The countable additivity property holds since the terms in a sum of nonnegative numbers can be rearranged in any way without altering the sum. Thus let $\{A_i : i \in I\}$ be a countable collection of disjoint subsets of S , and let $A = \bigcup_{i \in I} A_i$ then

$$\mu(A) = \sum_{x \in A} g(x) = \sum_{i \in I} \sum_{x \in A_i} g(x) = \sum_{i \in I} \mu(A_i) \quad (2.3.7)$$

If $0 < \mu(S) < \infty$ then \mathbb{P} is a probability measure by [scaling result](#) above.

In the context of our previous remarks, $f(x) = g(x)/\mu(S) = g(x)/\sum_{y \in S} g(y)$ for $x \in S$. Distributions of this type are said to be *discrete*. Discrete distributions are studied in detail in the chapter on Distributions.

If S is finite and g is a constant function, then the probability measure \mathbb{P} associated with g is the discrete uniform distribution on S .

Proof

Suppose that $g(x) = c$ for $x \in S$ where $c > 0$. Then $\mu(A) = c\#(A)$ and hence $\mathbb{P}(A) = \mu(A)/\mu(S) = \#(A)/\#(S)$ for $A \subseteq S$.

Continuous Distributions

The probability distributions that we will construct next are *continuous distributions* on \mathbb{R}^n for $n \in \mathbb{N}_+$ and require some calculus.

For $n \in \mathbb{N}_+$, the standard measure λ_n on \mathbb{R}^n is given by

$$\lambda_n(A) = \int_A 1 \, dx, \quad A \subseteq \mathbb{R}^n \quad (2.3.8)$$

In particular, $\lambda_1(A)$ is the length of $A \subseteq \mathbb{R}$, $\lambda_2(A)$ is the area of $A \subseteq \mathbb{R}^2$, and $\lambda_3(A)$ is the volume of $A \subseteq \mathbb{R}^3$.

Details

Technically, λ_n is Lebesgue measure on the measurable subsets of \mathbb{R}^n , named for Henri Lebesgue. The representation above in terms of the ordinary Riemann integral of calculus is valid for the subsets that typically occur in applications. As usual, all subsets of \mathbb{R}^n in the discussion below are assumed to be measurable.

When $n > 3$, $\lambda_n(A)$ is sometimes called the *n-dimensional volume* of $A \subseteq \mathbb{R}^n$. The probability measure associated with λ_n on a set with positive, finite *n-dimensional volume* is particularly important.

Suppose that $S \subseteq \mathbb{R}^n$ with $0 < \lambda_n(S) < \infty$. The *continuous uniform distribution* on S is defined by

$$\mathbb{P}(A) = \frac{\lambda_n(A)}{\lambda_n(S)}, \quad A \subseteq S \quad (2.3.9)$$

Note that the continuous uniform distribution is analogous to the discrete uniform distribution defined in (8), but with Lebesgue measure λ_n replacing counting measure $\#$. We can generalize this construction to produce many other distributions.

Suppose again that $S \subseteq \mathbb{R}^n$ and that $g : S \rightarrow [0, \infty)$. Then μ defined by $\mu(A) = \int_A g(x) \, dx$ for $A \subseteq S$ is a positive measure on S . If $0 < \mu(S) < \infty$, then \mathbb{P} defined as follows is a probability measure on S .

$$\mathbb{P}(A) = \frac{\mu(A)}{\mu(S)} = \frac{\int_A g(x) \, dx}{\int_S g(x) \, dx}, \quad A \in \mathcal{S} \quad (2.3.10)$$

Proof

Technically, the integral in the definition of $\mu(A)$ is the Lebesgue integral, but this integral agrees with the ordinary Riemann integral of calculus when g and A are sufficiently nice. The function g is assumed to be measurable and is the density function of μ with respect to λ_n . Technicalities aside, the proof is straightforward:

1. $\mu(A) \geq 0$ for $A \subseteq S$ since g is nonnegative.
2. If $\{A_i : i \in I\}$ is a countable disjoint collection of subsets of S and $A = \bigcup_{i \in I} A_i$, then by a basic property of the integral,

$$\mu(A) = \int_A g(x) dx = \sum_{i \in I} \int_{A_i} g(x) dx = \sum_{i \in I} \mu(A_i) \quad (2.3.11)$$

If $0 < \mu(S) < \infty$ then \mathbb{P} is a probability measure on S by the [scaling result](#) above.

Distributions of this type are said to be *continuous*. Continuous distributions are studied in detail in the chapter on Distributions. Note that the continuous distribution above is analogous to the discrete distribution in (9), but with integrals replacing sums. The general theory of integration allows us to unify these two special cases, and many others besides.

Rules of Probability

Basic Rules

Suppose again that we have a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$, so that S is the set of outcomes, \mathcal{S} the collection of events, and \mathbb{P} the probability measure. In the following theorems, A and B are events. The results follow easily from the axioms of probability in (1), so be sure to try the proofs yourself before reading the ones in the text.

$\mathbb{P}(A^c) = 1 - \mathbb{P}(A)$. This is known as the *complement rule*.

Proof




Figure 2.3.3: The complement rule

$\mathbb{P}(\emptyset) = 0$.

Proof

This follows from the the complement rule applied to $A = S$.

$\mathbb{P}(B \setminus A) = \mathbb{P}(B) - \mathbb{P}(A \cap B)$. This is known as the *difference rule*.

Proof



Figure 2.3.4: The difference rule

If $A \subseteq B$ then $\mathbb{P}(B \setminus A) = \mathbb{P}(B) - \mathbb{P}(A)$.

Proof

This result is a corollary of the difference rule. Note that $A \cap B = A$.

Recall that if $A \subseteq B$ we sometimes write $B - A$ for the set difference, rather than $B \setminus A$. With this notation, the difference rule has the nice form $\mathbb{P}(B - A) = \mathbb{P}(B) - \mathbb{P}(A)$.

If $A \subseteq B$ then $\mathbb{P}(A) \leq \mathbb{P}(B)$.

Proof

This result is a corollary of the previous result. Note that $\mathbb{P}(B \setminus A) \geq 0$ and hence $\mathbb{P}(B) - \mathbb{P}(A) \geq 0$.

Thus, \mathbb{P} is an increasing function, relative to the subset partial order on the collection of events \mathcal{S} , and the ordinary order on \mathbb{R} . In particular, it follows that $\mathbb{P}(A) \leq 1$ for any event A .




Figure 2.3.5: The increasing property

Suppose that $A \subseteq B$.

1. If $\mathbb{P}(B) = 0$ then $\mathbb{P}(A) = 0$.
2. If $\mathbb{P}(A) = 1$ then $\mathbb{P}(B) = 1$.

Proof

This follows immediately from the increasing property in the last theorem.

The Boole and Bonferroni Inequalities

The next result is known as *Boole's inequality*, named after George Boole. The inequality gives a simple upper bound on the probability of a union.

If $\{A_i : i \in I\}$ is a countable collection of events then

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) \leq \sum_{i \in I} \mathbb{P}(A_i) \quad (2.3.12)$$

Proof




Figure 2.3.6: Boole's inequality

Intuitively, Boole's inequality holds because parts of the union have been measured more than once in the sum of the probabilities on the right. Of course, the sum of the probabilities may be greater than 1, in which case Boole's inequality is not helpful. The following result is a simple consequence of *Boole's inequality*.

If $\{A_i : i \in I\}$ is a countable collection of events with $\mathbb{P}(A_i) = 0$ for each $i \in I$, then

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) = 0 \quad (2.3.13)$$

An event A with $\mathbb{P}(A) = 0$ is said to be *null*. Thus, a countable union of null events is still a null event.

The next result is known as *Bonferroni's inequality*, named after Carlo Bonferroni. The inequality gives a simple lower bound for the probability of an intersection.

If $\{A_i : i \in I\}$ is a countable collection of events then

$$\mathbb{P}\left(\bigcap_{i \in I} A_i\right) \geq 1 - \sum_{i \in I} [1 - \mathbb{P}(A_i)] \quad (2.3.14)$$

Proof

By De Morgan's law, $(\bigcap_{i \in I} A_i)^c = \bigcup_{i \in I} A_i^c$. Hence by *Boole's inequality*,

$$\mathbb{P} \left[\left(\bigcap_{i \in I} A_i \right)^c \right] \leq \sum_{i \in I} \mathbb{P}(A_i^c) = \sum_{i \in I} [1 - \mathbb{P}(A_i)] \quad (2.3.15)$$

Using the complement rule again gives Bonferroni's inequality.

Of course, the lower bound in Bonferroni's inequality may be less than or equal to 0, in which case it's not helpful. The following result is a simple consequence of Bonferroni's inequality.

If $\{A_i : i \in I\}$ is a countable collection of events with $\mathbb{P}(A_i) = 1$ for each $i \in I$, then

$$\mathbb{P} \left(\bigcap_{i \in I} A_i \right) = 1 \quad (2.3.16)$$

An event A with $\mathbb{P}(A) = 1$ is sometimes called *almost sure* or *almost certain*. Thus, a countable intersection of almost sure events is still almost sure.

Suppose that A and B are events in an experiment.

1. If $\mathbb{P}(A) = 0$, then $\mathbb{P}(A \cup B) = \mathbb{P}(B)$.
2. If $\mathbb{P}(A) = 1$, then $\mathbb{P}(A \cap B) = \mathbb{P}(B)$.

Proof

1. Using the [increasing property](#) and [Boole's inequality](#) we have $\mathbb{P}(B) \leq \mathbb{P}(A \cup B) \leq \mathbb{P}(A) + \mathbb{P}(B) = \mathbb{P}(B)$
2. Using the [increasing property](#) and [Bonferroni's inequality](#) we have $\mathbb{P}(B) = \mathbb{P}(A) + \mathbb{P}(B) - 1 \leq \mathbb{P}(A \cap B) \leq \mathbb{P}(B)$

The Partition Rule

Suppose that $\{A_i : i \in I\}$ is a countable collection of events that partition S . Recall that this means that the events are disjoint and their union is S . For any event B ,

$$\mathbb{P}(B) = \sum_{i \in I} \mathbb{P}(A_i \cap B) \quad (2.3.17)$$

Proof




Figure 2.3.7: The partition rule

Naturally, this result is useful when the probabilities of the intersections are known. Partitions usually arise in connection with a random variable. Suppose that X is a random variable taking values in a countable set T , and that B is an event. Then

$$\mathbb{P}(B) = \sum_{x \in T} \mathbb{P}(X = x, B) \quad (2.3.18)$$

In this formula, note that the comma acts like the intersection symbol in the previous formula.

The Inclusion-Exclusion Rule

The inclusion-exclusion formulas provide a method for computing the probability of a union of events in terms of the probabilities of the various intersections of the events. The formula is useful because often the probabilities of the intersections are easier to compute. Interestingly, however, the same formula works for computing the probability of an intersection of events in terms of the probabilities of the various unions of the events. This version is rarely stated, because it's simply not that useful. We start with two events.

If A, B are events then $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B)$.




Figure 2.3.8: The probability of the union of two events

Here is the complementary result for the intersection in terms of unions:

If A, B are events then $\mathbb{P}(A \cap B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cup B)$.

Proof

This follows immediately from the previous formula by rearranging the terms

Next we consider three events.

If A, B, C are events then $\mathbb{P}(A \cup B \cup C) = \mathbb{P}(A) + \mathbb{P}(B) + \mathbb{P}(C) - \mathbb{P}(A \cap B) - \mathbb{P}(A \cap C) - \mathbb{P}(B \cap C) + \mathbb{P}(A \cap B \cap C)$.

Analytic Proof

First note that $A \cup B \cup C = (A \cup B) \cup [C \setminus (A \cup B)]$. The event in parentheses and the event in square brackets are disjoint. Thus, using the additivity axiom and the difference rule,

$$\mathbb{P}(A \cup B \cup C) = \mathbb{P}(A \cup B) + \mathbb{P}(C) - \mathbb{P}[C \cap (A \cup B)] = \mathbb{P}(A \cup B) + \mathbb{P}(C) - \mathbb{P}[(C \cap A) \cup (C \cap B)] \quad (2.3.19)$$

Using the [inclusion-exclusion rule for two events](#) (twice) we have

$$\mathbb{P}(A \cup B \cup C) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B) + \mathbb{P}(C) - [\mathbb{P}(C \cap A) + \mathbb{P}(C \cap B) - \mathbb{P}(A \cap B \cap C)] \quad (2.3.20)$$

Proof by accounting



Figure 2.3.8: The probability of the union of three events

Here is the complementary result for the probability of an intersection in terms of the probabilities of the unions:

If A, B, C are events then $\mathbb{P}(A \cap B \cap C) = \mathbb{P}(A) + \mathbb{P}(B) + \mathbb{P}(C) - \mathbb{P}(A \cup B) - \mathbb{P}(A \cup C) - \mathbb{P}(B \cup C) + \mathbb{P}(A \cup B \cup C)$.

Proof

This follows from solving for $\mathbb{P}(A \cap B \cap C)$ in the previous result, and then using the [result for two events](#) on $\mathbb{P}(A \cap B)$, $\mathbb{P}(B \cap C)$, and $\mathbb{P}(A \cap C)$.

The inclusion-exclusion formulas for two and three events can be generalized to n events. For the remainder of this discussion, suppose that $\{A_i : i \in I\}$ is a collection of events where I is an index set with $\#(I) = n$.

The general *inclusion-exclusion formula* for the probability of a union.

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mathbb{P}\left(\bigcap_{j \in J} A_j\right) \quad (2.3.21)$$

Proof by induction

The proof is by induction on n . We have already established the formula for $n = 2$ and $n = 3$. Thus, suppose that the inclusion-exclusion formula holds for a given n , and suppose that $(A_1, A_2, \dots, A_{n+1})$ is a sequence of $n + 1$ events. Then

$$\bigcup_{i=1}^{n+1} A_i = \left(\bigcup_{i=1}^n A_i \right) \cup \left[A_{n+1} \setminus \left(\bigcup_{i=1}^n A_i \right) \right] \quad (2.3.22)$$

As before, the event in parentheses and the event in square brackets are disjoint. Thus using the additivity axiom, the difference rule, and the distributive rule we have

$$\mathbb{P} \left(\bigcup_{i=1}^{n+1} A_i \right) = \mathbb{P} \left(\bigcup_{i=1}^n A_i \right) + \mathbb{P}(A_{n+1}) - \mathbb{P} \left(\bigcup_{i=1}^n (A_{n+1} \cap A_i) \right) \quad (2.3.23)$$

By the induction hypothesis, the inclusion-exclusion formula holds for each union of n events on the right. Applying the formula and simplifying gives the inclusion-exclusion formula for $n + 1$ events.

Proof by accounting

This is the general version of the same argument we used [above](#) for 3 events. $\bigcup_{i \in I} A_i$ is the union of the disjoint events of the form $(\bigcap_{i \in K} A_i) \cap (\bigcap_{i \in K^c} A_i)$ where K is a nonempty subset of the index set I . In the inclusion-exclusion formula, the event corresponding to a given K is measured in $\mathbb{P}(\bigcap_{j \in J} A_j)$ for every nonempty $J \subseteq K$. Suppose that $\#(K) = k$. Accounting for the positive and negative signs, the net measurement is $\sum_{j=1}^k (-1)^{j-1} \binom{k}{j} = 1$.

Here is the complementary result for the probability of an intersection in terms of the probabilities of the various unions:

The general *inclusion-exclusion formula* for the probability of an intersection.

$$\mathbb{P} \left(\bigcap_{i \in I} A_i \right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mathbb{P} \left(\bigcup_{j \in J} A_j \right) \quad (2.3.24)$$

The general inclusion-exclusion formulas are not worth remembering *in detail*, but only *in pattern*. For the probability of a union, we start with the sum of the probabilities of the events, then subtract the probabilities of all of the paired intersections, then add the probabilities of the third-order intersections, and so forth, alternating signs, until we get to the probability of the intersection of all of the events.

The *general Bonferroni inequalities* (for a union) state that if sum on the right in the general inclusion-exclusion formula is truncated, then the truncated sum is an *upper bound* or a *lower bound* for the probability on the left, depending on whether the last term has a positive or negative sign. Here is the result stated explicitly:

Suppose that $m \in \{1, 2, \dots, n-1\}$. Then

1. $\mathbb{P} \left(\bigcup_{i \in I} A_i \right) \leq \sum_{k=1}^m (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mathbb{P} \left(\bigcap_{j \in J} A_j \right)$ if m is odd.
2. $\mathbb{P} \left(\bigcup_{i \in I} A_i \right) \geq \sum_{k=1}^m (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mathbb{P} \left(\bigcap_{j \in J} A_j \right)$ if m is even.

Proof

Let $P_k = \sum_{J \subseteq I, \#(J)=k} \mathbb{P} \left(\bigcap_{j \in J} A_j \right)$, the absolute value of the k th term in the inclusion-exclusion formula. The result follows since the inclusion-exclusion formula is an alternating series, and P_k is decreasing in k .

More elegant proofs of the inclusion-exclusion formula and the Bonferroni inequalities can be constructed using expected value.

Note that there is a probability term in the inclusion-exclusion formulas for every nonempty subset J of the index set I , with either a positive or negative sign, and hence there are $2^n - 1$ such terms. These probabilities suffice to compute the probability of *any* event that can be constructed from the given events, not just the union or the intersection.

The probability of any event that can be constructed from $\{A_i : i \in I\}$ can be computed from either of the following collections of $2^n - 1$ probabilities:

1. $\mathbb{P} \left(\bigcap_{j \in J} A_j \right)$ where J is a nonempty subset of I .
2. $\mathbb{P} \left(\bigcup_{j \in J} A_j \right)$ where J is a nonempty subset of I .

Remark

If you go back and look at your proofs of the [rules of probability](#) above, you will see that they hold for any finite measure μ , not just probability. The only change is that the number 1 is replaced by $\mu(S)$. In particular, the inclusion-exclusion rule is as important in combinatorics (the study of counting measure) as it is in probability.

Examples and Applications

Probability Rules

Suppose that A and B are events in an experiment with $\mathbb{P}(A) = \frac{1}{3}$, $\mathbb{P}(B) = \frac{1}{4}$, $\mathbb{P}(A \cap B) = \frac{1}{10}$. Express each of the following events in the language of the experiment and find its probability:

1. $A \setminus B$
2. $A \cup B$
3. $A^c \cup B^c$
4. $A^c \cap B^c$
5. $A \cup B^c$

Answer

1. A occurs but not B . $\frac{7}{30}$
2. A or B occurs. $\frac{29}{60}$
3. One of the events does not occur. $\frac{9}{10}$
4. Neither event occurs. $\frac{31}{60}$
5. Either A occurs or B does not occur. $\frac{17}{20}$

Suppose that A , B , and C are events in an experiment with $\mathbb{P}(A) = 0.3$, $\mathbb{P}(B) = 0.2$, $\mathbb{P}(C) = 0.4$, $\mathbb{P}(A \cap B) = 0.04$, $\mathbb{P}(A \cap C) = 0.1$, $\mathbb{P}(B \cap C) = 0.1$, $\mathbb{P}(A \cap B \cap C) = 0.01$. Express each of the following events in set notation and find its probability:

1. At least one of the three events occurs.
2. None of the three events occurs.
3. Exactly one of the three events occurs.
4. Exactly two of the three events occur.

Answer

1. $\mathbb{P}(A \cup B \cup C) = 0.67$
2. $\mathbb{P}[(A \cup B \cup C)^c] = 0.37$
3. $\mathbb{P}[(A \cap B^c \cap C^c) \cup (A^c \cap B \cap C^c) \cup (A^c \cap B^c \cap C)] = 0.45$
4. $\mathbb{P}[(A \cap B \cap C^c) \cup (A \cap B^c \cap C) \cup (A^c \cap B \cap C)] = 0.21$

Suppose that A and B are events in an experiment with $\mathbb{P}(A \setminus B) = \frac{1}{6}$, $\mathbb{P}(B \setminus A) = \frac{1}{4}$, and $\mathbb{P}(A \cap B) = \frac{1}{12}$. Find the probability of each of the following events:

1. A
2. B
3. $A \cup B$
4. $A^c \cup B^c$
5. $A^c \cap B^c$

Answer

1. $\frac{1}{4}$
2. $\frac{1}{3}$
3. $\frac{1}{2}$
4. $\frac{11}{12}$

5. $\frac{1}{2}$

Suppose that A and B are events in an experiment with $\mathbb{P}(A) = \frac{2}{5}$, $\mathbb{P}(A \cup B) = \frac{7}{10}$, and $\mathbb{P}(A \cap B) = \frac{1}{6}$. Find the probability of each of the following events:

1. B
2. $A \setminus B$
3. $B \setminus A$
4. $A^c \cup B^c$
5. $A^c \cap B^c$

Answer

1. $\frac{7}{15}$
2. $\frac{7}{30}$
3. $\frac{3}{10}$
4. $\frac{5}{6}$
5. $\frac{3}{10}$

Suppose that A , B , and C are events in an experiment with $\mathbb{P}(A) = \frac{1}{3}$, $\mathbb{P}(B) = \frac{1}{4}$, $\mathbb{P}(C) = \frac{1}{5}$.

1. Use [Boole's inequality](#) to find an upper bound for $\mathbb{P}(A \cup B \cup C)$.
2. Use [Bonferroni's inequality](#) to find a lower bound for $\mathbb{P}(A \cap B \cap C)$.

Answer

1. $\frac{47}{60}$
2. $-\frac{83}{60}$, not helpful.

Open the simple probability experiment.

1. Note the 16 events that can be constructed from A and B using the set operations of union, intersection, and complement.
2. Given $\mathbb{P}(A)$, $\mathbb{P}(B)$, and $\mathbb{P}(A \cap B)$ in the table, use the rules of probability to verify the probabilities of the other events.
3. Run the experiment 1000 times and compare the relative frequencies of the events with the probabilities of the events.

Suppose that A , B , and C are events in a random experiment with $\mathbb{P}(A) = 1/4$, $\mathbb{P}(B) = 1/3$, $\mathbb{P}(C) = 1/6$, $\mathbb{P}(A \cap B) = 1/18$, $\mathbb{P}(A \cap C) = 1/16$, $\mathbb{P}(B \cap C) = 1/12$, and $\mathbb{P}(A \cap B \cap C) = 1/24$. Find the probabilities of the various unions:

1. $A \cup B$
2. $A \cup C$
3. $B \cup C$
4. $A \cup B \cup C$

Answer

1. $19/36$
2. $17/48$
3. $5/12$
4. $85/144$

Suppose that A , B , and C are events in a random experiment with $\mathbb{P}(A) = 1/4$, $\mathbb{P}(B) = 1/4$, $\mathbb{P}(C) = 5/16$, $\mathbb{P}(A \cup B) = 7/16$, $\mathbb{P}(A \cup C) = 23/48$, $\mathbb{P}(B \cup C) = 11/24$, and $\mathbb{P}(A \cup B \cup C) = 7/12$. Find the probabilities of the various intersections:

1. $A \cap B$
2. $A \cap C$
3. $B \cap C$

4. $A \cap B \cap C$

Answer

1. $1/16$
2. $1/12$
3. $5/48$
4. $1/48$

Suppose that A , B , and C are events in a random experiment. Explicitly give all of the Bonferroni inequalities for $\mathbb{P}(A \cup B \cup C)$

Proof

1. $\mathbb{P}(A \cup B \cup C) \leq \mathbb{P}(A) + \mathbb{P}(B) + \mathbb{P}(C)$
2. $\mathbb{P}(A \cup B \cup C) \geq \mathbb{P}(A) + \mathbb{P}(B) + \mathbb{P}(C) - \mathbb{P}(A \cap B) - \mathbb{P}(A \cap C) - \mathbb{P}(B \cap C)$
3. $\mathbb{P}(A \cup B \cup C) = \mathbb{P}(A) + \mathbb{P}(B) + \mathbb{P}(C) - \mathbb{P}(A \cap B) - \mathbb{P}(A \cap C) - \mathbb{P}(B \cap C) + \mathbb{P}(A \cap B \cap C)$

Coins

Consider the random experiment of tossing a coin n times and recording the sequence of scores $\mathbf{X} = (X_1, X_2, \dots, X_n)$ (where 1 denotes heads and 0 denotes tails). This experiment is a generic example of n Bernoulli trials, named for Jacob Bernoulli. Note that the set of outcomes is $S = \{0, 1\}^n$, the set of *bit strings* of length n . If the coin is *fair*, then presumably, by the very meaning of the word, we have no reason to prefer one point in S over another. Thus, as a modeling assumption, it seems reasonable to give S the [uniform probability distribution](#) in which all outcomes are equally likely.

Suppose that a fair coin is tossed 3 times and the sequence of coin scores is recorded. Let A be the event that the first coin is heads and B the event that there are exactly 2 heads. Give each of the following events in list form, and then compute the probability of the event:

1. A
2. B
3. $A \cap B$
4. $A \cup B$
5. $A^c \cup B^c$
6. $A^c \cap B^c$
7. $A \cup B^c$

Answer

1. $\{100, 101, 110, 111\} \frac{1}{2}$
2. $\{110, 101, 011\} \frac{3}{8}$
3. $\{110, 101\} \frac{1}{4}$
4. $\{100, 101, 110, 111, 011\} \frac{5}{8}$
5. $\{000, 001, 010, 100, 011, 111\} \frac{3}{4}$
6. $\{000, 001, 010\} \frac{3}{8}$
7. $\{100, 101, 110, 111, 000, 010, 001\} \frac{7}{8}$

In the Coin experiment, select 3 coins. Run the experiment 1000 times, updating after every run, and compute the empirical probability of each event in the previous exercise.

Suppose that a fair coin is tossed 4 times and the sequence of scores is recorded. Let Y denote the number of heads. Give the event $\{Y = k\}$ (as a subset of the sample space) in list form, for each $k \in \{0, 1, 2, 3, 4\}$, and then give the probability of the event.

Answer

1. $\{Y = 0\} = \{0000\}, \mathbb{P}(Y = 0) = \frac{1}{16}$
2. $\{Y = 1\} = \{1000, 0100, 0010, 0001\}, \mathbb{P}(Y = 1) = \frac{4}{16}$

3. $\{Y = 2\} = \{1100, 1010, 1001, 0110, 0101, 0011\}$, $\mathbb{P}(Y = 2) = \frac{6}{16}$
4. $\{Y = 3\} = \{1110, 1101, 1011, 0111\}$, $\mathbb{P}(Y = 3) = \frac{4}{16}$
5. $\{Y = 4\} = \{1111\}$, $\mathbb{P}(Y = 4) = \frac{1}{16}$

Suppose that a fair coin is tossed n times and the sequence of scores is recorded. Let Y denote the number of heads.

$$\mathbb{P}(Y = k) = \binom{n}{k} \left(\frac{1}{2}\right)^n, \quad k \in \{0, 1, \dots, n\} \quad (2.3.25)$$

Proof

The number of bit strings of length n is 2^n , and since the coin is fair, these are equally likely. The number of bit strings of length n with exactly k 1's is $\binom{n}{k}$. Hence the probability of 1 occurring exactly k times is $\binom{n}{k}/2^n$.

The distribution of Y in the last exercise is a special case of the *binomial distribution*. The binomial distribution is studied in more detail in the chapter on Bernoulli Trials.

Dice

Consider the experiment of throwing n distinct, k -sided dice (with faces numbered from 1 to k) and recording the sequence of scores $\mathbf{X} = (X_1, X_2, \dots, X_n)$. We can record the outcome as a *sequence* because of the assumption that the dice are distinct; you can think of the dice as somehow labeled from 1 to n , or perhaps with different colors. The special case $k = 6$ corresponds to *standard dice*. In general, note that the set of outcomes is $S = \{1, 2, \dots, k\}^n$. If the dice are *fair*, then again, by the very meaning of the word, we have no reason to prefer one point in S over another, so as a modeling assumption it seems reasonable to give S the *uniform probability distribution*.

Suppose that two fair, standard dice are thrown and the sequence of scores recorded. Let A denote the event that the first die score is less than 3 and B the event that the sum of the dice scores is 6. Give each of the following events in list form and then find the probability of the event.

1. A
2. B
3. $A \cap B$
4. $A \cup B$
5. $B \setminus A$

Answer

1. $\{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6), (2, 1), (2, 2), (2, 3), (2, 4), (2, 5), (2, 6)\}$, $\frac{12}{36}$
2. $\{(1, 5), (5, 1), (2, 4), (4, 2), (3, 3)\}$, $\frac{5}{36}$
3. $\{(1, 5), (2, 4)\}$, $\frac{2}{36}$
4. $\{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6), (2, 1), (2, 2), (2, 3), (2, 4), (2, 5), (2, 6), (5, 1), (4, 2), (3, 3)\}$, $\frac{15}{36}$
5. $\{(5, 1), (4, 2), (3, 3)\}$, $\frac{3}{36}$

In the dice experiment, set $n = 2$. Run the experiment 100 times and compute the empirical probability of each event in the previous exercise.

Consider again the dice experiment with $n = 2$ fair dice. Let S denote the set of outcomes, Y the sum of the scores, U the minimum score, and V the maximum score.

1. Express Y as a function on S and give the set of values.
2. Find $\mathbb{P}(Y = y)$ for each y in the set in part (a).
3. Express U as a function on S and give the set of values.
4. Find $\mathbb{P}(U = u)$ for each u in the set in part (c).
5. Express V as a function on S and give the set of values.
6. Find $\mathbb{P}(V = v)$ for each v in the set in part (e).
7. Find the set of values of (U, V) .

8. Find $\mathbb{P}(U = u, V = v)$ for each (u, v) in the set in part (g).

Answer

Note that $S = \{1, 2, 3, 4, 5, 6\}^2$

1. $Y(x_1, x_2) = x_1 + x_2$ for $(x_1, x_2) \in S$. The set of values is $\{2, 3, \dots, 12\}$

2.	y	2	3	4	5	6	7	8	9	10	11	12
	$\mathbb{P}(Y = y)$	$\frac{1}{36}$	$\frac{2}{36}$	$\frac{3}{36}$	$\frac{4}{36}$	$\frac{5}{36}$	$\frac{6}{36}$	$\frac{5}{36}$	$\frac{4}{36}$	$\frac{3}{36}$	$\frac{2}{36}$	$\frac{1}{36}$

3. $U(x_1, x_2) = \min\{x_1, x_2\}$ for $(x_1, x_2) \in S$. The set of values is $\{1, 2, 3, 4, 5, 6\}$

4.	u	1	2	3	4	5	6
	$\mathbb{P}(U = u)$	$\frac{11}{36}$	$\frac{9}{36}$	$\frac{7}{36}$	$\frac{5}{36}$	$\frac{3}{36}$	$\frac{1}{36}$

5. $V(x_1, x_2) = \max\{x_1, x_2\}$ for $(x_1, x_2) \in S$. The set of values is $\{1, 2, 3, 4, 5, 6\}$

6.	v	1	2	3	4	5	6
	$\mathbb{P}(V = v)$	$\frac{1}{36}$	$\frac{3}{36}$	$\frac{5}{36}$	$\frac{7}{36}$	$\frac{9}{36}$	$\frac{11}{36}$

7. $\{(u, v) \in S : u \leq v\}$

$$8. \mathbb{P}(U = u, V = v) = \begin{cases} \frac{2}{36}, & u < v \\ \frac{1}{36}, & u = v \end{cases}$$

In the previous exercise, note that (U, V) could serve as the outcome vector for the experiment of rolling two standard, fair dice if we do not bother to distinguish the dice (so that we might as well record the smaller score first and then the larger score). Note that this random vector does *not* have a uniform distribution. On the other hand, we might have chosen at the beginning to just record the unordered set of scores and, as a modeling assumption, imposed the uniform distribution on the corresponding set of outcomes. Both models cannot be right, so which model (if either) describes real dice in the real world? It turns out that for real (fair) dice, the *ordered sequence* of scores is uniformly distributed, so real dice *behave* as distinct objects, whether you can tell them apart or not. In the early history of probability, gamblers sometimes got the wrong answers for events involving dice because they mistakenly applied the uniform distribution to the set of unordered scores. It's an important moral. If we are to impose the uniform distribution on a sample space, we need to make sure that it's the *right* sample space.

A pair of fair, standard dice are thrown repeatedly until the sum of the scores is either 5 or 7. Let A denote the event that the sum of the scores on the last throw is 5 rather than 7. Events of this type are important in the game of craps.

1. Suppose that we record the pair of scores on each throw. Give the set of outcomes S and express A as a subset of S .
2. Compute the probability of A in the setting of part (a).
3. Now suppose that we just record the pair of scores on the last throw. Give the set of outcomes T and express A as a subset of T .
4. Compute the probability of A in the setting of parts (c).

Answer

Let $D_5 = \{(1, 4), (2, 3), (3, 2), (4, 1)\}$ $D_7 = \{(1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\}$ $D = D_5 \cup D_7$,
 $C = \{1, 2, 3, 4, 5, 6\}^2 \setminus D$

1. $S = D \cup (C \times D) \cup (C^2 \times D) \cup \dots$, $A = D_5 \cup (C \times D_5) \cup (C^2 \times D_5) \cup \dots$
2. $\frac{2}{5}$
3. $T = D$, $A = D_5$
4. $\frac{2}{5}$

The previous problem shows the importance of defining the set of outcomes appropriately. Sometimes a clever choice of this set (and appropriate modeling assumptions) can turn a difficult problem into an easy one.

Sampling Models

Recall that many random experiments can be thought of as *sampling experiments*. For the general finite sampling model, we start with a population D with m (distinct) objects. We select a sample of n objects from the population, so that the sample space S is the set of possible samples. If we select a sample *at random* then the outcome \mathbf{X} (the random sample) is uniformly distributed on S :

$$\mathbb{P}(\mathbf{X} \in A) = \frac{\#(A)}{\#(S)}, \quad A \subseteq S \quad (2.3.26)$$

Recall from the section on Combinatorial Structures that there are four common types of sampling from a finite population, based on the criteria of *order* and *replacement*.

- If the sampling is with replacement and with regard to order, then the set of samples is the Cartesian power D^n . The number of samples is m^n .
- If the sampling is without replacement and with regard to order, then the set of samples is the set of all *permutations* of size n from D . The number of samples is $m^{(n)} = m(m-1) \cdots (m-n+1)$.
- If the sampling is without replacement and without regard to order, then the set of samples is the set of all *combinations* (or *subsets*) of size n from D . The number of samples is $\binom{m}{n} = m^{(n)}/n!$.
- If the sampling is with replacement and without regard to order, then the set of samples is the set of all *multisets* of size n from D . The number of samples is $\binom{m+n-1}{n}$.

If we sample with replacement, the sample size n can be any positive integer. If we sample without replacement, the sample size cannot exceed the population size, so we must have $n \in \{1, 2, \dots, m\}$.

The basic coin and dice experiments are examples of sampling *with* replacement. If we toss a fair coin n times and record the sequence of scores \mathbf{X} (where as usual, 0 denotes tails and 1 denotes heads), then \mathbf{X} is a random sample of size n chosen with order and with replacement from the population $\{0, 1\}$. Thus, \mathbf{X} is uniformly distributed on $\{0, 1\}^n$. If we throw n (distinct) standard fair dice and record the sequence of scores, then we generate a random sample \mathbf{X} of size n with order and with replacement from the population $\{1, 2, 3, 4, 5, 6\}$. Thus, \mathbf{X} is uniformly distributed on $\{1, 2, 3, 4, 5, 6\}^n$. Of an analogous result would hold for fair, k -sided dice.

Suppose that the sampling is without replacement (the most common case). If we record the *ordered* sample $\mathbf{X} = (X_1, X_2, \dots, X_n)$, then the *unordered* sample $\mathbf{W} = \{X_1, X_2, \dots\}$ is a random variable (that is, a function of \mathbf{X}). On the other hand, if we just record the unordered sample \mathbf{W} in the first place, then we cannot recover the ordered sample.

Suppose that \mathbf{X} is a random sample of size n chosen with order and without replacement from D , so that \mathbf{X} is uniformly distributed on the space of permutations of size n from D . Then \mathbf{W} , the unordered sample, is uniformly distributed on the space of combinations of size n from D . Thus, \mathbf{W} is also a random sample.

Proof

Let \mathbf{w} be a combination of size n from D . Then there are $n!$ permutations of the elements in \mathbf{w} . If A denotes this set of permutations, then $\mathbb{P}(\mathbf{W} = \mathbf{w}) = \mathbb{P}(\mathbf{X} \in A) = n! / m^{(n)} = 1 / \binom{m}{n}$.

The result in the last exercise does not hold if the sampling is with replacement (recall the [exercise above](#) and the discussion that follows). When sampling with replacement, there is no simple relationship between the number of ordered samples and the number of unordered samples.

Sampling From a Dichotomous Population

Suppose again that we have a population D with m (distinct) objects, but suppose now that each object is one of two types—either type 1 or type 0. Such populations are said to be *dichotomous*. Here are some specific examples:

- The population consists of persons, each either *male* or *female*.
- The population consists of voters, each either *democrat* or *republican*.
- The population consists of devices, each either *good* or *defective*.
- The population consists of balls, each either *red* or *green*.

Suppose that the population D has r type 1 objects and hence $m - r$ type 0 objects. Of course, we must have $r \in \{0, 1, \dots, m\}$. Now suppose that we select a sample of size n at random from the population. Note that this model has three parameters: the population size m , the number of type 1 objects in the population r , and the sample size n . Let Y denote the number of type 1 objects in the sample.

If the sampling is without replacement then

$$\mathbb{P}(Y = y) = \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}}, \quad y \in \{0, 1, \dots, n\} \quad (2.3.27)$$

Proof

Recall that the unordered sample is uniformly distributed over the set of combinations of size n chosen from the population. There are $\binom{m}{n}$ such samples. By the multiplication principle, the number of samples with exactly y type 1 objects and $n - y$ type 0 objects is $\binom{r}{y} \binom{m-r}{n-y}$.

In the previous exercise, random variable Y has the *hypergeometric distribution* with parameters m , r , and n . The hypergeometric distribution is studied in more detail in the chapter on Finite Sampling Models.

If the sampling is with replacement then

$$\mathbb{P}(Y = y) = \binom{n}{y} \frac{r^y (m-r)^{n-y}}{m^n} = \binom{n}{y} \left(\frac{r}{m}\right)^y \left(1 - \frac{r}{m}\right)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (2.3.28)$$

Proof

Recall that the ordered sample is uniformly distributed over the set D^n and there are m^n elements in this set. To count the number of samples with exactly y type 1 objects, we use a three-step procedure: first, select the coordinates for the type 1 objects; there are $\binom{n}{y}$ choices. Next select the y type 1 objects for these coordinates; there are r^y choices. Finally select the $n - y$ type 0 objects for the remaining coordinates of the sample; there are $(m - r)^{n-y}$ choices. The result now follows from the multiplication principle.

In the last exercise, random variable Y has the *binomial distribution* with parameters n and $p = \frac{r}{m}$. The binomial distribution is studied in more detail in the chapter on Bernoulli Trials.

Suppose that a group of voters consists of 40 democrats and 30 republicans. A sample 10 voters is chosen at random. Find the probability that the sample contains at least 4 democrats and at least 4 republicans, each of the following cases:

1. The sampling is without replacement.
2. The sampling is with replacement.

Answer

1. $\frac{1\,391\,351\,589}{2\,176\,695\,188} \approx 0.6382$
2. $\frac{24\,509\,952}{40\,353\,607} \approx 0.6074$

Look for other specialized sampling situations in the exercises below.

Urn Models

Drawing balls from an urn is a standard metaphor in probability for sampling from a finite population.

Consider an urn with 30 balls; 10 are red and 20 are green. A sample of 5 balls is chosen at random, without replacement. Let Y denote the number of red balls in the sample. Explicitly compute $\mathbb{P}(Y = y)$ for each $y \in \{0, 1, 2, 3, 4, 5\}$.

answer

y	0	1	2	3	4	5
$\mathbb{P}(Y = y)$	$\frac{2584}{23751}$	$\frac{8075}{23751}$	$\frac{8550}{23751}$	$\frac{3800}{23751}$	$\frac{700}{23751}$	$\frac{42}{23751}$

In the simulation of the ball and urn experiment, select 30 balls with 10 red and 20 green, sample size 5, and sampling without replacement. Run the experiment 1000 times and compare the empirical probabilities with the true probabilities that you computed in the previous exercise.

Consider again an urn with 30 balls; 10 are red and 20 are green. A sample of 5 balls is chosen at random, with replacement. Let Y denote the number of red balls in the sample. Explicitly compute $\mathbb{P}(Y = y)$ for each $y \in \{0, 1, 2, 3, 4, 5\}$.

Answer

y	0	1	2	3	4	5
$\mathbb{P}(Y = y)$	$\frac{32}{243}$	$\frac{80}{243}$	$\frac{80}{243}$	$\frac{40}{243}$	$\frac{10}{243}$	$\frac{1}{243}$

In the simulation of the ball and urn experiment, select 30 balls with 10 red and 20 green, sample size 5, and sampling with replacement. Run the experiment 1000 times and compare the empirical probabilities with the true probabilities that you computed in the previous exercise.

An urn contains 15 balls: 6 are red, 5 are green, and 4 are blue. Three balls are chosen at random, without replacement.

1. Find the probability that the chosen balls are all the same color.
2. Find the probability that the chosen balls are all different colors.

Answer

1. $\frac{34}{455}$
2. $\frac{120}{455}$

Suppose again that an urn contains 15 balls: 6 are red, 5 are green, and 4 are blue. Three balls are chosen at random, with replacement.

1. Find the probability that the chosen balls are all the same color.
2. Find the probability that the chosen balls are all different colors.

Answer

1. $\frac{405}{3375}$
2. $\frac{720}{3375}$

Cards

Recall that a standard *card deck* can be modeled by the product set

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (2.3.29)$$

where the first coordinate encodes the *denomination* or *kind* (ace, 2–10, jack, queen, king) and where the second coordinate encodes the *suit* (clubs, diamonds, hearts, spades). Sometimes we represent a card as a *string* rather than an ordered pair (for example $q\heartsuit$ for the queen of hearts).

Card games involve choosing a [random sample](#) without replacement from the deck D , which plays the role of the population. Thus, the basic *card experiment* consists of dealing n cards from a standard deck without replacement; in this special context, the sample of cards is often referred to as a *hand*. Just as in the general sampling model, if we record the *ordered* hand $\mathbf{X} = (X_1, X_2, \dots, X_n)$, then the *unordered* hand $\mathbf{W} = \{X_1, X_2, \dots, X_n\}$ is a random variable (that is, a function of \mathbf{X}). On the other hand, if we just record the unordered hand \mathbf{W} in the first place, then we cannot recover the ordered hand. Finally, recall that $n = 5$ is the *poker experiment* and $n = 13$ is the *bridge experiment*. The game of poker is treated in more detail in the chapter on Games of Chance. By the way, it takes about 7 standard riffle shuffles to randomize a deck of cards.

Suppose that 2 cards are dealt from a well-shuffled deck and the sequence of cards is recorded. For $i \in \{1, 2\}$, let H_i denote the event that card i is a heart. Find the probability of each of the following events.

1. H_1

2. $H_1 \cap H_2$
3. $H_2 \setminus H_1$
4. H_2
5. $H_1 \setminus H_2$
6. $H_1 \cup H_2$

Answer

1. $\frac{1}{4}$
2. $\frac{1}{17}$
3. $\frac{13}{68}$
4. $\frac{1}{4}$
5. $\frac{13}{68}$
6. $\frac{15}{34}$

Think about the results in the previous exercise, and suppose that we continue dealing cards. Note that in computing the probability of H_i , you could conceptually reduce the experiment to dealing a single card. Note also that the probabilities do not depend on the order in which the cards are dealt. For example, the probability of an event involving the 1st, 2nd and 3rd cards is the same as the probability of the corresponding event involving the 25th, 17th, and 40th cards. Technically, the cards are *exchangeable*. Here's another way to think of this concept: Suppose that the cards are dealt onto a table in some pattern, but you are not allowed to view the process. Then no experiment that you can devise will give you any information about the order in which the cards were dealt.

In the card experiment, set $n = 2$. Run the experiment 100 times and compute the empirical probability of each event in the previous exercise

In the poker experiment, find the probability of each of the following events:

1. The hand is a *full house* (3 cards of one kind and 2 cards of another kind).
2. The hand has *four of a kind* (4 cards of one kind and 1 of another kind).
3. The cards are all in the same suit (thus, the hand is either a *flush* or a *straight flush*).

Answer

1. $\frac{3744}{2\,598\,960} \approx 0.001441$
2. $\frac{624}{2\,598\,960} \approx 0.000240$
3. $\frac{5148}{2\,598\,960} \approx 0.001981$

Run the poker experiment 10000 times, updating every 10 runs. Compute the empirical probability of each event in the previous problem.

Find the probability that a bridge hand will contain no *honor cards* that is, no cards of denomination 10, jack, queen, king, or ace. Such a hand is called a *Yarborough*, in honor of the second Earl of Yarborough.

Answer

$$\frac{347\,373\,600}{635\,013\,559\,600} \approx 0.000547$$

Find the probability that a bridge hand will contain

1. Exactly 4 hearts.
2. Exactly 4 hearts and 3 spades.
3. Exactly 4 hearts, 3 spades, and 2 clubs.

Answer

1. $\frac{151\,519\,319\,380}{635\,013\,559\,600} \approx 0.2386$
2. $\frac{47\,079\,732\,700}{635\,013\,559\,600} \approx 0.0741$
3. $\frac{11\,404\,407\,300}{635\,013\,559\,600} \approx 0.0179$

A card hand that contains no cards in a particular suit is said to be *void* in that suit. Use the [inclusion-exclusion rule](#) to find the probability of each of the following events:

1. A poker hand is void in at least one suit.
2. A bridge hand is void in at least one suit.

Answer

1. $\frac{1\,913\,496}{2\,598\,960} \approx 0.7363$
2. $\frac{32\,427\,298\,180}{635\,013\,559\,600} \approx 0.051$

Birthdays

The following problem is known as the *birthday problem*, and is famous because the results are rather surprising at first.

Suppose that n persons are selected and their birthdays recorded (we will ignore leap years). Let A denote the event that the birthdays are distinct, so that A^c is the event that there is at least one duplication in the birthdays.

1. Define an appropriate sample space and probability measure. State the assumptions you are making.
2. Find $P(A)$ and $\mathbb{P}(A^c)$ in terms of the parameter n .
3. Explicitly compute $P(A)$ and $P(A^c)$ for $n \in \{10, 20, 30, 40, 50\}$

Answer

1. The set of outcomes is $S = D^n$ where D is the set of days of the year. We assume that the outcomes are equally likely, so that S has the uniform distribution.
2. $\#(A) = 365^{(n)}$, so $\mathbb{P}(A) = 365^{(n)}/365^n$ and $\mathbb{P}(A^c) = 1 - 365^{(n)}/365^n$

3. n	$\mathbb{P}(A)$	$\mathbb{P}(A^c)$
10	0.883	0.117
20	0.589	0.411
30	0.294	0.706
40	0.109	0.891
50	0.006	0.994

The small value of $\mathbb{P}(A)$ for relatively small sample sizes n is striking, but is due mathematically to the fact that 365^n grows much faster than $365^{(n)}$ as n increases. The birthday problem is treated in more generality in the chapter on Finite Sampling Models.

Suppose that 4 persons are selected and their birth *months* recorded. Assuming an appropriate uniform distribution, find the probability that the months are distinct.

Answer

$$\frac{11880}{20736} \approx 0.573$$

Continuous Uniform Distributions

Recall that in Buffon's coin experiment, a coin with radius $r \leq \frac{1}{2}$ is tossed “randomly” on a floor with square tiles of side length 1, and the coordinates (X, Y) of the center of the coin are recorded, relative to axes through the center of the square in which the coin lands (with the axes parallel to the sides of the square, of course). Let A denote the event that the coin does not touch the sides of the square.

1. Define the set of outcomes S mathematically and sketch S .
2. Argue that (X, Y) is uniformly distributed on S .
3. Express A in terms of the outcome variables (X, Y) and sketch A .
4. Find $\mathbb{P}(A)$.

5. Find $\mathbb{P}(A^c)$.

Answer

1. $S = [-\frac{1}{2}, \frac{1}{2}]^2$
2. Since the coin is tossed “randomly”, no region of S should be preferred over any other.
3. $\{r - \frac{1}{2} < X < \frac{1}{2} - r, r - \frac{1}{2} < Y < \frac{1}{2} - r\}$
4. $\mathbb{P}(A) = (1 - 2r)^2$
5. $\mathbb{P}(A^c) = 1 - (1 - 2r)^2$

In Buffon's coin experiment, set $r = 0.2$. Run the experiment 100 times and compute the empirical probability of each event in the previous exercise.

A point (X, Y) is chosen at random in the circular region $S \subset \mathbb{R}^2$ of radius 1, centered at the origin. Let A denote the event that the point is in the inscribed square region centered at the origin, with sides parallel to the coordinate axes, and let B denote the event that the point is in the inscribed square with vertices $(\pm 1, 0)$, $(0, \pm 1)$. Sketch each of the following events as a subset of S , and find the probability of the event.

1. A
2. B
3. $A \cap B^c$
4. $B \cap A^c$
5. $A \cap B$
6. $A \cup B$

Answer

1. $2/\pi$
2. $2/\pi$
3. $(6 - 4\sqrt{2})/\pi$
4. $(6 - 4\sqrt{2})/\pi$
5. $4(\sqrt{2} - 1)/\pi$
6. $4(2 - \sqrt{2})/\pi$

Suppose a point (X, Y) is chosen at random in the circular region $S \subseteq \mathbb{R}^2$ of radius 12, centered at the origin. Let R denote the distance from the origin to the point. Sketch each of the following events as a subset of S , and compute the probability of the event. Is R uniformly distributed on the interval $[0, 12]$?

1. $\{R \leq 3\}$
2. $\{3 < R \leq 6\}$
3. $\{6 < R \leq 9\}$
4. $\{9 < R \leq 12\}$

Answer

No, R is not uniformly distributed on $[0, 12]$.

1. $\frac{1}{16}$
2. $\frac{3}{16}$
3. $\frac{5}{16}$
4. $\frac{7}{16}$

In the simple probability experiment, points are generated according to the uniform distribution on a rectangle. Move and resize the events A and B and note how the probabilities of the various events change. Create each of the following configurations. In each case, run the experiment 1000 times and compare the relative frequencies of the events to the probabilities of the events.

1. A and B in general position

2. A and B disjoint
3. $A \subseteq B$
4. $B \subseteq A$

Genetics

Please refer to the discussion of genetics in the section on random experiments if you need to review some of the definitions in this subsection.

Recall first that the *ABO blood type* in humans is determined by three alleles: a , b , and o . Furthermore, a and b are co-dominant and o is recessive. Suppose that the probability distribution for the set of blood genotypes in a certain population is given in the following table:

Genotype	aa	ab	ao	bb	bo	oo
Probability	0.050	0.038	0.310	0.007	0.116	0.479

A person is chosen at random from the population. Let A , B , AB , and O be the events that the person is type A , type B , type AB , and type O respectively. Let H be the event that the person is homozygous and D the event that the person has an o allele. Find the probability of the following events:

1. A
2. B
3. AB
4. O
5. H
6. D
7. $H \cup D$
8. D^c

Answer

1. 0.360
2. 0.123
3. 0.038
4. 0.479
5. 0.536
6. 0.905
7. 0.962
8. 0.095

Suppose next that pod color in certain type of pea plant is determined by a gene with two alleles: g for green and y for yellow, and that g is dominant.

Let G be the event that a child plant has green pods. Find $\mathbb{P}(G)$ in each of the following cases:

1. At least one parent is type gg .
2. Both parents are type yy .
3. Both parents are type gy .
4. One parent is type yy and the other is type gy .

Answer

1. 1
2. 0
3. $\frac{3}{4}$
4. $\frac{1}{2}$

Next consider a sex-linked hereditary disorder in humans (such as colorblindness or hemophilia). Let h denote the healthy allele and d the defective allele for the gene linked to the disorder. Recall that d is recessive for women.

Let B be the event that a son has the disorder, C the event that a daughter is a healthy carrier, and D the event that a daughter has the disease. Find $\mathbb{P}(B)$, $\mathbb{P}(C)$ and $\mathbb{P}(D)$ in each of the following cases:

1. The mother and father are normal.
2. The mother is a healthy carrier and the father is normal.
3. The mother is normal and the father has the disorder.
4. The mother is a healthy carrier and the father has the disorder.
5. The mother has the disorder and the father is normal.
6. The mother and father both have the disorder.

Answer

1. 0, 0, 0
2. $1/2$, 0, $1/2$
3. 0, $1/2$, 0
4. $1/2$, $1/2$, $1/2$
5. 1, $1/2$, 0
6. 1, 0, 1

From this exercise, note that transmission of the disorder to a daughter can only occur if the mother is at least a carrier and the father has the disorder. In ordinary large populations, this is a unusual intersection of events, and thus sex-linked hereditary disorders are typically much less common in women than in men. In brief, women are protected by the extra X chromosome.

Radioactive Emissions

Suppose that T denotes the time between emissions (in milliseconds) for a certain type of radioactive material, and that T has the following probability distribution, defined for measurable $A \subseteq [0, \infty)$ by

$$\mathbb{P}(T \in A) = \int_A e^{-t} dt \quad (2.3.30)$$

1. Show that this really does define a probability distribution.
2. Find $\mathbb{P}(T > 3)$.
3. Find $\mathbb{P}(2 < T < 4)$.

Answer

1. Note that $\int_0^\infty e^{-t} dt = 1$
2. e^{-3}
3. $e^{-2} - e^{-4}$

Suppose that N denotes the number of emissions in a one millisecond interval for a certain type of radioactive material, and that N has the following probability distribution:

$$\mathbb{P}(N \in A) = \sum_{n \in A} \frac{e^{-1}}{n!}, \quad A \subseteq \mathbb{N} \quad (2.3.31)$$

1. Show that this really does define a probability distribution.
2. Find $\mathbb{P}(N \geq 3)$.
3. Find $\mathbb{P}(2 \leq N \leq 4)$.

Answer

1. Note that $\sum_{n=0}^\infty \frac{e^{-1}}{n!} = 1$
2. $1 - \frac{5}{2}e^{-1}$
3. $\frac{17}{24}e^{-1}$

The probability distribution that governs the [time between emissions](#) is a special case of the *exponential distribution*, while the probability distribution that governs the [number of emissions](#) is a special case of the *Poisson distribution*, named for Simeon Poisson. The exponential distribution and the Poisson distribution are studied in more detail in the chapter on the Poisson process.

Matching

Suppose that at an absented-minded secretary prepares 4 letters and matching envelopes to send to 4 different persons, but then stuffs the letters into the envelopes randomly. Find the probability of the event A that at least one letter is in the proper envelope.

Solution

Note first that the set of outcomes S can be taken to be the set of permutations of $\{1, 2, 3, 4\}$. For $x \in S$, x_i is the number of the envelope containing the i th letter. Clearly S should be given the uniform probability distribution. Next note that $A = A_1 \cup A_2 \cup A_3 \cup A_4$ where A_i is the event that the i th letter is inserted into the i th envelope. Using the inclusion-exclusion rule gives $\mathbb{P}(A) = \frac{5}{8}$.

This exercise is an example of the *matching problem*, originally formulated and studied by Pierre Remond Montmort. A complete analysis of the matching problem is given in the chapter on Finite Sampling Models.

In the simulation of the matching experiment select $n = 4$. Run the experiment 1000 times and compute the relative frequency of the event that at least one match occurs.

Data Analysis Exercises

For the M&M data set, let R denote the event that a bag has at least 10 red candies, T the event that a bag has at least 57 candies total, and W the event that a bag weighs at least 50 grams. Find the empirical probability the following events:

1. R
2. T
3. W
4. $R \cap T$
5. $T \setminus W$

Answer

1. $\frac{13}{30}$
2. $\frac{19}{30}$
3. $\frac{9}{30}$
4. $\frac{9}{30}$
5. $\frac{11}{30}$

For the cicada data, let W denote the event that a cicada weighs at least 0.20 grams, F the event that a cicada is female, and T the event that a cicada is type tredecia. Find the empirical probability of each of the following:

1. W
2. F
3. T
4. $W \cap F$
5. $F \cup T \cup W$

Answer

1. $\frac{37}{104}$
2. $\frac{59}{104}$
3. $\frac{44}{104}$
4. $\frac{34}{104}$
5. $\frac{85}{104}$

This page titled [2.3: Probability Measures](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.4: Conditional Probability

The purpose of this section is to study how probabilities are updated in light of new information, clearly an absolutely essential topic. If you are a new student of probability, you may want to skip the technical details.

Definitions and Interpretations

The Basic Definition

As usual, we start with a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$. Thus, S is the set of outcomes, \mathcal{S} the collection of events, and \mathbb{P} the probability measure on the sample space (S, \mathcal{S}) . Suppose now that we know that an event B has occurred. In general, this information should clearly alter the probabilities that we assign to other events. In particular, if A is another event then A occurs if and only if A and B occur; effectively, the sample space has been reduced to B . Thus, the probability of A , given that we know B has occurred, should be proportional to $\mathbb{P}(A \cap B)$.



Figure 2.4.1: Events A and B

However, conditional probability, given that B has occurred, should still be a probability measure, that is, it must satisfy the axioms of probability. This forces the proportionality constant to be $1/\mathbb{P}(B)$. Thus, we are led inexorably to the following definition:

Let A and B be events with $\mathbb{P}(B) > 0$. The *conditional probability* of A given B is defined to be

$$\mathbb{P}(A | B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \quad (2.4.1)$$

The Law of Large Numbers

The [definition](#) above was based on the axiomatic definition of probability. Let's explore the idea of conditional probability from the less formal and more intuitive notion of relative frequency (the law of large numbers). Thus, suppose that we run the experiment repeatedly. For $n \in \mathbb{N}_+$ and an event E , let $N_n(E)$ denote the number of times E occurs (the frequency of E) in the first n runs. Note that $N_n(E)$ is a random variable in the compound experiment that consists of replicating the original experiment. In particular, its value is unknown until we actually run the experiment n times.

If $N_n(B)$ is large, the conditional probability that A has occurred, given that B has occurred, should be close to the *conditional relative frequency* of A given B , namely the relative frequency of A for the runs on which B occurred: $N_n(A \cap B)/N_n(B)$. But note that

$$\frac{N_n(A \cap B)}{N_n(B)} = \frac{N_n(A \cap B)/n}{N_n(B)/n} \quad (2.4.2)$$

The numerator and denominator of the main fraction on the right are the relative frequencies of $A \cap B$ and B , respectively. So by the law of large numbers again, $N_n(A \cap B)/n \rightarrow \mathbb{P}(A \cap B)$ as $n \rightarrow \infty$ and $N_n(B)/n \rightarrow \mathbb{P}(B)$ as $n \rightarrow \infty$. Hence

$$\frac{N_n(A \cap B)}{N_n(B)} \rightarrow \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \text{ as } n \rightarrow \infty \quad (2.4.3)$$

and we are led again to the [definition](#) above.

In some cases, conditional probabilities can be computed *directly*, by effectively reducing the sample space to the given event. In other cases, the formula in the mathematical definition is better. In some cases, conditional probabilities are known from modeling assumptions, and then are used to compute other probabilities. We will see examples of all of these situations in the [computational exercises](#) below.

It's very important that you not confuse $\mathbb{P}(A | B)$, the probability of A given B , with $\mathbb{P}(B | A)$, the probability of B given A . Making that mistake is known as the *fallacy of the transposed conditional*. (How embarrassing!)

Conditional Distributions

Suppose that X is a random variable for the experiment with values in T . Mathematically, X is a function from S into T , and $\{X \in A\}$ denotes the event $\{s \in S : X(s) \in A\}$ for $A \subseteq T$. Intuitively, X is a variable of interest in the experiment, and every meaningful statement about X defines an event. Recall that the probability distribution of X is the probability measure on T given by

$$A \mapsto \mathbb{P}(X \in A), \quad A \subseteq T \quad (2.4.4)$$

This has a natural extension to a conditional distribution, given an event.

If B is an event with $\mathbb{P}(B) > 0$, then the *conditional distribution of X given B* is the probability measure on T given by

$$A \mapsto \mathbb{P}(X \in A \mid B), \quad A \subseteq T \quad (2.4.5)$$

Details

Recall that T will come with a σ -algebra of admissible subsets so that (T, \mathcal{T}) is a measurable space, just like the sample space (S, \mathcal{S}) . Random variable X is required to be measurable as a function from S into T . This ensures that $\{X \in A\}$ is a valid event for each $A \in \mathcal{T}$, so that the definition makes sense.

Basic Theory

Preliminary Results

Our first result is of fundamental importance, and indeed was a crucial part of the argument for the definition of conditional probability.

Suppose again that B is an event with $\mathbb{P}(B) > 0$. Then $A \mapsto \mathbb{P}(A \mid B)$ is a probability measure on S .

Proof

Clearly $\mathbb{P}(A \mid B) \geq 0$ for every event A , and $\mathbb{P}(S \mid B) = 1$. Thus, suppose that $\{A_i : i \in I\}$ is a countable collection of pairwise disjoint events. Then

$$\mathbb{P}\left(\bigcup_{i \in I} A_i \mid B\right) = \frac{1}{\mathbb{P}(B)} \mathbb{P}\left[\left(\bigcup_{i \in I} A_i\right) \cap B\right] = \frac{1}{\mathbb{P}(B)} \mathbb{P}\left(\bigcup_{i \in I} (A_i \cap B)\right) \quad (2.4.6)$$

But the collection of events $\{A_i \cap B : i \in I\}$ is also pairwise disjoint, so

$$\mathbb{P}\left(\bigcup_{i \in I} A_i \mid B\right) = \frac{1}{\mathbb{P}(B)} \sum_{i \in I} \mathbb{P}(A_i \cap B) = \sum_{i \in I} \frac{\mathbb{P}(A_i \cap B)}{\mathbb{P}(B)} = \sum_{i \in I} \mathbb{P}(A_i \mid B) \quad (2.4.7)$$

It's hard to overstate the importance of the [last result](#) because this theorem means that any result that holds for probability measures in general holds for conditional probability, as long as the conditioning event remains fixed. In particular the basic probability rules in the section on Probability Measure have analogs for conditional probability. To give two examples,

$$\mathbb{P}(A^c \mid B) = 1 - \mathbb{P}(A \mid B) \quad (2.4.8)$$

$$\mathbb{P}(A_1 \cup A_2 \mid B) = \mathbb{P}(A_1 \mid B) + \mathbb{P}(A_2 \mid B) - \mathbb{P}(A_1 \cap A_2 \mid B) \quad (2.4.9)$$

By the same token, it follows that the conditional distribution of a random variable with values in T , given in [above](#), really does define a probability distribution on T . No further proof is necessary. Our next results are very simple.

Suppose that A and B are events with $\mathbb{P}(B) > 0$.

1. If $B \subseteq A$ then $\mathbb{P}(A \mid B) = 1$.
2. If $A \subseteq B$ then $\mathbb{P}(A \mid B) = \mathbb{P}(A)/\mathbb{P}(B)$.
3. If A and B are disjoint then $\mathbb{P}(A \mid B) = 0$.

Proof

These results follow directly from the definition of conditional probability. In part (a), note that $A \cap B = B$. In part (b) note that $A \cap B = A$. In part (c) note that $A \cap B = \emptyset$.

Parts (a) and (c) certainly make sense. Suppose that we know that event B has occurred. If $B \subseteq A$ then A becomes a certain event. If $A \cap B = \emptyset$ then A becomes an impossible event. A conditional probability can be computed relative to a probability measure that is itself a conditional probability measure. The following result is a consistency condition.

Suppose that A , B , and C are events with $\mathbb{P}(B \cap C) > 0$. The probability of A given B , relative to $\mathbb{P}(\cdot | C)$, is the same as the probability of A given B and C (relative to \mathbb{P}). That is,

$$\frac{\mathbb{P}(A \cap B | C)}{\mathbb{P}(B | C)} = \mathbb{P}(A | B \cap C) \quad (2.4.10)$$

Proof

From the definition,

$$\frac{\mathbb{P}(A \cap B | C)}{\mathbb{P}(B | C)} = \frac{\mathbb{P}(A \cap B \cap C) / \mathbb{P}(C)}{\mathbb{P}(B \cap C) / \mathbb{P}(C)} = \frac{\mathbb{P}(A \cap B \cap C)}{\mathbb{P}(B \cap C)} = \mathbb{P}(A | B \cap C) \quad (2.4.11)$$

Correlation

Our next discussion concerns an important concept that deals with how two events are related, in a probabilistic sense.

Suppose that A and B are events with $\mathbb{P}(A) > 0$ and $\mathbb{P}(B) > 0$.

1. $\mathbb{P}(A | B) > \mathbb{P}(A)$ if and only if $\mathbb{P}(B | A) > \mathbb{P}(B)$ if and only if $\mathbb{P}(A \cap B) > \mathbb{P}(A)\mathbb{P}(B)$. In this case, A and B are *positively correlated*.
2. $\mathbb{P}(A | B) < \mathbb{P}(A)$ if and only if $\mathbb{P}(B | A) < \mathbb{P}(B)$ if and only if $\mathbb{P}(A \cap B) < \mathbb{P}(A)\mathbb{P}(B)$. In this case, A and B are *negatively correlated*.
3. $\mathbb{P}(A | B) = \mathbb{P}(A)$ if and only if $\mathbb{P}(B | A) = \mathbb{P}(B)$ if and only if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$. In this case, A and B are *uncorrelated or independent*.

Proof

These properties follow directly from the definition of conditional probability and simple algebra. Recall that multiplying or dividing an inequality by a positive number preserves the inequality.

Intuitively, if A and B are positively correlated, then the occurrence of either event means that the other event is more likely. If A and B are negatively correlated, then the occurrence of either event means that the other event is less likely. If A and B are uncorrelated, then the occurrence of either event does not change the probability of the other event. Independence is a fundamental concept that can be extended to more than two events and to random variables; these generalizations are studied in the next section on Independence. A much more general version of correlation, for random variables, is explored in the section on Covariance and Correlation in the chapter on Expected Value.

Suppose that A and B are events. Note from (4) that if $A \subseteq B$ or $B \subseteq A$ then A and B are positively correlated. If A and B are disjoint then A and B are negatively correlated.

Suppose that A and B are events in a random experiment.

1. A and B have the same correlation (positive, negative, or zero) as A^c and B^c .
2. A and B have the opposite correlation as A and B^c (that is, positive-negative, negative-positive, or 0-0).

Proof

1. Using DeMorgan's law and the complement law.

$$\mathbb{P}(A^c \cap B^c) - \mathbb{P}(A^c)\mathbb{P}(B^c) = \mathbb{P}[(A \cup B)^c] - \mathbb{P}(A^c)\mathbb{P}(B^c) = [1 - \mathbb{P}(A \cup B)] - [1 - \mathbb{P}(A)][1 - \mathbb{P}(B)] \quad (2.4.12)$$

Using the inclusion-exclusion law and algebra,

$$\mathbb{P}(A^c \cap B^c) - \mathbb{P}(A^c)\mathbb{P}(B^c) = \mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B) \quad (2.4.13)$$

2. Using the difference rule and the complement law:

$$\mathbb{P}(A \cap B^c) - \mathbb{P}(A)\mathbb{P}(B^c) = \mathbb{P}(A) - \mathbb{P}(A \cap B) - \mathbb{P}(A)[1 - \mathbb{P}(B)] = -[\mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)] \quad (2.4.14)$$

The Multiplication Rule

Sometimes conditional probabilities are known and can be used to find the probabilities of other events. Note first that if A and B are events with positive probability, then by the very definition of conditional probability,

$$\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B | A) = \mathbb{P}(B)\mathbb{P}(A | B) \quad (2.4.15)$$

The following generalization is known as the *multiplication rule* of probability. As usual, we assume that any event conditioned on has positive probability.

Suppose that (A_1, A_2, \dots, A_n) is a sequence of events. Then

$$\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_n) = \mathbb{P}(A_1)\mathbb{P}(A_2 | A_1)\mathbb{P}(A_3 | A_1 \cap A_2) \dots \mathbb{P}(A_n | A_1 \cap A_2 \cap \dots \cap A_{n-1}) \quad (2.4.16)$$

Proof

The product on the right is a collapsing product in which only the probability of the intersection of all n events survives. The product of the first two factors is $\mathbb{P}(A_1 \cap A_2)$, and hence the product of the first three factors is $\mathbb{P}(A_1 \cap A_2 \cap A_3)$, and so forth. The proof can be made more rigorous by induction on n .

The multiplication rule is particularly useful for experiments that consist of dependent stages, where A_i is an event in stage i . Compare the multiplication rule of probability with the multiplication rule of combinatorics.

As with any other result, the multiplication rule can be applied to a conditional probability measure. In the context above, if E is another event, then

$$\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_n | E) = \mathbb{P}(A_1 | E)\mathbb{P}(A_2 | A_1 \cap E)\mathbb{P}(A_3 | A_1 \cap A_2 \cap E) \dots \mathbb{P}(A_n | A_1 \cap A_2 \cap \dots \cap A_{n-1} \cap E) \quad (2.4.17)$$

Conditioning and Bayes' Theorem

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a countable collection of events that partition the sample space S , and that $\mathbb{P}(A_i) > 0$ for each $i \in I$.




Figure 2.4.2: A partition of S induces a partition of B .

The following theorem is known as the *law of total probability*.

If B is an event then

$$\mathbb{P}(B) = \sum_{i \in I} \mathbb{P}(A_i)\mathbb{P}(B | A_i) \quad (2.4.18)$$

Proof

Recall that $\{A_i \cap B : i \in I\}$ is a partition of B . Hence

$$\mathbb{P}(B) = \sum_{i \in I} \mathbb{P}(A_i \cap B) = \sum_{i \in I} \mathbb{P}(A_i)\mathbb{P}(B | A_i) \quad (2.4.19)$$

The following theorem is known as *Bayes' Theorem*, named after Thomas Bayes:

If B is an event then

$$\mathbb{P}(A_j | B) = \frac{\mathbb{P}(A_j)\mathbb{P}(B | A_j)}{\sum_{i \in I} \mathbb{P}(A_i)\mathbb{P}(B | A_i)}, \quad j \in I \quad (2.4.20)$$

Proof

Again the numerator is $\mathbb{P}(A_j \cap B)$ while the denominator is $\mathbb{P}(B)$ by the [law of total probability](#).

These two theorems are most useful, of course, when we know $\mathbb{P}(A_i)$ and $\mathbb{P}(B | A_i)$ for each $i \in I$. When we compute the probability of $\mathbb{P}(B)$ by the [law of total probability](#), we say that we are *conditioning* on the partition \mathcal{A} . Note that we can think of the sum as a *weighted average* of the conditional probabilities $\mathbb{P}(B | A_i)$ over $i \in I$, where $\mathbb{P}(A_i)$, $i \in I$ are the weight factors. In the context of [Bayes theorem](#), $\mathbb{P}(A_j)$ is the *prior probability* of A_j and $\mathbb{P}(A_j | B)$ is the *posterior probability* of A_j for $j \in I$. We will study more general versions of conditioning and Bayes theorem in the section on Discrete Distributions in the chapter on Distributions, and again in the section on Conditional Expected Value in the chapter on Expected Value.

Once again, the law of total probability and Bayes' theorem can be applied to a conditional probability measure. So, if E is another event with $\mathbb{P}(A_i \cap E) > 0$ for $i \in I$ then

$$\mathbb{P}(B | E) = \sum_{i \in I} \mathbb{P}(A_i | E) \mathbb{P}(B | A_i \cap E) \quad (2.4.21)$$

$$\mathbb{P}(A_j | B \cap E) = \frac{\mathbb{P}(A_j | E) \mathbb{P}(B | A_j \cap E)}{\sum_{i \in I} \mathbb{P}(A_i \cap E) \mathbb{P}(B | A_i \cap E)}, \quad j \in I \quad (2.4.22)$$

Examples and Applications

Basic Rules

Suppose that A and B are events in an experiment with $\mathbb{P}(A) = \frac{1}{3}$, $\mathbb{P}(B) = \frac{1}{4}$, $\mathbb{P}(A \cap B) = \frac{1}{10}$. Find each of the following:

1. $\mathbb{P}(A | B)$
2. $\mathbb{P}(B | A)$
3. $\mathbb{P}(A^c | B)$
4. $\mathbb{P}(B^c | A)$
5. $\mathbb{P}(A^c | B^c)$

Answer

1. $\frac{2}{5}$
2. $\frac{3}{10}$
3. $\frac{3}{5}$
4. $\frac{7}{10}$
5. $\frac{31}{45}$

Suppose that A , B , and C are events in a random experiment with $\mathbb{P}(A | C) = \frac{1}{2}$, $\mathbb{P}(B | C) = \frac{1}{3}$, and $\mathbb{P}(A \cap B | C) = \frac{1}{4}$. Find each of the following:

1. $\mathbb{P}(B \setminus A | C)$
2. $\mathbb{P}(A \cup B | C)$
3. $\mathbb{P}(A^c \cap B^c | C)$
4. $\mathbb{P}(A^c \cup B^c | C)$
5. $\mathbb{P}(A^c \cup B | C)$
6. $\mathbb{P}(A | B \cap C)$

Answer

1. $\frac{1}{12}$
2. $\frac{7}{12}$
3. $\frac{5}{12}$
4. $\frac{3}{4}$
5. $\frac{3}{4}$
6. $\frac{3}{4}$

Suppose that A and B are events in a random experiment with $\mathbb{P}(A) = \frac{1}{2}$, $\mathbb{P}(B) = \frac{1}{3}$, and $\mathbb{P}(A | B) = \frac{3}{4}$.

1. Find $\mathbb{P}(A \cap B)$
2. Find $\mathbb{P}(A \cup B)$

3. Find $\mathbb{P}(B \cup A^c)$
4. Find $\mathbb{P}(B | A)$
5. Are A and B positively correlated, negatively correlated, or independent?

Answer

1. $\frac{1}{4}$
2. $\frac{7}{12}$
3. $\frac{3}{4}$
4. $\frac{1}{2}$
5. positively correlated.

Open the conditional probability experiment.

1. Given $\mathbb{P}(A)$, $\mathbb{P}(B)$, and $\mathbb{P}(A \cap B)$, in the table, verify all of the other probabilities in the table.
2. Run the experiment 1000 times and compare the probabilities with the relative frequencies.

Simple Populations

In a certain population, 30% of the persons smoke cigarettes and 8% have COPD (*Chronic Obstructive Pulmonary Disease*). Moreover, 12% of the persons who smoke have COPD.

1. What percentage of the population smoke and have COPD?
2. What percentage of the population with COPD also smoke?
3. Are smoking and COPD positively correlated, negatively correlated, or independent?

Answer

1. 3.6%
2. 45%
3. positively correlated.

A company has 200 employees: 120 are women and 80 are men. Of the 120 female employees, 30 are classified as managers, while 20 of the 80 male employees are managers. Suppose that an employee is chosen at random.

1. Find the probability that the employee is female.
2. Find the probability that the employee is a manager.
3. Find the conditional probability that the employee is a manager given that the employee is female.
4. Find the conditional probability that the employee is female given that the employee is a manager.
5. Are the events *female* and *manager* positively correlated, negatively correlated, or independent?

Answer

1. $\frac{120}{200}$
2. $\frac{50}{200}$
3. $\frac{30}{120}$
4. $\frac{30}{50}$
5. independent

Dice and Coins

Consider the experiment that consists of rolling 2 standard, fair dice and recording the sequence of scores $\mathbf{X} = (X_1, X_2)$. Let Y denote the sum of the scores. For each of the following pairs of events, find the probability of each event and the conditional probability of each event given the other. Determine whether the events are positively correlated, negatively correlated, or independent.

1. $\{X_1 = 3\}, \{Y = 5\}$
2. $\{X_1 = 3\}, \{Y = 7\}$
3. $\{X_1 = 2\}, \{Y = 5\}$
4. $\{X_1 = 3\}, \{X_1 = 2\}$

Answer

In each case below, the answers are for $\mathbb{P}(A)$, $\mathbb{P}(B)$, $\mathbb{P}(A | B)$, and $\mathbb{P}(B | A)$

1. $\frac{1}{6}, \frac{1}{9}, \frac{1}{4}, \frac{1}{6}$. Positively correlated.
2. $\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}$. Independent.
3. $\frac{1}{6}, \frac{1}{9}, \frac{1}{4}, \frac{1}{6}$. Positively correlated.
4. $\frac{1}{6}, \frac{1}{6}, 0, 0$. Negatively correlated.

Note that positive correlation is not a transitive relation. From the previous exercise, for example, note that $\{X_1 = 3\}$ and $\{Y = 5\}$ are positively correlated, $\{Y = 5\}$ and $\{X_1 = 2\}$ are positively correlated, but $\{X_1 = 3\}$ and $\{X_1 = 2\}$ are negatively correlated (in fact, disjoint).

In dice experiment, set $n = 2$. Run the experiment 1000 times. Compute the empirical conditional probabilities corresponding to the conditional probabilities in the last exercise.

Consider again the experiment that consists of rolling 2 standard, fair dice and recording the sequence of scores $\mathbf{X} = (X_1, X_2)$. Let Y denote the sum of the scores, U the minimum score, and V the maximum score.

1. Find $\mathbb{P}(U = u | V = 4)$ for the appropriate values of u .
2. Find $\mathbb{P}(Y = y | V = 4)$ for the appropriate values of y .
3. Find $\mathbb{P}(V = v | Y = 8)$ for appropriate values of v .
4. Find $\mathbb{P}(U = u | Y = 8)$ for the appropriate values of u .
5. Find $\mathbb{P}[(X_1, X_2) = (x_1, x_2) | Y = 8]$ for the appropriate values of (x_1, x_2) .

Answer

1. $\frac{2}{7}$ for $u \in \{1, 2, 3\}$, $\frac{1}{7}$ for $u = 4$
2. $\frac{2}{7}$ for $y \in \{5, 6, 7\}$, $\frac{1}{7}$ for $y = 8$
3. $\frac{1}{5}$ for $v = 4$, $\frac{2}{5}$ for $v \in \{5, 6\}$
4. $\frac{2}{5}$ for $u \in \{2, 3\}$, $\frac{1}{5}$ for $u = 4$
5. $\frac{1}{5}$ for $(x_1, x_2) \in \{(2, 6), (6, 2), (3, 5), (5, 3), (4, 4)\}$

In the die-coin experiment, a standard, fair die is rolled and then a fair coin is tossed the number of times showing on the die. Let N denote the die score and H the event that all coin tosses result in heads.

1. Find $\mathbb{P}(H)$.
2. Find $\mathbb{P}(N = n | H)$ for $n \in \{1, 2, 3, 4, 5, 6\}$
3. Compare the results in (b) with $\mathbb{P}(N = n)$ for $n \in \{1, 2, 3, 4, 5, 6\}$. In each case, note whether the events H and $\{N = n\}$ are positively correlated, negatively correlated, or independent.

Answer

1. $\frac{21}{128}$
2. $\frac{64}{63} \frac{1}{2^n}$ for $n \in \{1, 2, 3, 4, 5, 6\}$
3. positively correlated for $n \in \{1, 2\}$ and negatively correlated for $n \in \{3, 4, 5, 6\}$

Run the die-coin experiment 1000 times. Let H and N be as defined in the previous exercise.

1. Compute the empirical probability of H . Compare with the true probability in the previous exercise.
2. Compute the empirical probability of $\{N = n\}$ given H , for $n \in \{1, 2, 3, 4, 5, 6\}$. Compare with the true probabilities in the previous exercise.

Suppose that a bag contains 12 coins: 5 are fair, 4 are biased with probability of heads $\frac{1}{3}$; and 3 are two-headed. A coin is chosen at random from the bag and tossed.

1. Find the probability that the coin is heads.
2. Given that the coin is heads, find the conditional probability of each coin type.

Answer

1. $\frac{41}{72}$
2. $\frac{15}{41}$ that the coin is fair, $\frac{8}{41}$ that the coin is biased, $\frac{18}{41}$ that the coin is two-headed

Compare [die-coin experiment](#) and [bag of coins experiment](#). In the die-coin experiment, we toss a coin with a *fixed* probability of heads a *random* number of times. In the bag of coins experiment, we effectively toss a coin with a *random* probability of heads a *fixed* number of times. The random experiment of tossing a coin with a fixed probability of heads p a fixed number of times n is known as the *binomial experiment* with parameters n and p . This is a very basic and important experiment that is studied in more detail in the section on the binomial distribution in the chapter on Bernoulli Trials. Thus, the die-coin and bag of coins experiments can be thought of as modifications of the binomial experiment in which a parameter has been *randomized*. In general, interesting new random experiments can often be constructed by randomizing one or more parameters in another random experiment.

In the coin-die experiment, a fair coin is tossed. If the coin lands tails, a fair die is rolled. If the coin lands heads, an ace-six flat die is tossed (faces 1 and 6 have probability $\frac{1}{4}$ each, while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each). Let H denote the event that the coin lands heads, and let Y denote the score when the chosen die is tossed.

1. Find $\mathbb{P}(Y = y)$ for $y \in \{1, 2, 3, 4, 5, 6\}$.
2. Find $\mathbb{P}(H | Y = y)$ for $y \in \{1, 2, 3, 4, 5, 6\}$.
3. Compare each probability in part (b) with $\mathbb{P}(H)$. In each case, note whether the events H and $\{Y = y\}$ are positively correlated, negatively correlated, or independent.

Answer

1. $\frac{5}{24}$ for $y \in \{1, 6\}$, $\frac{7}{48}$ for $y \in \{2, 3, 4, 5\}$
2. $\frac{3}{5}$ for $y \in \{1, 6\}$, $\frac{3}{7}$ for $y \in \{2, 3, 4, 5\}$
3. Positively correlated for $y \in \{1, 6\}$, negatively correlated for $y \in \{2, 3, 4, 5\}$

Run the coin-die experiment 1000 times. Let H and Y be as defined in the previous exercise.

1. Compute the empirical probability of $\{Y = y\}$, for each y , and compare with the true probability in the previous exercise
2. Compute the empirical probability of H given $\{Y = y\}$ for each y , and compare with the true probability in the previous exercise.

Cards

Consider the card experiment that consists of dealing 2 cards from a standard deck and recording the sequence of cards dealt. For $i \in \{1, 2\}$, let Q_i be the event that card i is a queen and H_i the event that card i is a heart. For each of the following pairs of events, compute the probability of each event, and the conditional probability of each event given the other. Determine whether the events are positively correlated, negatively correlated, or independent.

1. Q_1, H_1
2. Q_1, Q_2
3. Q_2, H_2
4. Q_1, H_2

Answer

The answers below are for $\mathbb{P}(A)$, $\mathbb{P}(B)$, $\mathbb{P}(A | B)$, and $\mathbb{P}(B | A)$ where A and B are the given events

1. $\frac{1}{13}, \frac{1}{4}, \frac{1}{13}, \frac{1}{4}$, independent.
2. $\frac{1}{13}, \frac{1}{13}, \frac{3}{51}, \frac{3}{51}$, negatively correlated.
3. $\frac{1}{13}, \frac{1}{4}, \frac{1}{13}, \frac{1}{4}$, independent.
4. $\frac{1}{13}, \frac{1}{4}, \frac{1}{13}, \frac{1}{4}$, independent.

In the card experiment, set $n = 2$. Run the experiment 500 times. Compute the conditional relative frequencies corresponding to the conditional probabilities in the last exercise.

Consider the card experiment that consists of dealing 3 cards from a standard deck and recording the sequence of cards dealt. Find the probability of the following events:

1. All three cards are all hearts.
2. The first two cards are hearts and the third is a spade.
3. The first and third cards are hearts and the second is a spade.

Proof

1. $\frac{11}{850}$
2. $\frac{13}{850}$
3. $\frac{13}{850}$

In the card experiment, set $n = 3$ and run the simulation 1000 times. Compute the empirical probability of each event in the previous exercise and compare with the true probability.

Bivariate Uniform Distributions

Recall that Buffon's coin experiment consists of tossing a coin with radius $r \leq \frac{1}{2}$ randomly on a floor covered with square tiles of side length 1. The coordinates (X, Y) of the center of the coin are recorded relative to axes through the center of the square, parallel to the sides. Since the needle is dropped randomly, the basic modeling assumption is that (X, Y) is uniformly distributed on the square $[-1/2, 1/2]^2$.




Figure 2.4.3: Buffon's coin experiment

In Buffon's coin experiment,

1. Find $\mathbb{P}(Y > 0 \mid X < Y)$
2. Find the conditional distribution of (X, Y) given that the coin does not touch the sides of the square.

Answer

1. $\frac{3}{4}$
2. Given $(X, Y) \in [r - \frac{1}{2}, \frac{1}{2} - r]^2$, (X, Y) is uniformly distributed on this set.

Run Buffon's coin experiment 500 times. Compute the empirical probability that $Y > 0$ given that $X < Y$ and compare with the probability in the last exercise.

In the conditional probability experiment, the random points are uniformly distributed on the rectangle S . Move and resize events A and B and note how the probabilities change. For each of the following configurations, run the experiment 1000 times and compare the relative frequencies with the true probabilities.

1. A and B in general position
2. A and B disjoint
3. $A \subseteq B$
4. $B \subseteq A$

Reliability

A plant has 3 assembly lines that produces memory chips. Line 1 produces 50% of the chips and has a defective rate of 4%; line 2 has produces 30% of the chips and has a defective rate of 5%; line 3 produces 20% of the chips and has a defective rate of 1%. A chip is chosen at random from the plant.

1. Find the probability that the chip is defective.
2. Given that the chip is defective, find the conditional probability for each line.

Answer

1. 0.037
2. 0.541 for line 1, 0.405 for line 2, 0.054 for line 3

Suppose that a bit (0 or 1) is sent through a noisy communications channel. Because of the noise, the bit sent may be received incorrectly as the complementary bit. Specifically, suppose that if 0 is sent, then the probability that 0 is received is 0.9 and the probability that 1 is received is 0.1. If 1 is sent, then the probability that 1 is received is 0.8 and the probability that 0 is received is 0.2. Finally, suppose that 1 is sent with probability 0.6 and 0 is sent with probability 0.4. Find the probability that

1. 1 was sent given that 1 was received
2. 0 was sent given that 0 was received

Answer

1. 12/13
2. 3/4

Suppose that T denotes the lifetime of a light bulb (in 1000 hour units), and that T has the following exponential distribution, defined for measurable $A \subseteq [0, \infty)$:

$$\mathbb{P}(T \in A) = \int_A e^{-t} dt \quad (2.4.23)$$

1. Find $\mathbb{P}(T > 3)$
2. Find $\mathbb{P}(T > 5 \mid T > 2)$

Answer

1. e^{-3}
2. e^{-3}

Suppose again that T denotes the lifetime of a light bulb (in 1000 hour units), but that T is uniformly distributed on the interval $[0, 10]$.

1. Find $\mathbb{P}(T > 3)$
2. Find $\mathbb{P}(T > 5 \mid T > 2)$

Answer

1. $\frac{7}{10}$
2. $\frac{5}{8}$

Genetics

Please refer to the discussion of genetics in the section on random experiments if you need to review some of the definitions in this section.

Recall first that the *ABO blood type* in humans is determined by three alleles: a , b , and o . Furthermore, a and b are co-dominant and o is recessive. Suppose that the probability distribution for the set of blood genotypes in a certain population is given in the following table:

Genotype	aa	ab	ao	bb	bo	oo
Probability	0.050	0.038	0.310	0.007	0.116	0.479

Suppose that a person is chosen at random from the population. Let A , B , AB , and O be the events that the person is type A , type B , type AB , and type O respectively. Let H be the event that the person is homozygous, and let D denote the event that the person has an o allele. Find each of the following:

1. $\mathbb{P}(A)$, $\mathbb{P}(B)$, $\mathbb{P}(AB)$, $\mathbb{P}(O)$, $\mathbb{P}(H)$, $\mathbb{P}(D)$

2. $P(A \cap H)$, $P(A | H)$, $P(H | A)$. Are the events A and H positively correlated, negatively correlated, or independent?
3. $P(B \cap H)$, $P(B | H)$, $P(H | B)$. Are the events B and H positively correlated, negatively correlated, or independent?
4. $P(A \cap D)$, $P(A | D)$, $P(D | A)$. Are the events A and D positively correlated, negatively correlated, or independent?
5. $P(B \cap D)$, $P(B | D)$, $P(D | B)$. Are the events B and D positively correlated, negatively correlated, or independent?
6. $P(H \cap D)$, $P(H | D)$, $P(D | H)$. Are the events H and D positively correlated, negatively correlated, or independent?

Answer

1. 0.360, 0.123, 0.038, 0.479, 0.536, 0.905
2. 0.050, 0.093, 0.139. A and H are negatively correlated.
3. 0.007, 0.013, 0.057. B and H are negatively correlated.
4. 0.310, 0.343, 0.861. A and D are negatively correlated.
5. 0.116, 0.128, 0.943. B and D are positively correlated.
6. 0.479, 0.529, 0.894. H and D are negatively correlated.

Suppose next that pod color in certain type of pea plant is determined by a gene with two alleles: g for green and y for yellow, and that g is dominant and y recessive.

Suppose that a green-pod plant and a yellow-pod plant are bred together. Suppose further that the green-pod plant has a $\frac{1}{4}$ chance of carrying the recessive yellow-pod allele.

1. Find the probability that a child plant will have green pods.
2. Given that a child plant has green pods, find the updated probability that the green-pod parent has the recessive allele.

Answer

1. $\frac{7}{8}$
2. $\frac{1}{7}$

Suppose that two green-pod plants are bred together. Suppose further that with probability $\frac{1}{3}$ neither plant has the recessive allele, with probability $\frac{1}{2}$ one plant has the recessive allele, and with probability $\frac{1}{6}$ both plants have the recessive allele.

1. Find the probability that a child plant has green pods.
2. Given that a child plant has green pods, find the updated probability that both parents have the recessive gene.

Answer

1. $\frac{23}{24}$
2. $\frac{3}{23}$

Next consider a sex-linked hereditary disorder in humans (such as colorblindness or hemophilia). Let h denote the healthy allele and d the defective allele for the gene linked to the disorder. Recall that h is dominant and d recessive for women.

Suppose that in a certain population, 50% are male and 50% are female. Moreover, suppose that 10% of males are color blind but only 1% of females are color blind.

1. Find the percentage of color blind persons in the population.
2. Find the percentage of color blind persons that are male.

Answer

1. 5.5%
2. 90.9%

Since color blindness is a sex-linked hereditary disorder, note that it's reasonable in the previous exercise that the probability that a female is color blind is the square of the probability that a male is color blind. If p is the probability of the defective allele on the X chromosome, then p is also the probability that a male will be color blind. But since the defective allele is recessive, a woman would need two copies of the defective allele to be color blind, and assuming independence, the probability of this event is p^2 .

A man and a woman do not have a certain sex-linked hereditary disorder, but the woman has a $\frac{1}{3}$ chance of being a carrier.

1. Find the probability that a son born to the couple will be normal.

2. Find the probability that a daughter born to the couple will be a carrier.
3. Given that a son born to the couple is normal, find the updated probability that the mother is a carrier.

Answer

1. $\frac{5}{6}$
2. $\frac{1}{6}$
3. $\frac{1}{5}$

Urn Models

Urn 1 contains 4 red and 6 green balls while urn 2 contains 7 red and 3 green balls. An urn is chosen at random and then a ball is chosen at random from the selected urn.

1. Find the probability that the ball is green.
2. Given that the ball is green, find the conditional probability that urn 1 was selected.

Answer

1. $\frac{9}{20}$
2. $\frac{2}{3}$

Urn 1 contains 4 red and 6 green balls while urn 2 contains 6 red and 3 green balls. A ball is selected at random from urn 1 and transferred to urn 2. Then a ball is selected at random from urn 2.

1. Find the probability that the ball from urn 2 is green.
2. Given that the ball from urn 2 is green, find the conditional probability that the ball from urn 1 was green.

Answer

1. $\frac{9}{25}$
2. $\frac{2}{3}$

An urn initially contains 6 red and 4 green balls. A ball is chosen at random from the urn and its color is recorded. It is then replaced in the urn and 2 new balls of the same color are added to the urn. The process is repeated. Find the probability of each of the following events:

1. Balls 1 and 2 are red and ball 3 is green.
2. Balls 1 and 3 are red and ball 2 is green.
3. Ball 1 is green and balls 2 and 3 are red.
4. Ball 2 is red.
5. Ball 1 is red given that ball 2 is red.

Answer

1. $\frac{4}{35}$
2. $\frac{4}{35}$
3. $\frac{4}{35}$
4. $\frac{3}{5}$
5. $\frac{2}{3}$

Think about the results in the previous exercise. Note in particular that the answers to parts (a), (b), and (c) are the same, and that the probability that the second ball is red in part (d) is the same as the probability that the first ball is red. More generally, the probabilities of events do not depend on the order of the draws. For example, the probability of an event involving the first, second, and third draws is the same as the probability of the corresponding event involving the seventh, tenth and fifth draws. Technically, the sequence of events (R_1, R_2, \dots) is exchangeable. The random process described in this exercise is a special case of *Pólya's urn scheme*, named after George Pólya. We will study Pólya's urn in more detail in the chapter on Finite Sampling Models

An urn initially contains 6 red and 4 green balls. A ball is chosen at random from the urn and its color is recorded. It is then replaced in the urn and two new balls of the *other* color are added to the urn. The process is repeated. Find the probability of

each of the following events:

1. Balls 1 and 2 are red and ball 3 is green.
2. Balls 1 and 3 are red and ball 2 is green.
3. Ball 1 is green and balls 2 and 3 are red.
4. Ball 2 is red.
5. Ball 1 is red given that ball 2 is red.

Answer

1. $\frac{6}{35}$
2. $\frac{6}{35}$
3. $\frac{16}{105}$
4. $\frac{17}{30}$
5. $\frac{9}{17}$

Think about the results in the previous exercise, and compare with [Pólya's urn](#). Note that the answers to parts (a), (b), and (c) are not all the same, and that the probability that the second ball is red in part (d) is not the same as the probability that the first ball is red. In short, the sequence of events (R_1, R_2, \dots) is *not* exchangeable.

Diagnostic Testing

Suppose that we have a random experiment with an event A of interest. When we run the experiment, of course, event A will either occur or not occur. However, suppose that we are not able to observe the occurrence or non-occurrence of A directly. Instead we have a *diagnostic test* designed to indicate the occurrence of event A ; thus the test that can be either *positive* for A or *negative* for A . The test also has an element of randomness, and in particular can be in error. Here are some typical examples of the type of situation we have in mind:

- The event is that a person has a certain disease and the test is a blood test for the disease.
- The event is that a woman is pregnant and the test is a home pregnancy test.
- The event is that a person is lying and the test is a lie-detector test.
- The event is that a device is defective and the test consists of a sensor reading.
- The event is that a missile is in a certain region of airspace and the test consists of radar signals.
- The event is that a person has committed a crime, and the test is a jury trial with evidence presented for and against the event.

Let T be the event that the test is positive for the occurrence of A . The conditional probability $\mathbb{P}(T | A)$ is called the *sensitivity* of the test. The complementary probability

$$\mathbb{P}(T^c | A) = 1 - \mathbb{P}(T | A) \quad (2.4.24)$$

is the *false negative* probability. The conditional probability $\mathbb{P}(T^c | A^c)$ is called the *specificity* of the test. The complementary probability

$$\mathbb{P}(T | A^c) = 1 - \mathbb{P}(T^c | A^c) \quad (2.4.25)$$

is the *false positive* probability. In many cases, the sensitivity and specificity of the test are known, as a result of the development of the test. However, the *user* of the test is interested in the opposite conditional probabilities, namely $\mathbb{P}(A | T)$, the probability of the event of interest, given a positive test, and $\mathbb{P}(A^c | T^c)$, the probability of the complementary event, given a negative test. Of course, if we know $\mathbb{P}(A | T)$ then we also have $\mathbb{P}(A^c | T) = 1 - \mathbb{P}(A | T)$, the probability of the complementary event given a positive test. Similarly, if we know $\mathbb{P}(A^c | T^c)$ then we also have $\mathbb{P}(A | T^c)$, the probability of the event given a negative test. Computing the probabilities of interest is simply a special case of Bayes' theorem.

The probability that the event occurs, given a positive test is

$$\mathbb{P}(A | T) = \frac{\mathbb{P}(A)\mathbb{P}(T | A)}{\mathbb{P}(A)\mathbb{P}(T | A) + \mathbb{P}(A^c)\mathbb{P}(T | A^c)} \quad (2.4.26)$$

The probability that the event does not occur, given a negative test is

$$\mathbb{P}(A^c | T^c) = \frac{\mathbb{P}(A^c)\mathbb{P}(T^c | A^c)}{\mathbb{P}(A)\mathbb{P}(T^c | A) + \mathbb{P}(A^c)\mathbb{P}(T^c | A^c)} \quad (2.4.27)$$

There is often a trade-off between sensitivity and specificity. An attempt to make a test more sensitive may result in the test being less specific, and an attempt to make a test more specific may result in the test being less sensitive. As an extreme example, consider the worthless test that always returns positive, no matter what the evidence. Then $T = S$ so the test has sensitivity 1, but specificity 0. At the opposite extreme is the worthless test that always returns negative, no matter what the evidence. Then $T = \emptyset$ so the test has specificity 1 but sensitivity 0. In between these extremes are helpful tests that are actually based on evidence of some sort.

Suppose that the sensitivity $a = \mathbb{P}(T | A) \in (0, 1)$ and the specificity $b = \mathbb{P}(T^c | A^c) \in (0, 1)$ are fixed. Let $p = \mathbb{P}(A)$ denote the prior probability of the event A and $P = \mathbb{P}(A | T)$ the posterior probability of A given a positive test.

P as a function of p is given by

$$P = \frac{ap}{(a+b-1)p + (1-b)}, \quad p \in [0, 1] \quad (2.4.28)$$

1. P increases continuously from 0 to 1 as p increases from 0 to 1.
2. P is concave downward if $a + b > 1$. In this case A and T are positively correlated.
3. P is concave upward if $a + b < 1$. In this case A and T are negatively correlated.
4. $P = p$ if $a + b = 1$. In this case, A and T are uncorrelated (independent).

Proof

The formula for P in terms of p follows from (42) and algebra. For part (a), note that

$$\frac{dP}{dp} = \frac{a(1-b)}{[(a+b-1)p + (1-b)]^2} > 0 \quad (2.4.29)$$

For parts (b)-(d), note that

$$\frac{d^2P}{dp^2} = \frac{-2a(1-b)(a+b-1)}{[(1+b-1)p + (1-b)]^3} \quad (2.4.30)$$

If $a + b > 1$, $d^2P/dp^2 < 0$ so P is concave downward on $[0, 1]$ and hence $P > p$ for $0 < p < 1$. If $a + b < 1$, $d^2P/dp^2 > 0$ so P is concave upward on $[0, 1]$ and hence $P < p$ for $0 < p < 1$. Trivially if $a + b = 1$, $P = p$ for $0 \leq p \leq 1$.

Of course, part (b) is the typical case, where the test is useful. In fact, we would hope that the sensitivity and specificity are close to 1. In case (c), the test is worse than useless since it gives the wrong information about A . But this case could be turned into a useful test by simply reversing the roles of *positive* and *negative*. In case (d), the test is worthless and gives no information about A . It's interesting that the broad classification above depends only on the *sum* of the sensitivity and specificity.




Figure 2.4.4: $P = \mathbb{P}(A | T)$ as a function of $p = \mathbb{P}(A)$ in the three cases

Suppose that a diagnostic test has sensitivity 0.99 and specificity 0.95. Find $\mathbb{P}(A | T)$ for each of the following values of $\mathbb{P}(A)$:

1. 0.001
2. 0.01
3. 0.2
4. 0.5
5. 0.7
6. 0.9

Answer

1. 0.0194
2. 0.1667
3. 0.8319
4. 0.9519
5. 0.9788
6. 0.9944

With sensitivity 0.99 and specificity 0.95, the test in the last exercise superficially looks good. However the small value of $\mathbb{P}(A | T)$ for small values of $\mathbb{P}(A)$ is striking (but inevitable given the [properties above](#)). The moral, of course, is that $\mathbb{P}(A | T)$ depends critically on $\mathbb{P}(A)$ not just on the sensitivity and specificity of the test. Moreover, the correct comparison is $\mathbb{P}(A | T)$ with $\mathbb{P}(A)$, as in the exercise, not $\mathbb{P}(A | T)$ with $\mathbb{P}(T | A)$ —Beware of the fallacy of the transposed conditional! In terms of the correct comparison, the test does indeed work well; $\mathbb{P}(A | T)$ is significantly larger than $\mathbb{P}(A)$ in all cases.

A woman initially believes that there is an even chance that she is or is not pregnant. She takes a home pregnancy test with sensitivity 0.95 and specificity 0.90 (which are reasonable values for a home pregnancy test). Find the updated probability that the woman is pregnant in each of the following cases.

1. The test is positive.
2. The test is negative.

Answer

1. 0.905
2. 0.053

Suppose that 70% of defendants brought to trial for a certain type of crime are guilty. Moreover, historical data show that juries convict guilty persons 80% of the time and convict innocent persons 10% of the time. Suppose that a person is tried for a crime of this type. Find the updated probability that the person is guilty in each of the following cases:

1. The person is convicted.
2. The person is acquitted.

Answer

1. 0.949
2. 0.341

The “Check Engine” light on your car has turned on. Without the information from the light, you believe that there is a 10% chance that your car has a serious engine problem. You learn that if the car has such a problem, the light will come on with probability 0.99, but if the car does not have a serious problem, the light will still come on, under circumstances similar to yours, with probability 0.3. Find the updated probability that you have an engine problem.

Answer

0.268

The standard test for HIV is the *ELISA* (*Enzyme-Linked Immunosorbent Assay*) test. It has sensitivity and specificity of 0.999. Suppose that a person is selected at random from a population in which 1% are infected with HIV, and given the ELISA test. Find the probability that the person has HIV in each of the following cases:

1. The test is positive.
2. The test is negative.

Answer

1. 0.9098
2. 0.00001

The ELISA test for HIV is a very good one. Let's look another test, this one for prostate cancer, that's rather bad.

The PSA test for prostate cancer is based on a blood marker known as the *Prostate Specific Antigen*. An elevated level of PSA is evidence for prostate cancer. To have a diagnostic test, in the sense that we are discussing here, we must decide on a definite level of PSA, above which we declare the test to be positive. A positive test would typically lead to other more invasive tests (such as biopsy) which, of course, carry risks and cost. The PSA test with cutoff 2.6 ng/ml has sensitivity 0.40 and specificity 0.81. The overall incidence of prostate cancer among males is 156 per 100000. Suppose that a man, with no particular risk factors, has the PSA test. Find the probability that the man has prostate cancer in each of the following cases:

1. The test is positive.
2. The test is negative.

Answer

1. 0.00328
2. 0.00116

Diagnostic testing is closely related to a general statistical procedure known as *hypothesis testing*. A separate chapter on hypothesis testing explores this procedure in detail.

Data Analysis Exercises

For the M&M data set, find the empirical probability that a bag has at least 10 reds, given that the weight of the bag is at least 48 grams.

Answer

$$\frac{10}{23}.$$

Consider the Cicada data.

1. Find the empirical probability that a cicada weighs at least 0.25 grams given that the cicada is male.
2. Find the empirical probability that a cicada weighs at least 0.25 grams given that the cicada is the tredecula species.

Answer

1. $\frac{2}{45}$
2. $\frac{7}{44}$

This page titled [2.4: Conditional Probability](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.5: Independence

In this section, we will discuss independence, one of the fundamental concepts in probability theory. Independence is frequently invoked as a modeling assumption, and moreover, (classical) probability itself is based on the idea of independent replications of the experiment. As usual, if you are a new student of probability, you may want to skip the technical details.

Basic Theory

As usual, our starting point is a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$ so that S is the set of outcomes, \mathcal{S} the collection of events, and \mathbb{P} the probability measure on the sample space (S, \mathcal{S}) . We will define independence for two events, then for collections of events, and then for collections of random variables. In each case, the basic idea is the same.

Independence of Two Events

Two events A and B are independent if

$$\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B) \quad (2.5.1)$$

If both of the events have positive probability, then independence is equivalent to the statement that the conditional probability of one event given the other is the same as the unconditional probability of the event:

$$\mathbb{P}(A | B) = \mathbb{P}(A) \iff \mathbb{P}(B | A) = \mathbb{P}(B) \iff \mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B) \quad (2.5.2)$$

This is how you should think of independence: knowledge that one event has occurred does not change the probability assigned to the other event. Independence of two events was discussed in the last section in the context of correlation. In particular, for two events, *independent* and *uncorrelated* mean the same thing.

The terms *independent* and *disjoint* sound vaguely similar but they are actually very different. First, note that disjointness is purely a set-theory concept while independence is a probability (measure-theoretic) concept. Indeed, two events can be independent relative to one probability measure and dependent relative to another. But most importantly, two disjoint events can *never* be independent, except in the trivial case that one of the events is null.

Suppose that A and B are disjoint events, each with positive probability. Then A and B are dependent, and in fact are negatively correlated.

Proof

Note that $\mathbb{P}(A \cap B) = \mathbb{P}(\emptyset) = 0$ but $\mathbb{P}(A)\mathbb{P}(B) > 0$.

If A and B are independent events then intuitively it seems clear that any event that can be constructed from A should be independent of any event that can be constructed from B . This is the case, as the next result shows. Moreover, this basic idea is essential for the generalization of independence that we will consider shortly.

If A and B are independent events, then each of the following pairs of events is independent:

1. A^c, B
2. B, A^c
3. A^c, B^c

Proof

Suppose that A and B are independent. Then by the difference rule and the complement rule,

$$\mathbb{P}(A^c \cap B) = \mathbb{P}(B) - \mathbb{P}(A \cap B) = \mathbb{P}(B) - \mathbb{P}(A)\mathbb{P}(B) = \mathbb{P}(B)[1 - \mathbb{P}(A)] = \mathbb{P}(B)\mathbb{P}(A^c) \quad (2.5.3)$$

Hence A^c and B are equivalent. Parts (b) and (c) follow from (a).

An event that is “essentially deterministic”, that is, has probability 0 or 1, is independent of any other event, even itself.

Suppose that A and B are events.

1. If $\mathbb{P}(A) = 0$ or $\mathbb{P}(A) = 1$, then A and B are independent.
2. A is independent of itself if and only if $\mathbb{P}(A) = 0$ or $\mathbb{P}(A) = 1$.

Proof

1. Recall that if $\mathbb{P}(A) = 0$ then $\mathbb{P}(A \cap B) = 0$, and if $\mathbb{P}(A) = 1$ then $\mathbb{P}(A \cap B) = \mathbb{P}(B)$. In either case we have $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$.
2. The independence of A with itself gives $\mathbb{P}(A) = [\mathbb{P}(A)]^2$ and hence either $\mathbb{P}(A) = 0$ or $\mathbb{P}(A) = 1$.

General Independence of Events

To extend the definition of independence to more than two events, we might think that we could just require *pairwise independence*, the independence of each pair of events. However, this is not sufficient for the *strong* type of independence that we have in mind. For example, suppose that we have three events A , B , and C . *Mutual* independence of these events should not only mean that each pair is independent, but also that an event that can be constructed from A and B (for example $A \cup B^c$) should be independent of C . Pairwise independence does not achieve this; an [exercise](#) below gives three events that are pairwise independent, but the intersection of two of the events is related to the third event in the strongest possible sense.

Another possible generalization would be to simply require the probability of the intersection of the events to be the product of the probabilities of the events. However, this condition does not even guarantee pairwise independence. An [exercise](#) below gives an example. However, the definition of independence for two events does generalize in a natural way to an arbitrary collection of events.

Suppose that A_i is an event for each i in an index set I . Then the collection $\mathcal{A} = \{A_i : i \in I\}$ is *independent* if for every finite $J \subseteq I$,

$$\mathbb{P}\left(\bigcap_{j \in J} A_j\right) = \prod_{j \in J} \mathbb{P}(A_j) \quad (2.5.4)$$

Independence of a collection of events is much stronger than mere pairwise independence of the events in the collection. The basic *inheritance property* in the following result follows immediately from the definition.

Suppose that \mathcal{A} is a collection of events.

1. If \mathcal{A} is independent, then \mathcal{B} is independent for every $\mathcal{B} \subseteq \mathcal{A}$.
2. If \mathcal{B} is independent for every finite $\mathcal{B} \subseteq \mathcal{A}$ then \mathcal{A} is independent.

For a finite collection of events, the number of conditions required for mutual independence grows exponentially with the number of events.

There are $2^n - n - 1$ non-trivial conditions in the definition of the independence of n events.

1. Explicitly give the 4 conditions that must be satisfied for events A , B , and C to be independent.
2. Explicitly give the 11 conditions that must be satisfied for events A , B , C , and D to be independent.

Answer

There are 2^n subcollections of the n events. One is empty and n involve a single event. The remaining $2^n - n - 1$ subcollections involve two or more events and correspond to non-trivial conditions.

1. A , B , C are independent if and only if

$$\begin{aligned} \mathbb{P}(A \cap B) &= \mathbb{P}(A)\mathbb{P}(B) \\ \mathbb{P}(A \cap C) &= \mathbb{P}(A)\mathbb{P}(C) \\ \mathbb{P}(B \cap C) &= \mathbb{P}(B)\mathbb{P}(C) \\ \mathbb{P}(A \cap B \cap C) &= \mathbb{P}(A)\mathbb{P}(B)\mathbb{P}(C) \end{aligned}$$

2. A , B , C , D are independent if and only if

$$\begin{aligned}
 \mathbb{P}(A \cap B) &= \mathbb{P}(A)\mathbb{P}(B) \\
 \mathbb{P}(A \cap C) &= \mathbb{P}(A)\mathbb{P}(C) \\
 \mathbb{P}(A \cap D) &= \mathbb{P}(A)\mathbb{P}(D) \\
 \mathbb{P}(B \cap C) &= \mathbb{P}(B)\mathbb{P}(C) \\
 \mathbb{P}(B \cap D) &= \mathbb{P}(B)\mathbb{P}(D) \\
 \mathbb{P}(C \cap D) &= \mathbb{P}(C)\mathbb{P}(D) \\
 \mathbb{P}(A \cap B \cap C) &= \mathbb{P}(A)\mathbb{P}(B)\mathbb{P}(C) \\
 \mathbb{P}(A \cap B \cap D) &= \mathbb{P}(A)\mathbb{P}(B)\mathbb{P}(D) \\
 \mathbb{P}(A \cap C \cap D) &= \mathbb{P}(A)\mathbb{P}(C)\mathbb{P}(D) \\
 \mathbb{P}(B \cap C \cap D) &= \mathbb{P}(B)\mathbb{P}(C)\mathbb{P}(D) \\
 \mathbb{P}(A \cap B \cap C \cap D) &= \mathbb{P}(A)\mathbb{P}(B)\mathbb{P}(C)\mathbb{P}(D)
 \end{aligned}$$

If the events A_1, A_2, \dots, A_n are independent, then it follows immediately from the definition that

$$\mathbb{P}\left(\bigcap_{i=1}^n A_i\right) = \prod_{i=1}^n \mathbb{P}(A_i) \quad (2.5.5)$$

This is known as the *multiplication rule* for independent events. Compare this with the general multiplication rule for conditional probability.

The collection of *essentially deterministic* events $\mathcal{D} = \{A \in \mathcal{S} : \mathbb{P}(A) = 0 \text{ or } \mathbb{P}(A) = 1\}$ is independent.

Proof

Suppose that $\{A_1, A_2, \dots, A_n\} \subseteq \mathcal{D}$. If $\mathbb{P}(A_i) = 0$ for some $i \in \{1, 2, \dots, n\}$ then $\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_n) = 0$. If $\mathbb{P}(A_i) = 1$ for every $i \in \{1, 2, \dots, n\}$ then $\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_n) = 1$. In either case, $\mathbb{P}(A_1 \cap A_2 \cap \dots \cap A_n) = \mathbb{P}(A_1)\mathbb{P}(A_2) \cdots \mathbb{P}(A_n)$.

The next result generalizes the [theorem above](#) on the complements of two independent events.

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ and $\mathcal{B} = \{B_i : i \in I\}$ are two collections of events with the property that for each $i \in I$, either $B_i = A_i$ or $B_i = A_i^c$. Then \mathcal{A} is independent if and only if \mathcal{B} is an independent.

Proof

The proof is actually very similar to the proof for [two events](#), except for more complicated notation. First, by the symmetry of the relation between \mathcal{A} and \mathcal{B} , it suffices to show \mathcal{A} independent implies \mathcal{B} independent. Next, by the [inheritance property](#), it suffices to consider the case where the index set I is finite.

1. Fix $k \in I$ and define $B_k = A_k^c$ and $B_i = A_i$ for $i \in I \setminus \{k\}$. Suppose now that $J \subseteq I$. If $k \notin J$ then trivially,

$$\mathbb{P}\left(\bigcap_{j \in J} B_j\right) = \prod_{j \in J} \mathbb{P}(B_j).$$

$$\begin{aligned}
 \mathbb{P}\left(\bigcap_{j \in J} B_j\right) &= \mathbb{P}\left(\bigcap_{j \in J \setminus \{k\}} A_j\right) - \mathbb{P}\left(\bigcap_{j \in J} A_j\right) \\
 &= \prod_{j \in J \setminus \{k\}} \mathbb{P}(A_j) - \prod_{j \in J} \mathbb{P}(A_j) = \left[\prod_{j \in J \setminus \{k\}} \mathbb{P}(A_j) \right] [1 - \mathbb{P}(A_k)] = \prod_{j \in J} \mathbb{P}(B_j)
 \end{aligned}$$

Hence $\{B_i : i \in I\}$ is a collection of independent events.

2. Suppose now that $\mathcal{B} = \{B_i : i \in I\}$ is a general collection of events where $B_i = A_i$ or $B_i = A_i^c$ for each $i \in I$. Then \mathcal{B} can be obtained from \mathcal{A} by a finite sequence of complement changes of the type in (a), each of which preserves independence.

The last theorem in turn leads to the type of strong independence that we want. The following exercise gives examples.

If A, B, C , and D are independent events, then

1. $A \cup B, C^c, D$ are independent.
2. $A \cup B^c, C^c \cup D^c$ are independent.

Proof

We will give proofs that use the [complement theorem](#), but to do so, some additional notation is helpful. If E is an event, let $E^1 = E$ and $E^0 = E^c$.

1. Note that $A \cup B = \bigcup_{(i,j) \in I} A^i \cap B^j$ where $I = \{(1, 0), (0, 1), (1, 1)\}$ and note that the events in the union are disjoint. By the distributive property, $(A \cup B) \cap C^c = \bigcup_{(i,j) \in I} A^i \cap B^j \cap C^0$ and again the events in the union are disjoint. By additivity and [complement theorem](#),

$$\mathbb{P}[(A \cup B) \cap C^c] = \sum_{(i,j) \in I} \mathbb{P}(A^i) \mathbb{P}(B^j) \mathbb{P}(C^0) = \left(\sum_{(i,j) \in I} \mathbb{P}(A^i) \mathbb{P}(B^j) \right) \mathbb{P}(C^0) = \mathbb{P}(A \cup B) \mathbb{P}(C^c) \quad (2.5.6)$$

By exactly the same type of argument, $\mathbb{P}[(A \cup B) \cap D] = \mathbb{P}(A \cup B) \mathbb{P}(D)$ and $\mathbb{P}[(A \cup B) \cap C^c \cap D] = \mathbb{P}(A \cup B) \mathbb{P}(C^c) \mathbb{P}(D)$. Directly from the result above on [complements](#), $\mathbb{P}(C^c \cap D) = \mathbb{P}(C^c) \mathbb{P}(D)$.

2. Note that $A \cup B^c = \bigcup_{(i,j) \in I} A^i \cap B^j$ where $I = \{(0, 0), (1, 0), (1, 1)\}$ and note that the events in the union are disjoint. Similarly $C^c \cup D^c = \bigcup_{(k,l) \in J} C^k \cap D^l$ where $J = \{(0, 0), (1, 0), (0, 1)\}$ and again the events in the union are disjoint. By the distributive rule for set operations,

$$(A \cup B^c) \cap (C^c \cup D^c) = \bigcup_{(i,j,k,l) \in I \times J} A^i \cap B^j \cap C^k \cap D^l \quad (2.5.7)$$

and once again, the events in the union are disjoint. By additivity and the [complement theorem](#),

$$\mathbb{P}[(A \cup B^c) \cap (C^c \cup D^c)] = \sum_{(i,j,k,l) \in I \times J} \mathbb{P}(A^i) \mathbb{P}(B^j) \mathbb{P}(C^k) \mathbb{P}(D^l) \quad (2.5.8)$$

But also by additivity, the [complement theorem](#), and the distributive property of arithmetic,

$$\mathbb{P}(A \cup B^c) \mathbb{P}(C^c \cup D^c) = \left(\sum_{(i,j) \in I} \mathbb{P}(A^i) \mathbb{P}(B^j) \right) \left(\sum_{(k,l) \in J} \mathbb{P}(C^k) \mathbb{P}(D^l) \right) = \sum_{(i,j,k,l) \in I \times J} \mathbb{P}(A^i) \mathbb{P}(B^j) \mathbb{P}(C^k) \mathbb{P}(D^l) \quad (2.5.9)$$

The complete generalization of these results is a bit complicated, but roughly means that if we start with a collection of independent events, and form new events from disjoint subcollections (using the set operations of union, intersection, and complement), then the new events are independent. For a precise statement, see the section on measure spaces. The importance of the [complement theorem](#) lies in the fact that any event that can be defined in terms of a finite collection of events $\{A_i : i \in I\}$ can be written as a disjoint union of events of the form $\bigcap_{i \in I} B_i$ where $B_i = A_i$ or $B_i = A_i^c$ for each $i \in I$.

Another consequence of the general [complement theorem](#) is a formula for the probability of the union of a collection of independent events that is much nicer than the inclusion-exclusion formula.

If A_1, A_2, \dots, A_n are independent events, then

$$\mathbb{P}\left(\bigcup_{i=1}^n A_i\right) = 1 - \prod_{i=1}^n [1 - \mathbb{P}(A_i)] \quad (2.5.10)$$

Proof

From DeMorgan's law and the independence of $A_1^c, A_2^c, \dots, A_n^c$ we have

$$\mathbb{P}\left(\bigcup_{i=1}^n A_i\right) = 1 - \mathbb{P}\left(\bigcap_{i=1}^n A_i^c\right) = 1 - \prod_{i=1}^n \mathbb{P}(A_i^c) = 1 - \prod_{i=1}^n [1 - \mathbb{P}(A_i)] \quad (2.5.11)$$

Independence of Random Variables

Suppose now that X_i is a random variable for the experiment with values in a set T_i for each i in a nonempty index set I . Mathematically, X_i is a function from S into T_i , and recall that $\{X_i \in B\}$ denotes the event $\{s \in S : X_i(s) \in B\}$ for $B \subseteq T_i$. Intuitively, X_i is a variable of interest in the experiment, and every meaningful statement about X_i defines an event. Intuitively, the random variables are independent if information about some of the variables tells us nothing about the other variables. Mathematically, independence of a collection of random variables can be reduced to the independence of collections of events.

The collection of random variables $\mathcal{X} = \{X_i : i \in I\}$ is *independent* if the collection of events $\{\{X_i \in B_i\} : i \in I\}$ is independent for every choice of $B_i \subseteq T_i$ for $i \in I$. Equivalently then, \mathcal{X} is independent if for every finite $J \subseteq I$, and for every choice of $B_j \subseteq T_j$

for $j \in J$ we have

$$\mathbb{P}\left(\bigcap_{j \in J} \{X_j \in B_j\}\right) = \prod_{j \in J} \mathbb{P}(X_j \in B_j) \quad (2.5.12)$$

Details

Recall that T_i will have a σ -algebra \mathcal{T}_i of admissible subsets so that (T_i, \mathcal{T}_i) is a measurable space just like the sample space (S, \mathcal{S}) for each $i \in I$. Also X_i is measurable as a function from S into T_i for each $i \in I$. These technical assumptions ensure that the definition makes sense.

Suppose that \mathcal{X} is a collection of random variables.

1. If \mathcal{X} is independent, then \mathcal{Y} is independent for every $\mathcal{Y} \subseteq \mathcal{X}$
2. If \mathcal{Y} is independent for every finite $\mathcal{Y} \subseteq \mathcal{X}$ then \mathcal{X} is independent.

It would seem almost obvious that if a collection of random variables is independent, and we transform each variable in deterministic way, then the new collection of random variables should still be independent.

Suppose now that g_i is a function from T_i into a set U_i for each $i \in I$. If $\{X_i : i \in I\}$ is independent, then $\{g_i(X_i) : i \in I\}$ is also independent.

Proof

Except for the abstract setting, the proof of independence is easy. Suppose that $C_i \subseteq U_i$ for each $i \in I$. Then $\{g_i(X_i) \in C_i\} = \{X_i \in g_i^{-1}(C_i)\}$ for $i \in I$. By the independence of $\{X_i : i \in I\}$, the collection of events $\{\{X_i \in g_i^{-1}(C_i)\} : i \in I\}$ is independent.

Technically, the set U_i will have a σ -algebra \mathcal{U}_i of admissible subsets so that (U_i, \mathcal{U}_i) is a measurable space just like (T_i, \mathcal{T}_i) and just like the sample space (S, \mathcal{S}) . The function g_i is required to be measurable as a function from T_i into U_i just as X_i is measurable as a function from S into T_i . In the proof above, $C_i \in \mathcal{U}_i$ so that $g_i^{-1}(C_i) \in \mathcal{T}_i$ and hence $\{X_i \in g_i^{-1}(C_i)\} \in \mathcal{S}$.

As with events, the (mutual) independence of random variables is a very strong property. If a collection of random variables is independent, then any subcollection is also independent. New random variables formed from disjoint subcollections are independent. For a simple example, suppose that X , Y , and Z are independent real-valued random variables. Then

1. $\sin(X)$, $\cos(Y)$, and e^Z are independent.
2. (X, Y) and Z are independent.
3. $X^2 + Y^2$ and $\arctan(Z)$ are independent.
4. X and Z are independent.
5. Y and Z are independent.

In particular, note that statement 2 in the list above is much stronger than the conjunction of statements 4 and 5. Contrapositively, if X and Z are dependent, then (X, Y) and Z are also dependent. Independence of random variables subsumes independence of events.

A collection of events \mathcal{A} is independent if and only if the corresponding collection of indicator variables $\{\mathbf{1}_{A_i} : A_i \in \mathcal{A}\}$ is independent.

Proof

Let $\mathcal{A} = \{A_i : i \in I\}$ where I is a nonempty index set. For $i \in I$, the only non-trivial events that can be defined in terms of $\mathbf{1}_{A_i}$ are $\{\mathbf{1}_{A_i} = 1\} = A_i$ and $\{\mathbf{1}_{A_i} = 0\} = A_i^c$. So $\{\mathbf{1}_{A_i} : i \in I\}$ is independent if and only if every collection of the form $\{B_i : i \in I\}$ is independent, where for each $i \in I$, either $B_i = A_i$ or $B_i = A_i^c$. But by the [complement theorem](#), this is equivalent to the independence of $\{A_i : i \in I\}$.

Many of the concepts that we have been using informally can now be made precise. A compound experiment that consists of “independent stages” is essentially just an experiment whose outcome is a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots)$ where X_i is the outcome of the i th stage.

In particular, suppose that we have a basic experiment with outcome variable X . By definition, the outcome of the experiment that consists of “independent replications” of the basic experiment is a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots)$ each with the same probability distribution as X . This is fundamental to the very concept of probability, as expressed in the law of large numbers. From a statistical point of view, suppose that we have a population of objects and a vector of measurements \mathbf{X} of interest for the objects in the

sample. The sequence \mathbf{X} above corresponds to *sampling from the distribution of X* ; that is, X_i is the vector of measurements for the i th object drawn from the sample. When we sample from a finite population, sampling *with replacement* generates independent random variables while sampling *without replacement* generates dependent random variables.

Conditional Independence and Conditional Probability

As noted at the beginning of our discussion, independence of events or random variables depends on the underlying probability measure. Thus, suppose that B is an event with positive probability. A collection of events or a collection of random variables is *conditionally independent given B* if the collection is independent relative to the conditional probability measure $A \mapsto \mathbb{P}(A | B)$. For example, a collection of events $\{A_i : i \in I\}$ is conditionally independent given B if for every finite $J \subseteq I$,

$$\mathbb{P}\left(\bigcap_{j \in J} A_j \mid B\right) = \prod_{j \in J} \mathbb{P}(A_j | B) \quad (2.5.13)$$

Note that the definitions and theorems of this section would still be true, but with all probabilities conditioned on B .

Conversely, conditional probability has a nice interpretation in terms of independent replications of the experiment. Thus, suppose that we start with a basic experiment with S as the set of outcomes. We let X denote the outcome random variable, so that mathematically X is simply the identity function on S . In particular, if A is an event then trivially, $\mathbb{P}(X \in A) = \mathbb{P}(A)$. Suppose now that we replicate the experiment independently. This results in a new, compound experiment with a sequence of independent random variables (X_1, X_2, \dots) , each with the same distribution as X . That is, X_i is the outcome of the i th repetition of the experiment.

Suppose now that A and B are events in the basic experiment with $\mathbb{P}(B) > 0$. In the compound experiment, the event that “when B occurs for the first time, A also occurs” has probability

$$\frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} = \mathbb{P}(A | B) \quad (2.5.14)$$

Proof

In the compound experiment, if we record (X_1, X_2, \dots) then the new set of outcomes is $S^\infty = S \times S \times \dots$. The event that “when B occurs for the first time, A also occurs” is

$$\bigcup_{n=1}^{\infty} \{X_1 \notin B, X_2 \notin B, \dots, X_{n-1} \notin B, X_n \in A \cap B\} \quad (2.5.15)$$

The events in the union are disjoint. Also, since (X_1, X_2, \dots) is a sequence of independent variables, each with the distribution of X we have

$$\mathbb{P}(X_1 \notin B, X_2 \notin B, \dots, X_{n-1} \notin B, X_n \in A \cap B) = [\mathbb{P}(B^c)]^{n-1} \mathbb{P}(A \cap B) = [1 - \mathbb{P}(B)]^{n-1} \mathbb{P}(A \cap B) \quad (2.5.16)$$

Hence, using geometric series, the probability of the union is

$$\sum_{n=1}^{\infty} [1 - \mathbb{P}(B)]^{n-1} \mathbb{P}(A \cap B) = \frac{\mathbb{P}(A \cap B)}{1 - [1 - \mathbb{P}(B)]} = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \quad (2.5.17)$$

Heuristic Argument

Suppose that we create a new experiment by repeating the basic experiment until B occurs for the first time, and then record the outcome of just the last repetition of the basic experiment. Now the set of outcomes is simply B and the appropriate probability measure on the new experiment is $A \mapsto \mathbb{P}(A | B)$.

Suppose that A and B are disjoint events in a basic experiment with $\mathbb{P}(A) > 0$ and $\mathbb{P}(B) > 0$. In the compound experiment obtained by replicating the basic experiment, the event that “ A occurs before B ” has probability

$$\frac{\mathbb{P}(A)}{\mathbb{P}(A) + \mathbb{P}(B)} \quad (2.5.18)$$

Proof

Note that the event “ A occurs before B ” is the same as the event “when $A \cup B$ occurs for the first time, A occurs”.

Examples and Applications

Basic Rules

Suppose that A , B , and C are independent events in an experiment with $\mathbb{P}(A) = 0.3$, $\mathbb{P}(B) = 0.4$, and $\mathbb{P}(C) = 0.8$. Express each of the following events in set notation and find its probability:

1. All three events occur.
2. None of the three events occurs.
3. At least one of the three events occurs.
4. At least one of the three events does not occur.
5. Exactly one of the three events occurs.
6. Exactly two of the three events occurs.

Answer

1. $\mathbb{P}(A \cap B \cap C) = 0.096$
2. $\mathbb{P}(A^c \cap B^c \cap C^c) = 0.084$
3. $\mathbb{P}(A \cup B \cup C) = 0.916$
4. $\mathbb{P}(A^c \cup B^c \cup C^c) = 0.904$
5. $\mathbb{P}[(A \cap B^c \cap C^c) \cup (A^c \cap B \cap C^c) \cup (A^c \cap B^c \cap C)] = 0.428$
6. $\mathbb{P}[(A \cap B \cap C^c) \cup (A \cap B^c \cap C) \cup (A^c \cap B \cap C)] = 0.392$

Suppose that A , B , and C are independent events for an experiment with $\mathbb{P}(A) = \frac{1}{3}$, $\mathbb{P}(B) = \frac{1}{4}$, and $\mathbb{P}(C) = \frac{1}{5}$. Find the probability of each of the following events:

1. $(A \cap B) \cup C$
2. $A \cup B^c \cup C$
3. $(A^c \cap B^c) \cup C^c$

Answer

1. $\frac{4}{15}$
2. $\frac{13}{15}$
3. $\frac{9}{10}$

Simple Populations

A small company has 100 employees; 40 are men and 60 are women. There are 6 male executives. How many female executives should there be if gender and rank are independent? The underlying experiment is to choose an employee at random.

Answer

9

Suppose that a farm has four orchards that produce peaches, and that peaches are classified by size as small, medium, and large. The table below gives total number of peaches in a recent harvest by orchard and by size. Fill in the body of the table with counts for the various intersections, so that orchard and size are independent variables. The underlying experiment is to select a peach at random from the farm.

Frequency	Size Small	Medium	Large	Total
Orchard 1			Total2000	400
2			Total2000	600
3			Total2000	300
4			Total2000	700
Total	400	1000	600	2000

Answer

Frequency	Size Small	Medium	Large	Total
-----------	------------	--------	-------	-------

Orchard 1	80	200	120	400
2	120	300	180	600
3	60	150	90	300
4	140	350	210	700
total	400	1000	600	2000

Note from the last two exercises that you cannot “see” independence in a Venn diagram. Again, independence is a measure-theoretic concept, not a set-theoretic concept.

Bernoulli Trials

A *Bernoulli trials sequence* is a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent, identically distributed indicator variables. Random variable X_i is the outcome of trial i , where in the usual terminology of reliability theory, 1 denotes *success* and 0 denotes *failure*. The canonical example is the sequence of scores when a coin (not necessarily fair) is tossed repeatedly. Another basic example arises whenever we start with an basic experiment and an event A of interest, and then repeat the experiment. In this setting, X_i is the indicator variable for event A on the i th run of the experiment. The Bernoulli trials process is named for Jacob Bernoulli, and has a single basic *parameter* $p = \mathbb{P}(X_i = 1)$. This random process is studied in detail in the chapter on Bernoulli trials.

For $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$,

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = p^{x_1 + x_2 + \dots + x_n} (1 - p)^{n - (x_1 + x_2 + \dots + x_n)} \quad (2.5.19)$$

Proof

If X is a generic Bernoulli trial, then by definition, $\mathbb{P}(X = 1) = p$ and $\mathbb{P}(X = 0) = 1 - p$. Equivalently, $\mathbb{P}(X = x) = p^x (1 - p)^{1-x}$ for $x \in \{0, 1\}$. Thus the result follows by independence.

Note that the sequence of indicator random variables \mathbf{X} is exchangeable. That is, if the sequence (x_1, x_2, \dots, x_n) in the previous result is permuted, the probability does not change. On the other hand, there are exchangeable sequences of indicator random variables that are dependent, as Pólya's urn model so dramatically illustrates.

Let Y denote the number of successes in the first n trials. Then

$$\mathbb{P}(Y = y) = \binom{n}{y} p^y (1 - p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (2.5.20)$$

Proof

Note that $Y = \sum_{i=1}^n X_i$, where X_i is the outcome of trial i , as in the previous result. For $y \in \{0, 1, \dots, n\}$, the event $\{Y = y\}$ occurs if and only if exactly y of the n trials result in success (1). The number of ways to choose the y trials that result in success is $\binom{n}{y}$, and by the previous result, the probability of any particular sequence of y successes and $n - y$ failures is $p^y (1 - p)^{n-y}$. Thus the result follows by the additivity of probability.

The distribution of Y is called the *binomial distribution* with parameters n and p . The binomial distribution is studied in more detail in the chapter on Bernoulli Trials.

More generally, a *multinomial trials sequence* is a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent, identically distributed random variables, each taking values in a finite set S . The canonical example is the sequence of scores when a k -sided die (not necessarily fair) is thrown repeatedly. Multinomial trials are also studied in detail in the chapter on Bernoulli trials.

Cards

Consider the experiment that consists of dealing 2 cards at random from a standard deck and recording the sequence of cards dealt. For $i \in \{1, 2\}$, let Q_i be the event that card i is a queen and H_i the event that card i is a heart. Compute the appropriate probabilities to verify the following results. Reflect on these results.

1. Q_1 and H_1 are independent.
2. Q_2 and H_2 are independent.
3. Q_1 and Q_2 are negatively correlated.

4. H_1 and H_2 are negatively correlated.
5. Q_1 and H_2 are independent.
6. H_1 and Q_2 are independent.

Answer

1. $\mathbb{P}(Q_1) = \mathbb{P}(Q_1 | H_1) = \frac{1}{13}$
2. $\mathbb{P}(Q_2) = \mathbb{P}(Q_2 | H_2) = \frac{1}{13}$
3. $\mathbb{P}(Q_1) = \frac{1}{13}$, $\mathbb{P}(Q_1 | Q_2) = \frac{1}{17}$
4. $\mathbb{P}(H_1) = \frac{1}{4}$, $\mathbb{P}(H_1 | H_2) = \frac{4}{17}$
5. $\mathbb{P}(Q_1) = \mathbb{P}(Q_1 | H_2) = \frac{1}{13}$
6. $\mathbb{P}(Q_2) = \mathbb{P}(Q_2 | H_1) = \frac{1}{13}$

In the card experiment, set $n = 2$. Run the simulation 500 times. For each pair of events in the previous exercise, compute the product of the empirical probabilities and the empirical probability of the intersection. Compare the results.

Dice

The following exercise gives three events that are *pairwise independent*, but not (*mutually*) *independent*.

Consider the dice experiment that consists of rolling 2 standard, fair dice and recording the sequence of scores. Let A denote the event that first score is 3, B the event that the second score is 4, and C the event that the sum of the scores is 7. Then

1. A , B , C are pairwise independent.
2. $A \cap B$ implies (is a subset of) C and hence these events are dependent in the strongest possible sense.

Answer

Note that $A \cap B = A \cap C = B \cap C = \{(3, 4)\}$, and the probability of the common intersection is $\frac{1}{36}$. On the other hand, $\mathbb{P}(A) = \mathbb{P}(B) = \mathbb{P}(C) = \frac{6}{36} = \frac{1}{6}$.

In the dice experiment, set $n = 2$. Run the experiment 500 times. For each pair of events in the previous exercise, compute the product of the empirical probabilities and the empirical probability of the intersection. Compare the results.

The following exercise gives an example of three events with the property that the probability of the intersection is the product of the probabilities, but the events are not pairwise independent.

Suppose that we throw a standard, fair die one time. Let $A = \{1, 2, 3, 4\}$, $B = C = \{4, 5, 6\}$. Then

1. $\mathbb{P}(A \cap B \cap C) = \mathbb{P}(A)\mathbb{P}(B)\mathbb{P}(C)$.
2. B and C are the same event, and hence are dependent in the strongest possible sense.

Answer

Note that $A \cap B \cap C = \{4\}$, so $\mathbb{P}(A \cap B \cap C) = \frac{1}{6}$. On the other hand, $\mathbb{P}(A) = \frac{4}{6}$ and $\mathbb{P}(B) = \mathbb{P}(C) = \frac{3}{6}$.

Suppose that a standard, fair die is thrown 4 times. Find the probability of the following events.

1. Six does not occur.
2. Six occurs at least once.
3. The sum of the first two scores is 5 and the sum of the last two scores is 7.

Answer

1. $\left(\frac{5}{6}\right)^4 \approx 0.4823$
2. $1 - \left(\frac{5}{6}\right)^4 \approx 0.5177$
3. $\frac{1}{54}$

Suppose that a pair of standard, fair dice are thrown 8 times. Find the probability of each of the following events.

1. Double six does not occur.
2. Double six occurs at least once.

3. Double six does not occur on the first 4 throws but occurs at least once in the last 4 throws.

Answer

1. $\left(\frac{35}{36}\right)^8 \approx 0.7982$
2. $1 - \left(\frac{35}{36}\right)^8 \approx 0.2018$
3. $\left(\frac{35}{36}\right)^4 \left[1 - \left(\frac{35}{36}\right)^4\right] \approx 0.0952$

Consider the dice experiment that consists of rolling n , k -sided dice and recording the sequence of scores $\mathbf{X} = (X_1, X_2, \dots, X_n)$. The following conditions are equivalent (and correspond to the assumption that the dice are fair):

1. \mathbf{X} is uniformly distributed on $\{1, 2, \dots, k\}^n$.
2. \mathbf{X} is a sequence of independent variables, and X_i is uniformly distributed on $\{1, 2, \dots, k\}$ for each i .

Proof

Let $S = \{1, 2, \dots, k\}$ and note that S^n has k^n points. Suppose that \mathbf{X} is uniformly distributed on S^n . Then $\mathbb{P}(\mathbf{X} = \mathbf{x}) = 1/k^n$ for each $\mathbf{x} \in S^n$ so $\mathbb{P}(X_i = x) = k^{n-1}/k^n = 1/k$ for each $x \in S$. Hence X_i is uniformly distributed on S . Moreover,

$$\mathbb{P}(\mathbf{X} = \mathbf{x}) = \mathbb{P}(X_1 = x_1)\mathbb{P}(X_2 = x_2) \cdots \mathbb{P}(X_n = x_n), \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in S^n \quad (2.5.21)$$

so \mathbf{X} is an independent sequence. Conversely, if \mathbf{X} is an independent sequence and X_i is uniformly distributed on S for each i then $\mathbb{P}(X_i = x) = 1/k$ for each $x \in S$ and hence $\mathbb{P}(\mathbf{X} = \mathbf{x}) = 1/k^n$ for each $\mathbf{x} \in S^n$. Thus \mathbf{X} is uniformly distributed on S^n .

A pair of standard, fair dice are thrown repeatedly. Find the probability of each of the following events.

1. A sum of 4 occurs before a sum of 7.
2. A sum of 5 occurs before a sum of 7.
3. A sum of 6 occurs before a sum of 7.
4. When a sum of 8 occurs the first time, it occurs “the hard way” as $(4, 4)$.

Answer

1. $\frac{3}{9}$
2. $\frac{4}{10}$
3. $\frac{5}{11}$
4. $\frac{1}{5}$

Problems of the type in the last exercise are important in the game of *craps*. Craps is studied in more detail in the chapter on Games of Chance.

Coins

A biased coin with probability of heads $\frac{1}{3}$ is tossed 5 times. Let \mathbf{X} denote the outcome of the tosses (encoded as a bit string) and let Y denote the number of heads. Find each of the following:

1. $\mathbb{P}(\mathbf{X} = \mathbf{x})$ for each $\mathbf{x} \in \{0, 1\}^5$.
2. $\mathbb{P}(Y = y)$ for each $y \in \{0, 1, 2, 3, 4, 5\}$
3. $\mathbb{P}(1 \leq Y \leq 3)$

Answer

1. $\frac{32}{243}$ if $\mathbf{x} = 00000$, $\frac{16}{243}$ if \mathbf{x} has exactly one 1 (there are 5 of these), $\frac{8}{243}$ if \mathbf{x} has exactly two 1s (there are 10 of these), $\frac{4}{243}$ if \mathbf{x} has exactly three 1s (there are 10 of these), $\frac{2}{243}$ if \mathbf{x} has exactly four 1s (there are 5 of these), $\frac{1}{243}$ if $\mathbf{x} = 11111$
2. $\frac{32}{243}$ if $y = 0$, $\frac{80}{243}$ if $y = 1$, $\frac{80}{243}$ if $y = 2$, $\frac{40}{243}$ if $y = 3$, $\frac{10}{243}$ if $y = 4$, $\frac{1}{243}$ if $y = 5$
3. $\frac{200}{243}$

A box contains a fair coin and a two-headed coin. A coin is chosen at random from the box and tossed repeatedly. Let F denote the event that the fair coin is chosen, and let H_i denote the event that the i th toss results in heads. Then

1. (H_1, H_2, \dots) are conditionally independent given F , with $\mathbb{P}(H_i | F) = \frac{1}{2}$ for each i .
2. (H_1, H_2, \dots) are conditionally independent given F^c , with $\mathbb{P}(H_i | F^c) = 1$ for each i .

3. $\mathbb{P}(H_i) = \frac{3}{4}$ for each i .
4. $\mathbb{P}(H_1 \cap H_2 \cap \dots \cap H_n) = \frac{1}{2^{n+1}} + \frac{1}{2}$.
5. (H_1, H_2, \dots) are dependent.
6. $\mathbb{P}(F | H_1 \cap H_2 \cap \dots \cap H_n) = \frac{1}{2^{n+1}}$.
7. $\mathbb{P}(F | H_1 \cap H_2 \cap \dots \cap H_n) \rightarrow 0$ as $n \rightarrow \infty$.

Proof

Parts (a) and (b) are essentially modeling assumptions, based on the design of the experiment. If we know what kind of coin we have, then the tosses are independent. Parts (c) and (d) follow by conditioning on the type of coin and using parts (a) and (b). Part (e) follows from (c) and (d). Note that the expression in (d) is not $(3/4)^n$. Part (f) follows from part (d) and Bayes' theorem. Finally part (g) follows from part (f).

Consider again the box in the previous exercise, but we change the experiment as follows: a coin is chosen at random from the box and tossed and the result recorded. The coin is returned to the box and the process is repeated. As before, let H_i denote the event that toss i results in heads. Then

1. (H_1, H_2, \dots) are independent.
2. $\mathbb{P}(H_i) = \frac{3}{4}$ for each i .
3. $\mathbb{P}(H_1 \cap H_2 \cap \dots \cap H_n) = \left(\frac{3}{4}\right)^n$.

Proof

Again, part (a) is essentially a modeling assumption. Since we return the coin and draw a new coin at random each time, the results of the tosses should be independent. Part (b) follows by conditioning on the type of the i th coin. Part (c) follows from parts (a) and (b).

Think carefully about the results in the previous two exercises, and the differences between the two models. Tossing a coin produces independent random variables *if* the probability of heads is fixed (that is, non-random even if unknown). Tossing a coin with a random probability of heads generally does not produce independent random variables; the result of a toss gives information about the probability of heads which in turn gives information about subsequent tosses.

Uniform Distributions

Recall that Buffon's coin experiment consists of tossing a coin with radius $r \leq \frac{1}{2}$ randomly on a floor covered with square tiles of side length 1. The coordinates (X, Y) of the center of the coin are recorded relative to axes through the center of the square in which the coin lands. The following conditions are equivalent:

1. (X, Y) is uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]^2$.
2. X and Y are independent and each is uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$.




Figure 2.5.1: Buffon's coin experiment

Proof

Let $S = [-\frac{1}{2}, \frac{1}{2}]$, and let λ_1 denote length measure on S and λ_2 area measure on S^2 . Note that $\lambda_1(S) = \lambda_2(S^2) = 1$. Suppose that (X, Y) is uniformly distributed on S^2 , so that $\mathbb{P}[(X, Y) \in C] = \lambda_2(C)$ for $C \subseteq S^2$. For $A \subseteq S$,

$$\mathbb{P}(X \in A) = \mathbb{P}[(X, Y) \in A \times S] = \lambda_2(A \times S) = \lambda_1(A) \quad (2.5.22)$$

Hence X is uniformly distributed on S . By a similar argument, Y is also uniformly distributed on S . Moreover, for $A \subseteq S$ and $B \subseteq S$,

$$\mathbb{P}(X \in A, Y \in B) = \mathbb{P}[(X, Y) \in A \times B] = \lambda_2(A \times B) = \lambda_1(A)\lambda_1(B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B) \quad (2.5.23)$$

so X and Y are independent. Conversely, if X and Y are independent and each is uniformly distributed on S , then for $A \subseteq S$ and $B \subseteq S$,

$$\mathbb{P}[(X, Y) \in A \times B] = \mathbb{P}(X \in A)\mathbb{P}(Y \in B) = \lambda_1(A)\lambda_1(B) = \lambda_2(A \times B) \quad (2.5.24)$$

It then follows that $\mathbb{P}[(X, Y) \in C] = \lambda_2(C)$ for every $C \subseteq S^2$. For more details about this last step, see the advanced section on existence and uniqueness of measures.

Compare this result with the result above for [fair dice](#).

In Buffon's coin experiment, set $r = 0.3$. Run the simulation 500 times. For the events $\{X > 0\}$ and $\{Y < 0\}$, compute the product of the empirical probabilities and the empirical probability of the intersection. Compare the results.

The arrival time X of the A train is uniformly distributed on the interval $(0, 30)$, while the arrival time Y of the B train is uniformly distributed on the interval $(15, 30)$. (The arrival times are in minutes, after 8:00 AM). Moreover, the arrival times are independent. Find the probability of each of the following events:

1. The A train arrives first.
2. Both trains arrive sometime after 20 minutes.

Answer

1. $\frac{3}{4}$
2. $\frac{2}{9}$

Reliability

Recall the simple model of structural reliability in which a system is composed of n components. Suppose in addition that the components operate independently of each other. As before, let X_i denote the state of component i , where 1 means working and 0 means failure. Thus, our basic assumption is that the *state vector* $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent indicator random variables. We assume that the state of the system (either working or failed) depends only on the states of the components. Thus, the state of the system is an indicator random variable

$$Y = y(X_1, X_2, \dots, X_n) \quad (2.5.25)$$

where $y : \{0, 1\}^n \rightarrow \{0, 1\}$ is the *structure function*. Generally, the probability that a device is working is the *reliability* of the device. Thus, we will denote the reliability of component i by $p_i = \mathbb{P}(X_i = 1)$ so that the vector of component reliabilities is $\mathbf{p} = (p_1, p_2, \dots, p_n)$. By independence, the system reliability r is a function of the component reliabilities:

$$r(p_1, p_2, \dots, p_n) = \mathbb{P}(Y = 1) \quad (2.5.26)$$

Appropriately enough, this function is known as the *reliability function*. Our challenge is usually to find the reliability function, given the structure function. When the components all have the same probability p then of course the system reliability r is just a function of p . In this case, the state vector $\mathbf{X} = (X_1, X_2, \dots, X_n)$ forms a sequence of [Bernoulli trials](#).

Comment on the independence assumption for real systems, such as your car or your computer.

Recall that a *series system* is working if and only if each component is working.

1. The state of the system is $U = X_1 X_2 \cdots X_n = \min\{X_1, X_2, \dots, X_n\}$.
2. The reliability is $\mathbb{P}(U = 1) = p_1 p_2 \cdots p_n$.

Recall that a *parallel system* is working if and only if at least one component is working.

1. The state of the system is $V = 1 - (1 - X_1)(1 - X_2) \cdots (1 - X_n) = \max\{X_1, X_2, \dots, X_n\}$.
2. The reliability is $\mathbb{P}(V = 1) = 1 - (1 - p_1)(1 - p_2) \cdots (1 - p_n)$.

Recall that a *k out of n system* is working if and only if at least k of the n components are working. Thus, a parallel system is a 1 out of n system and a series system is an n out of n system. A k out of $2k - 1$ system is a *majority rules system*. The reliability function of a general k out of n system is a mess. However, if the component reliabilities are the same, the function has a reasonably simple form.

For a k out of n system with common component reliability p , the system reliability is

$$r(p) = \sum_{i=k}^n \binom{n}{i} p^i (1-p)^{n-i} \quad (2.5.27)$$

Consider a system of 3 independent components with common reliability $p = 0.8$. Find the reliability of each of the following:

1. The parallel system.
2. The 2 out of 3 system.
3. The series system.

Answer

1. 0.992
2. 0.896
3. 0.512

Consider a system of 3 independent components with reliabilities $p_1 = 0.8$, $p_2 = 0.8$, $p_3 = 0.7$. Find the reliability of each of the following:

1. The parallel system.
2. The 2 out of 3 system.
3. The series system.

Answer

1. 0.994
2. 0.902
3. 0.504

Consider an airplane with an odd number of engines, each with reliability p . Suppose that the airplane is a majority rules system, so that the airplane needs a majority of working engines in order to fly.

1. Find the reliability of a 3 engine plane as a function of p .
2. Find the reliability of a 5 engine plane as a function of p .
3. For what values of p is a 5 engine plane preferable to a 3 engine plane?

Answer

1. $r_3(p) = 3p^2 - 2p^3$
2. $r_5(p) = 6p^5 - 15p^4 + 10p^3$
3. The 5-engine plane would be preferable if $p > \frac{1}{2}$ (which one would hope would be the case). The 3-engine plane would be preferable if $p < \frac{1}{2}$. If $p = \frac{1}{2}$, the 3-engine and 5-engine planes are equally reliable.

The graph below is known as the *Wheatstone bridge network* and is named for Charles Wheatstone. The edges represent components, and the system works if and only if there is a working path from vertex a to vertex b .

1. Find the structure function.
2. Find the reliability function.




Figure 2.5.2: The Wheatstone bridge network

Answer

1. $Y = X_3(X_1 + X_2 - X_1X_2)(X_4 + X_5 - X_4X_5) + (1 - X_3)(X_1X_4 + X_2X_5 - X_1X_2X_4X_5)$
2. $r(p_1, p_2, p_3, p_4, p_5) = p_3(p_1 + p_2 - p_1p_2)(p_4 + p_5 - p_4p_5) + (1 - p_3)(p_1p_4 + p_2p_5 - p_1p_2p_4p_5)$

A system consists of 3 components, connected in parallel. Because of environmental factors, the components do not operate independently, so our usual assumption does not hold. However, we will assume that under low stress conditions, the components are independent, each with reliability 0.9; under medium stress conditions, the components are independent with reliability 0.8; and under high stress conditions, the components are independent, each with reliability 0.7. The probability of low stress is 0.5, of medium stress is 0.3, and of high stress is 0.2.

1. Find the reliability of the system.
2. Given that the system works, find the conditional probability of each stress level.

Answer

1. 0.9917. Condition on the stress level.
2. 0.5037 for low, 0.3001 for medium, 0.1962 for high. Use Bayes' theorem and part (a).

Suppose that bits are transmitted across a noisy communications channel. Each bit that is sent, independently of the others, is received correctly with probability 0.9 and changed to the complementary bit with probability 0.1. Using redundancy to improve reliability, suppose that a given bit will be sent 3 times. We naturally want to compute the probability that we correctly identify the bit that was sent. Assume we have no prior knowledge of the bit, so we assign probability $\frac{1}{2}$ each to the event that 000 was sent and the event that 111 was sent. Now find the conditional probability that 111 was sent given each of the 8 possible bit strings received.

Answer

Let \mathbf{X} denote the string sent and \mathbf{Y} the string received.

\mathbf{y}	$\mathbb{P}(\mathbf{X} = 111 \mid \mathbf{Y} = \mathbf{y})$
111	729/730
110	9/10
101	9/10
011	9/10
100	1/10
010	1/10
001	1/10
000	1/730

Diagnostic Testing

Recall the discussion of diagnostic testing in the section on Conditional Probability. Thus, we have an event A for a random experiment whose occurrence or non-occurrence we cannot observe directly. Suppose now that we have n tests for the occurrence of A , labeled from 1 to n . We will let T_i denote the event that test i is positive for A . The tests are independent in the following sense:

- If A occurs, then (T_1, T_2, \dots, T_n) are (conditionally) independent and test i has sensitivity $a_i = \mathbb{P}(T_i \mid A)$.
- If A does not occur, then (T_1, T_2, \dots, T_n) are (conditionally) independent and test i has specificity $b_i = \mathbb{P}(T_i^c \mid A^c)$.

Note that *unconditionally*, it is not reasonable to assume that the tests are independent. For example, a positive result for a given test presumably is evidence that the condition A has occurred, which in turn is evidence that a subsequent test will be positive. In short, we expect that T_i and T_j should be positively correlated.

We can form a new, compound test by giving a *decision rule* in terms of the individual test results. In other words, the event T that the compound test is positive for A is a function of (T_1, T_2, \dots, T_n) . The typical decision rules are very similar to the [reliability structures](#) discussed above. A special case of interest is when the n tests are independent applications of a given basic test. In this case, $a_i = a$ and $b_i = b$ for each i .

Consider the compound test that is positive for A if and only if each of the n tests is positive for A .

1. $T = T_1 \cap T_2 \cap \dots \cap T_n$
2. The sensitivity is $\mathbb{P}(T \mid A) = a_1 a_2 \dots a_n$.
3. The specificity is $\mathbb{P}(T^c \mid A^c) = 1 - (1 - b_1)(1 - b_2) \dots (1 - b_n)$

Consider the compound test that is positive for A if and only if each at least one of the n tests is positive for A .

1. $T = T_1 \cup T_2 \cup \dots \cup T_n$
2. The sensitivity is $\mathbb{P}(T \mid A) = 1 - (1 - a_1)(1 - a_2) \dots (1 - a_n)$.
3. The specificity is $\mathbb{P}(T^c \mid A^c) = b_1 b_2 \dots b_n$.

More generally, we could define the compound k out of n test that is positive for A if and only if at least k of the individual tests are positive for A . The [series test](#) is the n out of n test, while the [parallel test](#) is the 1 out of n test. The k out of $2k - 1$ test is the *majority rules test*.

Suppose that a woman initially believes that there is an even chance that she is or is not pregnant. She buys three identical pregnancy tests with sensitivity 0.95 and specificity 0.90. Tests 1 and 3 are positive and test 2 is negative.

1. Find the updated probability that the woman is pregnant.
2. Can we just say that tests 2 and 3 cancel each other out? Find the probability that the woman is pregnant given just one positive test, and compare the answer with the answer to part (a).

Answer

1. 0.834
2. No: 0.905.

Suppose that 3 independent, identical tests for an event A are applied, each with sensitivity a and specificity b . Find the sensitivity and specificity of the following tests:

1. 1 out of 3 test
2. 2 out of 3 test
3. 3 out of 3 test

Answer

1. sensitivity $1 - (1 - a)^3$, specificity b^3
2. sensitivity $3a^2$, specificity $b^3 + 3b^2(1 - b)$
3. sensitivity a^3 , specificity $1 - (1 - b)^3$

In a criminal trial, the defendant is convicted if and only if all 6 jurors vote guilty. Assume that if the defendant really is guilty, the jurors vote guilty, independently, with probability 0.95, while if the defendant is really innocent, the jurors vote not guilty, independently with probability 0.8. Suppose that 70% of defendants brought to trial are guilty.

1. Find the probability that the defendant is convicted.
2. Given that the defendant is convicted, find the probability that the defendant is guilty.
3. Comment on the assumption that the jurors act independently.

Answer

1. 0.5148
2. 0.99996
3. The independence assumption is not reasonable since jurors collaborate.

Genetics

Please refer to the discussion of genetics in the section on random experiments if you need to review some of the definitions in this section.

Recall first that the *ABO blood type* in humans is determined by three alleles: a , b , and o . Furthermore, a and b are co-dominant and o is recessive. Suppose that in a certain population, the proportion of a , b , and o alleles are p , q , and r respectively. Of course we must have $p > 0$, $q > 0$, $r > 0$ and $p + q + r = 1$.

Suppose that the blood genotype in a person is the result of independent alleles, chosen with probabilities p , q , and r as above.

1. The probability distribution of the genotypes is given in the following table:

Genotype	aa	ab	ao	bb	bo	oo
Probability	p^2	$2pq$	$2pr$	q^2	$2qr$	r^2

2. The probability distribution of the blood types is given in the following table:

Blood type	A	B	AB	O
Probability	$p^2 + 2pr$	$q^2 + 2qr$	$2pq$	r^2

Proof

Part (a) follows from the independence assumption and basic rules of probability. Even though genotypes are listed as unordered pairs, note that there are two ways that a heterozygous genotype can occur, since either parent could contribute either of the two distinct alleles. Part (b) follows from part (a) and basic rules of probability.

The discussion above is related to the *Hardy-Weinberg model* of genetics. The model is named for the English mathematician Godfrey Hardy and the German physician Wilhelm Weiberg

Suppose that the probability distribution for the set of blood types in a certain population is given in the following table:

Blood type	A	B	AB	O
Probability	0.360	0.123	0.038	0.479

Find p , q , and r .

Answer

$$p = 0.224, q = 0.084, r = 0.692$$

Suppose next that pod color in certain type of pea plant is determined by a gene with two alleles: g for green and y for yellow, and that g is dominant and o recessive.

Suppose that 2 green-pod plants are bred together. Suppose further that each plant, independently, has the recessive yellow-pod allele with probability $\frac{1}{4}$.

1. Find the probability that 3 offspring plants will have green pods.
2. Given that the 3 offspring plants have green pods, find the updated probability that both parents have the recessive allele.

Answer

1. $\frac{987}{1024}$
2. $\frac{27}{987}$

Next consider a sex-linked hereditary disorder in humans (such as colorblindness or hemophilia). Let h denote the healthy allele and d the defective allele for the gene linked to the disorder. Recall that h is dominant and d recessive for women.

Suppose that a healthy woman initially has a $\frac{1}{2}$ chance of being a carrier. (This would be the case, for example, if her mother and father are healthy but she has a brother with the disorder, so that her mother must be a carrier).

1. Find the probability that the first two sons of the women will be healthy.
2. Given that the first two sons are healthy, compute the updated probability that she is a carrier.
3. Given that the first two sons are healthy, compute the conditional probability that the third son will be healthy.

Answer

1. $\frac{5}{8}$
2. $\frac{1}{5}$
3. $\frac{9}{10}$

Laplace's Rule of Succession

Suppose that we have $m + 1$ coins, labeled $0, 1, \dots, m$. Coin i lands heads with probability $\frac{i}{m}$ for each i . The experiment is to choose a coin at random (so that each coin is equally likely to be chosen) and then toss the chosen coin repeatedly.

1. The probability that the first n tosses are all heads is $p_{m,n} = \frac{1}{m+1} \sum_{i=0}^m \left(\frac{i}{m}\right)^n$
2. $p_{m,n} \rightarrow \frac{1}{n+1}$ as $m \rightarrow \infty$
3. The conditional probability that toss $n + 1$ is heads given that the previous n tosses were all heads is $\frac{p_{m,n+1}}{p_{m,n}}$
4. $\frac{p_{m,n+1}}{p_{m,n}} \rightarrow \frac{n+1}{n+2}$ as $m \rightarrow \infty$

Proof

Part (a) follows by conditioning on the chosen coin. For part (b), note that $p_{m,n}$ is an approximating sum for $\int_0^1 x^n dx = \frac{1}{n+1}$. Part (c) follows from the definition of conditional probability, and part (d) is a trivial consequence of (b), (c).

Note that coin 0 is two-tailed, the probability of heads increases with i , and coin m is two-headed. The limiting conditional probability in part (d) is called *Laplace's Rule of Succession*, named after Simon Laplace. This rule was used by Laplace and others as a general principle for estimating the conditional probability that an event will occur on time $n + 1$, given that the event has occurred n times in succession.

Suppose that a missile has had 10 successful tests in a row. Compute Laplace's estimate that the 11th test will be successful. Does this make sense?

Answer

$\frac{11}{12}$. No, not really.

This page titled [2.5: Independence](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.6: Convergence

This is the first of several sections in this chapter that are more advanced than the basic topics in the first five sections. In this section we discuss several topics related to convergence of events and random variables, a subject of fundamental importance in probability theory. In particular the results that we obtain will be important for:

- Properties of distribution functions,
- The weak law of large numbers,
- The strong law of large numbers.

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) .

Basic Theory

Sequences of events

Our first discussion deals with sequences of events and various types of limits of such sequences. The limits are also event. We start with two simple definitions.

Suppose that (A_1, A_2, \dots) is a sequence of events.

1. The sequence is *increasing* if $A_n \subseteq A_{n+1}$ for every $n \in \mathbb{N}_+$.
2. The sequence is *decreasing* if $A_{n+1} \subseteq A_n$ for every $n \in \mathbb{N}_+$.

Note that these are the standard definitions of increasing and decreasing, relative to the ordinary total order \leq on the index set \mathbb{N}_+ and the subset partial order \subseteq on the collection of events. The terminology is also justified by the corresponding indicator variables.

Suppose that (A_1, A_2, \dots) is a sequence of events, and let $I_n = \mathbf{1}_{A_n}$ denote the indicator variable of the event A_n for $n \in \mathbb{N}_+$.

1. The sequence of events is increasing if and only if the sequence of indicator variables is increasing in the ordinary sense. That is, $I_n \leq I_{n+1}$ for each $n \in \mathbb{N}_+$.
2. The sequence of events is decreasing if and only if the sequence of indicator variables is decreasing in the ordinary sense. That is, $I_{n+1} \leq I_n$ for each $n \in \mathbb{N}_+$.

Proof




Figure 2.6.1: A sequence of increasing events and their union




Figure 2.6.2: A sequence of decreasing events and their intersection

If a sequence of events is either increasing or decreasing, we can define the *limit* of the sequence in a way that turns out to be quite natural.

Suppose that (A_1, A_2, \dots) is a sequence of events.

1. If the sequence is increasing, we define $\lim_{n \rightarrow \infty} A_n = \bigcup_{n=1}^{\infty} A_n$.
2. If the sequence is decreasing, we define $\lim_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} A_n$.

Once again, the terminology is clarified by the corresponding indicator variables.

Suppose again that (A_1, A_2, \dots) is a sequence of events, and let $I_n = \mathbf{1}_{A_n}$ denote the indicator variable of A_n for $n \in \mathbb{N}_+$.

1. If the sequence of events is increasing, then $\lim_{n \rightarrow \infty} I_n$ is the indicator variable of $\bigcup_{n=1}^{\infty} A_n$
2. If the sequence of events is decreasing, then $\lim_{n \rightarrow \infty} I_n$ is the indicator variable of $\bigcap_{n=1}^{\infty} A_n$

Proof

1. If $s \in \bigcup_{n=1}^{\infty} A_n$ then $s \in A_k$ for some $k \in \mathbb{N}_+$. Since the events are increasing, $s \in A_n$ for every $n \geq k$. In this case, $I_n(s) = 1$ for every $n \geq k$ and hence $\lim_{n \rightarrow \infty} I_n(s) = 1$. On the other hand, if $s \notin \bigcup_{n=1}^{\infty} A_n$ then $s \notin A_n$ for every $n \in \mathbb{N}_+$. In this case, $I_n(s) = 0$ for every $n \in \mathbb{N}_+$ and hence $\lim_{n \rightarrow \infty} I_n(s) = 0$.
2. If $s \in \bigcap_{n=1}^{\infty} A_n$ then $s \in A_n$ for each $n \in \mathbb{N}_+$. In this case, $I_n(s) = 1$ for each $n \in \mathbb{N}_+$ and hence $\lim_{n \rightarrow \infty} I_n(s) = 1$. If $s \notin \bigcap_{n=1}^{\infty} A_n$ then $s \notin A_k$ for some $k \in \mathbb{N}_+$. Since the events are decreasing, $s \notin A_n$ for all $n \geq k$. In this case, $I_n(s) = 0$ for $n \geq k$ and hence $\lim_{n \rightarrow \infty} I_n(s) = 0$.

An arbitrary union of events can always be written as a union of increasing events, and an arbitrary intersection of events can always be written as an intersection of decreasing events:

Suppose that (A_1, A_2, \dots) is a sequence of events. Then

1. $\bigcup_{i=1}^n A_i$ is increasing in $n \in \mathbb{N}_+$ and $\bigcup_{i=1}^{\infty} A_i = \lim_{n \rightarrow \infty} \bigcup_{i=1}^n A_i$.
2. $\bigcap_{i=1}^n A_i$ is decreasing in $n \in \mathbb{N}_+$ and $\bigcap_{i=1}^{\infty} A_i = \lim_{n \rightarrow \infty} \bigcap_{i=1}^n A_i$.

Proof

1. Trivially $\bigcup_{i=1}^n A_i \subseteq \bigcup_{i=1}^{n+1} A_i$. The second statement simply means that $\bigcup_{n=1}^{\infty} \bigcup_{i=1}^n A_i = \bigcup_{i=1}^{\infty} A_i$.
2. Trivially $\bigcap_{i=1}^{n+1} A_i \subseteq \bigcap_{i=1}^n A_i$. The second statement simply means that $\bigcap_{n=1}^{\infty} \bigcap_{i=1}^n A_i = \bigcap_{i=1}^{\infty} A_i$.

There is a more interesting and useful way to generate increasing and decreasing sequences from an arbitrary sequence of events, using the *tail* segment of the sequence rather than the *initial* segment.

Suppose that (A_1, A_2, \dots) is a sequence of events. Then

1. $\bigcup_{i=n}^{\infty} A_i$ is decreasing in $n \in \mathbb{N}_+$.
2. $\bigcap_{i=n}^{\infty} A_i$ is increasing in $n \in \mathbb{N}_+$.

Proof

1. Clearly $\bigcup_{i=n+1}^{\infty} A_i \subseteq \bigcup_{i=n}^{\infty} A_i$
2. Clearly $\bigcap_{i=n}^{\infty} A_i \subseteq \bigcap_{i=n+1}^{\infty} A_i$

Since the new sequences defined in the previous results are decreasing and increasing, respectively, we can take their limits. These are the *limit superior* and *limit inferior*, respectively, of the original sequence.

Suppose that (A_1, A_2, \dots) is a sequence of events. Define

1. $\limsup_{n \rightarrow \infty} A_n = \lim_{n \rightarrow \infty} \bigcup_{i=n}^{\infty} A_i = \bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} A_i$. This is the event that occurs if and only if A_n occurs for infinitely many values of n .
2. $\liminf_{n \rightarrow \infty} A_n = \lim_{n \rightarrow \infty} \bigcap_{i=n}^{\infty} A_i = \bigcup_{n=1}^{\infty} \bigcap_{i=n}^{\infty} A_i$. This is the event that occurs if and only if A_n occurs for all but finitely many values of n .

Proof

1. From the definition, the event $\limsup_{n \rightarrow \infty} A_n$ occurs if and only if for each $n \in \mathbb{N}_+$ there exists $i \geq n$ such that A_i occurs.
2. From the definition, the event $\liminf_{n \rightarrow \infty} A_n$ occurs if and only if there exists $n \in \mathbb{N}_+$ such that A_i occurs for every $i \geq n$.

Once again, the terminology and notation are clarified by the corresponding indicator variables. You may need to review limit inferior and limit superior for sequences of real numbers in the section on Partial Orders.

Suppose that (A_1, A_2, \dots) is a sequence of events, and let $I_n = \mathbf{1}_{A_n}$ denote the indicator variable of A_n for $n \in \mathbb{N}_+$. Then

1. $\limsup_{n \rightarrow \infty} I_n$ is the indicator variable of $\limsup_{n \rightarrow \infty} A_n$.
2. $\liminf_{n \rightarrow \infty} I_n$ is the indicator variable of $\liminf_{n \rightarrow \infty} A_n$.

Proof

1. By the result above, $\lim_{n \rightarrow \infty} \mathbf{1}(\bigcup_{i=n}^{\infty} A_i)$ is the indicator variable of $\limsup_{n \rightarrow \infty} A_n$. But $\mathbf{1}(\bigcup_{i=n}^{\infty} A_i) = \max\{I_i : i \geq n\}$ and hence $\lim_{n \rightarrow \infty} \mathbf{1}(\bigcup_{i=n}^{\infty} A_i) = \limsup_{n \rightarrow \infty} I_n$.
2. By the result above, $\lim_{n \rightarrow \infty} \mathbf{1}(\bigcap_{i=n}^{\infty} A_i)$ is the indicator variable of $\liminf_{n \rightarrow \infty} A_n$. But $\mathbf{1}(\bigcap_{i=n}^{\infty} A_i) = \min\{I_i : i \geq n\}$ and hence $\lim_{n \rightarrow \infty} \mathbf{1}(\bigcap_{i=n}^{\infty} A_i) = \liminf_{n \rightarrow \infty} I_n$.

Suppose that (A_1, A_2, \dots) is a sequence of events. Then $\liminf_{n \rightarrow \infty} A_n \subseteq \limsup_{n \rightarrow \infty} A_n$.

Proof

If A_n occurs for all but finitely many $n \in \mathbb{N}_+$ then certainly A_n occurs for infinitely many $n \in \mathbb{N}_+$.

Suppose that (A_1, A_2, \dots) is a sequence of events. Then

1. $(\limsup_{n \rightarrow \infty} A_n)^c = \liminf_{n \rightarrow \infty} A_n^c$
2. $(\liminf_{n \rightarrow \infty} A_n)^c = \limsup_{n \rightarrow \infty} A_n^c$.

Proof

These results follow from DeMorgan's laws.

The Continuity Theorems

Generally speaking, a function is *continuous* if it preserves limits. Thus, the following results are the *continuity theorems of probability*. Part (a) is the continuity theorem for increasing events and part (b) the continuity theorem for decreasing events.

Suppose that (A_1, A_2, \dots) is a sequence of events.

1. If the sequence is increasing then $\lim_{n \rightarrow \infty} \mathbb{P}(A_n) = \mathbb{P}(\lim_{n \rightarrow \infty} A_n) = \mathbb{P}(\bigcup_{n=1}^{\infty} A_n)$
2. If the sequence is decreasing then $\lim_{n \rightarrow \infty} \mathbb{P}(A_n) = \mathbb{P}(\lim_{n \rightarrow \infty} A_n) = \mathbb{P}(\bigcap_{n=1}^{\infty} A_n)$


Proof

1. Let $B_1 = A_1$ and let $B_i = A_i \setminus A_{i-1}$ for $i \in \{2, 3, \dots\}$. Note that the collection of events $\{B_1, B_2, \dots\}$ is pairwise disjoint and has the same union as $\{A_1, A_2, \dots\}$. From countable additivity and the definition of infinite series,

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \mathbb{P}\left(\bigcup_{i=1}^{\infty} B_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(B_i) = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{P}(B_i) \quad (2.6.1)$$

But $\mathbb{P}(B_1) = \mathbb{P}(A_1)$ and $\mathbb{P}(B_i) = \mathbb{P}(A_i) - \mathbb{P}(A_{i-1})$ for $i \in \{2, 3, \dots\}$. Therefore $\sum_{i=1}^n \mathbb{P}(B_i) = \mathbb{P}(A_n)$ and hence we have $\mathbb{P}(\bigcup_{i=1}^{\infty} A_i) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n)$.

The construction in the continuity theorem for increasing events

 The construction in the continuity theorem

2. The sequence of complements (A_1^c, A_2^c, \dots) is increasing. Hence using part (a), DeMorgan's law, and the complement rule we have

$$\mathbb{P}\left(\bigcap_{i=1}^{\infty} A_i\right) = 1 - \mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i^c\right) = 1 - \lim_{n \rightarrow \infty} \mathbb{P}(A_n^c) = \lim_{n \rightarrow \infty} [1 - \mathbb{P}(A_n^c)] = \lim_{n \rightarrow \infty} \mathbb{P}(A_n) \quad (2.6.2)$$

The continuity theorems can be applied to the increasing and decreasing sequences that we constructed earlier from an arbitrary sequence of events.

Suppose that (A_1, A_2, \dots) is a sequence of events.

1. $\mathbb{P}(\bigcup_{i=1}^{\infty} A_i) = \lim_{n \rightarrow \infty} \mathbb{P}(\bigcup_{i=1}^n A_i)$
2. $\mathbb{P}(\bigcap_{i=1}^{\infty} A_i) = \lim_{n \rightarrow \infty} \mathbb{P}(\bigcap_{i=1}^n A_i)$

Proof

These results follow immediately from the [continuity theorems](#).

Suppose that (A_1, A_2, \dots) is a sequence of events. Then

1. $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(\bigcup_{i=n}^{\infty} A_i)$
2. $\mathbb{P}(\liminf_{n \rightarrow \infty} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(\bigcap_{i=n}^{\infty} A_i)$

Proof

These results follows directly from the [definitions](#), and the continuity theorems.

The next result shows that the countable additivity axiom for a probability measure is equivalent to *finite additivity* and the continuity property for increasing events.

Temporarily, suppose that \mathbb{P} is only finitely additive, but satisfies the [continuity property for increasing events](#). Then \mathbb{P} is countably additive.

Proof

Suppose that (A_1, A_2, \dots) is a sequence of pairwise disjoint events. Since we are assuming that \mathbb{P} is finitely additive we have

$$\mathbb{P}\left(\bigcup_{i=1}^n A_i\right) = \sum_{i=1}^n \mathbb{P}(A_i) \quad (2.6.3)$$

If we let $n \rightarrow \infty$, the left side converges to $\mathbb{P}(\bigcup_{i=1}^{\infty} A_i)$ by the continuity assumption and the result [above](#), while the right side converges to $\sum_{i=1}^{\infty} \mathbb{P}(A_i)$ by the definition of an infinite series.

There are a few mathematicians who reject the countable additivity axiom of probability measure in favor of the weaker finite additivity axiom. Whatever the philosophical arguments may be, life is certainly much harder without the continuity theorems.

The Borel-Cantelli Lemmas

The *Borel-Cantelli Lemmas*, named after Emil Borel and Francesco Cantelli, are very important tools in probability theory. The first lemma gives a condition that is sufficient to conclude that infinitely many events occur with probability 0.

First Borel-Cantelli Lemma. Suppose that (A_1, A_2, \dots) is a sequence of events. If $\sum_{n=1}^{\infty} \mathbb{P}(A_n) < \infty$ then $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = 0$.

Proof

From the result above on [limit superiors](#), we have $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(\bigcup_{i=n}^{\infty} A_i)$. But from Boole's inequality, $\mathbb{P}(\bigcup_{i=n}^{\infty} A_i) \leq \sum_{i=n}^{\infty} \mathbb{P}(A_i)$. Since $\sum_{i=1}^{\infty} \mathbb{P}(A_i) < \infty$, we have $\sum_{i=n}^{\infty} \mathbb{P}(A_i) \rightarrow 0$ as $n \rightarrow \infty$.

The second lemma gives a condition that is sufficient to conclude that infinitely many *independent* events occur with probability 1.

Second Borel-Cantelli Lemma. Suppose that (A_1, A_2, \dots) is a sequence of independent events. If $\sum_{n=1}^{\infty} \mathbb{P}(A_n) = \infty$ then $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = 1$.

Proof

Note first that $1 - x \leq e^{-x}$ for every $x \in \mathbb{R}$, and hence $1 - \mathbb{P}(A_i) \leq \exp[-\mathbb{P}(A_i)]$ for each $i \in \mathbb{N}_+$. From the results above on [limit superiors](#) and [complements](#),

$$\mathbb{P}\left[\left(\limsup_{n \rightarrow \infty} A_n\right)^c\right] = \mathbb{P}\left(\liminf_{n \rightarrow \infty} A_n^c\right) = \lim_{n \rightarrow \infty} \mathbb{P}\left(\bigcap_{i=n}^{\infty} A_i^c\right) \quad (2.6.4)$$

But by independence and the inequality above,

$$\mathbb{P}\left(\bigcap_{i=n}^{\infty} A_i^c\right) = \prod_{i=n}^{\infty} \mathbb{P}(A_i^c) = \prod_{i=n}^{\infty} [1 - \mathbb{P}(A_i)] \leq \prod_{i=n}^{\infty} \exp[-\mathbb{P}(A_i)] = \exp\left(-\sum_{i=n}^{\infty} \mathbb{P}(A_i)\right) = 0 \quad (2.6.5)$$

For independent events, both Borel-Cantelli lemmas apply of course, and lead to a *zero-one law*.

If (A_1, A_2, \dots) is a sequence of independent events then $\limsup_{n \rightarrow \infty} A_n$ has probability 0 or 1:

1. If $\sum_{n=1}^{\infty} \mathbb{P}(A_n) < \infty$ then $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = 0$.
2. If $\sum_{n=1}^{\infty} \mathbb{P}(A_n) = \infty$ then $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = 1$.

This result is actually a special case of a more general zero-one law, known as the *Kolmogorov zero-one law*, and named for Andrei Kolmogorov. This law is studied in the more advanced section on measure. Also, we can use the zero-one law to derive a calculus theorem that relates infinite series and infinite products. This derivation is an example of the *probabilistic method*—the use of probability to obtain results, seemingly unrelated to probability, in other areas of mathematics.

Suppose that $p_i \in (0, 1)$ for each $i \in \mathbb{N}_+$. Then

$$\prod_{i=1}^{\infty} p_i > 0 \text{ if and only if } \sum_{i=1}^{\infty} (1 - p_i) < \infty \quad (2.6.6)$$

Proof

We can easily construct a probability space with a sequence of independent events (A_1, A_2, \dots) such that $\mathbb{P}(A_i) = 1 - p_i$ for each $i \in \mathbb{N}_+$. The result then follows from the proofs of the two Borel-Cantelli lemmas.

Our next result is a simple application of the second Borel-Cantelli lemma to independent replications of a basic experiment.

Suppose that A is an event in a basic random experiment with $\mathbb{P}(A) > 0$. In the compound experiment that consists of independent replications of the basic experiment, the event “ A occurs infinitely often” has probability 1.

Proof

Let p denote the probability of A in the basic experiment. In the compound experiment, we have a sequence of independent events (A_1, A_2, \dots) with $\mathbb{P}(A_n) = p$ for each $n \in \mathbb{N}_+$ (these are “independent copies” of A). But $\sum_{n=1}^{\infty} \mathbb{P}(A_n) = \infty$ since $p > 0$ so the result follows from the second Borel-Cantelli lemma.

Convergence of Random Variables

Our next discussion concerns two ways that a sequence of random variables defined for our experiment can “converge”. These are fundamentally important concepts, since some of the deepest results in probability theory are limit theorems involving random variables. The most important special case is when the random variables are real valued, but the proofs are essentially the same for variables with values in a metric space, so we will use the more general setting.

Thus, suppose that (S, d) is a metric space, and that \mathcal{S} is the corresponding Borel σ -algebra (that is, the σ -algebra generated by the topology), so that our measurable space is (S, \mathcal{S}) . Here is the most important special case:

For $n \in \mathbb{N}_+$, is the n -dimensional *Euclidean space* is (\mathbb{R}^n, d_n) where

$$d_n(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_{i=1}^n (y_i - x_i)^2}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n), \mathbf{y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n \quad (2.6.7)$$

Euclidean spaces are named for Euclid, of course. As noted above, the one-dimensional case where $d(x, y) = |y - x|$ for $x, y \in \mathbb{R}$ is particularly important. Returning to the general metric space, recall that if (x_1, x_2, \dots) is a sequence in S and $x \in S$, then $x_n \rightarrow x$ as $n \rightarrow \infty$ means that $d(x_n, x) \rightarrow 0$ as $n \rightarrow \infty$ (in the usual calculus sense). For the rest of our discussion, we assume that (X_1, X_2, \dots) is a sequence of random variable with values in S and X is a random variable with values in S , all defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

We say that $X_n \rightarrow X$ as $n \rightarrow \infty$ *with probability 1* if the event that $X_n \rightarrow X$ as $n \rightarrow \infty$ has probability 1. That is,

$$\mathbb{P}\{\omega \in S : X_n(\omega) \rightarrow X(\omega) \text{ as } n \rightarrow \infty\} = 1 \quad (2.6.8)$$

Details

We need to make sure that the definition makes sense, in that the statement that X_n converges to X as $n \rightarrow \infty$ defines a valid event. Note that X_n does *not* converge to X as $n \rightarrow \infty$ if and only if for some $\epsilon > 0$, $d(X_n, X) > \epsilon$ for infinitely many $n \in \mathbb{N}_+$. Note that if this condition holds for a given $\epsilon > 0$ then it holds for all smaller $\epsilon > 0$. Moreover, there are arbitrarily small rational $\epsilon > 0$ so X_n does not converge to X as $n \rightarrow \infty$ if and only if for some rational $\epsilon > 0$, $d(X_n, X) > \epsilon$ for infinitely many $n \in \mathbb{N}_+$. Hence

$$\{X_n \rightarrow X \text{ as } n \rightarrow \infty\}^c = \bigcup_{\epsilon \in \mathbb{Q}_+} \limsup_{n \rightarrow \infty} \{d(X_n, X) > \epsilon\} \quad (2.6.9)$$

where \mathbb{Q}_+ is the set of positive rational numbers. A critical point to remember is that this set is countable. So, building a little at a time, note that $\{d(X_n, X) > \epsilon\}$ is an event for each $\epsilon \in \mathbb{Q}_+$ and $n \in \mathbb{N}_+$ since X_n and X are random variables. Next, the limit superior of a sequence of events is an event. Finally, a countable union of events is an event.

As good probabilists, we usually suppress references to the sample space and write the definition simply as $\mathbb{P}(X_n \rightarrow X \text{ as } n \rightarrow \infty) = 1$. The statement that an event has probability 1 is usually the strongest affirmative statement that we can make in probability theory. Thus, convergence with probability 1 is the strongest form of convergence. The phrases *almost surely* and *almost everywhere* are sometimes used instead of the phrase *with probability 1*.

Recall that metrics d and e on S are *equivalent* if they generate the same topology on S . Recall also that convergence of a sequence is a *topological property*. That is, if (x_1, x_2, \dots) is a sequence in S and $x \in S$, and if d, e are equivalent metrics on S , then $x_n \rightarrow x$ as $n \rightarrow \infty$ relative to d if and only if $x_n \rightarrow x$ as $n \rightarrow \infty$ relative to e . So for our random variables as defined above, it follows that $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1 relative to d if and only if $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1 relative to e .

The following statements are equivalent:

1. $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1.
2. $\mathbb{P}[d(X_n, X) > \epsilon \text{ for infinitely many } n \in \mathbb{N}_+] = 0$ for every rational $\epsilon > 0$.
3. $\mathbb{P}[d(X_n, X) > \epsilon \text{ for infinitely many } n \in \mathbb{N}_+] = 0$ for every $\epsilon > 0$.
4. $\mathbb{P}[d(X_k, X) > \epsilon \text{ for some } k \geq n] \rightarrow 0$ as $n \rightarrow \infty$ for every $\epsilon > 0$.

Proof

From the details in the [definition above](#), $\mathbb{P}(X_n \rightarrow X \text{ as } n \rightarrow \infty) = 1$ if and only if

$$\mathbb{P}\left(\bigcup_{\epsilon \in \mathbb{Q}_+} \{d(X_n, X) > \epsilon \text{ for infinitely many } n \in \mathbb{N}_+\}\right) = 0 \quad (2.6.10)$$

where again \mathbb{Q}_+ is the set of positive rational numbers. But by Boole's inequality, a countable union of events has probability 0 if and only if every event in the union has probability 0. Thus, (a) is equivalent to (b). Statement (b) is clearly equivalent to (c) since there are arbitrarily small positive rational numbers. Finally, (c) is equivalent to (d) by the continuity result in [above](#).

Our next result gives a fundamental criterion for convergence with probability 1:

If $\sum_{n=1}^{\infty} \mathbb{P}[d(X_n, X) > \epsilon] < \infty$ for every $\epsilon > 0$ then $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1.

Proof

By the [first Borel-Cantelli lemma](#), if $\sum_{n=1}^{\infty} \mathbb{P}[d(X_n, X) > \epsilon] < \infty$ then $\mathbb{P}[d(X_n, X) > \epsilon \text{ for infinitely many } n \in \mathbb{N}_+] = 0$. Hence the result follows from the previous theorem.

Here is our next mode of convergence.

We say that $X_n \rightarrow X$ as $n \rightarrow \infty$ *in probability* if

$$\mathbb{P}[d(X_n, X) > \epsilon] \rightarrow 0 \text{ as } n \rightarrow \infty \text{ for each } \epsilon > 0 \quad (2.6.11)$$

The phrase *in probability* sounds superficially like the phrase *with probability 1*. However, as we will soon see, convergence in probability is much weaker than convergence with probability 1. Indeed, convergence with probability 1 is often called *strong*

convergence, while convergence in probability is often called *weak convergence*.

If $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1 then $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability.

Proof

Let $\epsilon > 0$. Then $\mathbb{P}[d(X_n, X) > \epsilon] \leq \mathbb{P}[d(X_k, X) > \epsilon \text{ for some } k \geq n]$. But if $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1, then the expression on the right converges to 0 as $n \rightarrow \infty$ by part (d) of the result [above](#). Hence $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability.

The converse fails with a passion. A simple counterexample is given [below](#). However, there is a partial converse that is very useful.

If $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability, then there exists a subsequence $(n_1, n_2, n_3 \dots)$ of \mathbb{N}_+ such that $X_{n_k} \rightarrow X$ as $k \rightarrow \infty$ with probability 1.

Proof

Suppose that $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability. Then for each $k \in \mathbb{N}_+$ there exists $n_k \in \mathbb{N}_+$ such that $\mathbb{P}[d(X_{n_k}, X) > 1/k] < 1/k^2$. We can make the choices so that $n_k < n_{k+1}$ for each k . It follows that $\sum_{k=1}^{\infty} \mathbb{P}[d(X_{n_k}, X) > \epsilon] < \infty$ for every $\epsilon > 0$. By the result [above](#), $X_{n_k} \rightarrow X$ as $n \rightarrow \infty$ with probability 1.

Note that the proof works because $1/k \rightarrow 0$ as $k \rightarrow \infty$ and $\sum_{k=1}^{\infty} 1/k^2 < \infty$. Any two sequences with these properties would work just as well.

There are two other modes of convergence that we will discuss later:

- Convergence in distribution.
- Convergence in mean,

Examples and Applications

Coins

Suppose that we have an infinite sequence of coins labeled $1, 2, \dots$. Moreover, coin n has probability of heads $1/n^a$ for each $n \in \mathbb{N}_+$, where $a > 0$ is a parameter. We toss each coin in sequence one time. In terms of a , find the probability of the following events:

1. infinitely many heads occur
2. infinitely many tails occur

Answer

Let H_n be the event that toss n results in heads, and T_n the event that toss n results in tails.

1. $\mathbb{P}(\limsup_{n \rightarrow \infty} H_n) = 1$, $\mathbb{P}(\limsup_{n \rightarrow \infty} T_n) = 1$ if $a \in (0, 1]$
2. $\mathbb{P}(\limsup_{n \rightarrow \infty} H_n) = 0$, $\mathbb{P}(\limsup_{n \rightarrow \infty} T_n) = 1$ if $a \in (1, \infty)$

The following exercise gives a simple example of a sequence of random variables that converge in probability but not with probability 1. Naturally, we are assuming the standard metric on \mathbb{R} .

Suppose again that we have a sequence of coins labeled $1, 2, \dots$, and that coin n lands heads up with probability $\frac{1}{n}$ for each n . We toss the coins in order to produce a sequence (X_1, X_2, \dots) of independent indicator random variables with

$$\mathbb{P}(X_n = 1) = \frac{1}{n}, \quad \mathbb{P}(X_n = 0) = 1 - \frac{1}{n}; \quad n \in \mathbb{N}_+ \quad (2.6.12)$$

1. $\mathbb{P}(X_n = 0 \text{ for infinitely many } n) = 1$, so that infinitely many tails occur with probability 1.
2. $\mathbb{P}(X_n = 1 \text{ for infinitely many } n) = 1$, so that infinitely many heads occur with probability 1.
3. $\mathbb{P}(X_n \text{ does not converge as } n \rightarrow \infty) = 1$.
4. $X_n \rightarrow 0$ as $n \rightarrow \infty$ in probability.

Proof

1. This follows from the second Borel-Cantelli lemma, since $\sum_{n=1}^{\infty} \mathbb{P}(X_n = 0) = \infty$

2. This also follows from the second Borel-Cantelli lemma, since $\sum_{n=1}^{\infty} \mathbb{P}(X_n = 1) = \infty$.
3. This follows from parts (a) and (b). Recall that the intersection of two events with probability 1 still has probability 1.
4. Suppose $0 < \epsilon < 1$. Then $\mathbb{P}(|X_n - 0| > \epsilon) = \mathbb{P}(X_n = 1) = \frac{1}{n} \rightarrow 0$ as $n \rightarrow \infty$.

Discrete Spaces

Recall that a measurable space (S, \mathcal{S}) is *discrete* if S is countable and \mathcal{S} is the collection of all subsets of S (the power set of S). Moreover, \mathcal{S} is the Borel σ -algebra corresponding to the *discrete metric* d on S given by $d(x, x) = 0$ for $x \in S$ and $d(x, y) = 1$ for distinct $x, y \in S$. How do convergence with probability 1 and convergence in probability work for the discrete metric?

Suppose that (S, \mathcal{S}) is a discrete space. Suppose further that (X_1, X_2, \dots) is a sequence of random variables with values in S and X is a random variable with values in S , all defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Relative to the discrete metric d ,

1. $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1 if and only if $\mathbb{P}(X_n = X \text{ for all but finitely many } n \in \mathbb{N}_+) = 1$.
2. $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability if and only if $\mathbb{P}(X_n \neq X) \rightarrow 0$ as $n \rightarrow \infty$.

Proof

1. If (x_1, x_2, \dots) is a sequence of points in S and $x \in S$, then relative to metric d , $x_n \rightarrow x$ as $n \rightarrow \infty$ if and only if $x_n = x$ for all but finitely many $n \in \mathbb{N}_+$.
2. If $\epsilon \geq 1$ then $\mathbb{P}[d(X_n, X) > \epsilon] = 0$. If $\epsilon \in (0, 1)$ then $\mathbb{P}[d(X_n, X) > \epsilon] = \mathbb{P}(X_n \neq X)$.

Of course, it's important to realize that a discrete space can be the Borel space for metrics other than the discrete metric.

This page titled [2.6: Convergence](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.7: Measure Spaces

In this section we discuss positive measure spaces (which include probability spaces) from a more advanced point of view. The sections on Measure Theory and Special Set Structures in the chapter on Foundations are essential prerequisites. On the other hand, if you are not interested in the measure-theoretic aspects of probability, you can safely skip this section.

Positive Measure

Definitions

Suppose that S is a set, playing the role of a universal set for a mathematical theory. As we have noted before, S usually comes with a σ -algebra \mathcal{S} of admissible subsets of S , so that (S, \mathcal{S}) is a *measurable space*. In particular, this is the case for the model of a random experiment, where S is the set of outcomes and \mathcal{S} the σ -algebra of events, so that the measurable space (S, \mathcal{S}) is the *sample space* of the experiment. A probability measure is a special case of a more general object known as a positive measure.

A *positive measure* on (S, \mathcal{S}) is a function $\mu : \mathcal{S} \rightarrow [0, \infty]$ that satisfies the following axioms:

1. $\mu(\emptyset) = 0$
2. If $\{A_i : i \in I\}$ is a countable, pairwise disjoint collection of sets in \mathcal{S} then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i) \quad (2.7.1)$$

The triple (S, \mathcal{S}, μ) is a *measure space*.

Axiom (b) is called *countable additivity*, and is the essential property. The measure of a set that consists of a countable union of disjoint pieces is the sum of the measures of the pieces. Note also that since the terms in the sum are positive, there is no issue with the order of the terms in the sum, although of course, ∞ is a possible value.




Figure 2.7.1: A union of four disjoint sets

So perhaps the term *measurable space* for (S, \mathcal{S}) makes a little more sense now—a measurable space is one that can have a positive measure defined on it.

Suppose that (S, \mathcal{S}, μ) is a measure space.

1. If $\mu(S) < \infty$ then (S, \mathcal{S}, μ) is a *finite measure space*.
2. If $\mu(S) = 1$ then (S, \mathcal{S}, μ) is a *probability space*.

So probability measures are positive measures, but positive measures are important beyond the application to probability. The standard measures on the Euclidean spaces are all positive measures: the extension of *length* for measurable subsets of \mathbb{R} , the extension of *area* for measurable subsets of \mathbb{R}^2 , the extension of *volume* for measurable subsets of \mathbb{R}^3 , and the higher dimensional analogues. We will actually construct these measures in the next section on Existence and Uniqueness. In addition, counting measure $\#$ is a positive measure on the subsets of a set S . Even more general measures that can take positive and negative values are explored in the chapter on Distributions.

Properties

The following results give some simple properties of a positive measure space (S, \mathcal{S}, μ) . The proofs are essentially identical to the proofs of the corresponding properties of probability, except that the measure of a set may be infinite so we must be careful to avoid the dreaded indeterminate form $\infty - \infty$.

If $A, B \in \mathcal{S}$, then $\mu(B) = \mu(A \cap B) + \mu(B \setminus A)$.

Proof

Note that $B = (A \cap B) \cup (B \setminus A)$, and the sets in the union are disjoint.

If $A, B \in \mathcal{S}$ and $A \subseteq B$ then

1. $\mu(B) = \mu(A) + \mu(B \setminus A)$
2. $\mu(A) \leq \mu(B)$

Proof

Part (a) follows from the previous theorem, since $A \cap B = A$. Part (b) follows from part (a).

Thus μ is an *increasing function*, relative to the subset partial order \subseteq on \mathcal{S} and the ordinary order \leq on $[0, \infty]$. In particular, if μ is a finite measure, then $\mu(A) < \infty$ for every $A \in \mathcal{S}$. Note also that if $A, B \in \mathcal{S}$ and $\mu(B) < \infty$ then $\mu(B \setminus A) = \mu(B) - \mu(A \cap B)$. In the special case that $A \subseteq B$, this becomes $\mu(B \setminus A) = \mu(B) - \mu(A)$. In particular, these results holds for a finite measure and are just like the difference rules for probability. If μ is a finite measure, then $\mu(A^c) = \mu(S) - \mu(A)$. This is the analogue of the complement rule in probability, with but with $\mu(S)$ replacing 1.

The following result is the analogue of Boole's inequality for probability. For a general positive measure, the result is referred to as the *subadditive property*.

Suppose that $A_i \in \mathcal{S}$ for i in a countable index set I . Then

$$\mu\left(\bigcup_{i \in I} A_i\right) \leq \sum_{i \in I} \mu(A_i) \quad (2.7.2)$$

Proof

The proof is exactly like the one for Boole's inequality. Assume that $I = \mathbb{N}_+$. Let $B_1 = A_1$ and $B_i = A_i \setminus (A_1 \cup \dots \cup A_{i-1})$ for $i \in \{2, 3, \dots\}$. Then $\{B_i : i \in I\}$ is a disjoint collection of sets in \mathcal{S} with the same union as $\{A_i : i \in I\}$. Also $B_i \subseteq A_i$ for each i so $\mu(B_i) \leq \mu(A_i)$. Hence

$$\mu\left(\bigcup_{i \in I} A_i\right) = \mu\left(\bigcup_{i \in I} B_i\right) = \sum_{i \in I} \mu(B_i) \leq \sum_{i \in I} \mu(A_i) \quad (2.7.3)$$

For a union of sets with finite measure, the inclusion-exclusion formula holds, and the proof is just like the one for probability.

Suppose that $A_i \in \mathcal{S}$ for each $i \in I$ where $\#(I) = n$, and that $\mu(A_i) < \infty$ for $i \in I$. Then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mu\left(\bigcap_{j \in J} A_j\right) \quad (2.7.4)$$

Proof

The proof is by induction on n . The proof for $n = 2$ is simple: $A_1 \cup A_2 = A_1 \cup (A_2 \setminus A_1)$. The union on the right is disjoint, so using additivity and the difference rule,

$$\mu(A_1 \cup A_2) = \mu(A_1) + \mu(A_2 \setminus A_1) = \mu(A_1) + \mu(A_2) - \mu(A_1 \cap A_2) \quad (2.7.5)$$

Suppose now that the inclusion-exclusion formula holds for a given $n \in \mathbb{N}_+$, and consider the case $n + 1$. Then

$$\bigcup_{i=1}^{n+1} A_i = \left(\bigcup_{i=1}^n A_i\right) \cup \left[A_{n+1} \setminus \left(\bigcup_{i=1}^n A_i\right)\right] \quad (2.7.6)$$

As before, the set in parentheses and the set in square brackets are disjoint. Thus using the additivity axiom, the difference rule, and the distributive rule we have

$$\mu\left(\bigcup_{i=1}^{n+1} A_i\right) = \mu\left(\bigcup_{i=1}^n A_i\right) + \mu(A_{n+1}) - \mu\left(\bigcup_{i=1}^n (A_{n+1} \cap A_i)\right) \quad (2.7.7)$$

By the induction hypothesis, the inclusion-exclusion formula holds for each union of n sets on the right. Applying the formula and simplifying gives the inclusion-exclusion formula for $n + 1$ sets.

The continuity theorem for increasing sets holds for a positive measure. The continuity theorem for decreasing events holds also, if the sets have finite measure. Again, the proofs are similar to the ones for a probability measure, except for considerations of infinite measure.

Suppose that (A_1, A_2, \dots) is a sequence of sets in \mathcal{S} .

1. If the sequence is increasing then $\mu(\bigcup_{i=1}^{\infty} A_i) = \lim_{n \rightarrow \infty} \mu(A_n)$.
2. If sequence is decreasing and $\mu(A_1) < \infty$ then $\mu(\bigcap_{i=1}^{\infty} A_i) = \lim_{n \rightarrow \infty} \mu(A_n)$.

Proof

1. Note that if $\mu(A_k) = \infty$ for some k then $\mu(A_n) = \infty$ for $n \geq k$ and $\mu(\bigcup_{i=1}^{\infty} A_i) = \infty$. Thus, suppose that $\mu(A_i) < \infty$ for each i . Let $B_1 = A_1$ and $B_i = A_i \setminus A_{i-1}$ for $i \in \{2, 3, \dots\}$. Then (B_1, B_2, \dots) is a disjoint sequence with the same union as (A_1, A_2, \dots) . Also, $\mu(B_1) = \mu(A_1)$ and by the [proper difference rule](#), $\mu(B_i) = \mu(A_i) - \mu(A_{i-1})$ for $i \in \{2, 3, \dots\}$. Hence

$$\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \mu\left(\bigcup_{i=1}^{\infty} B_i\right) = \sum_{i=1}^{\infty} \mu(B_i) = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mu(B_i) \quad (2.7.8)$$

But $\sum_{i=1}^n \mu(B_i) = \mu(A_1) + \sum_{i=2}^n [\mu(A_i) - \mu(A_{i-1})] = \mu(A_n)$.

2. Note that $A_1 \setminus A_n$ is increasing in n . Hence using the continuity result for increasing sets,

$$\mu\left(\bigcap_{i=1}^{\infty} A_i\right) = \mu\left[A_1 \setminus \bigcup_{i=1}^{\infty} (A_1 \setminus A_i)\right] = \mu(A_1) - \mu\left[\bigcup_{i=1}^{\infty} (A_1 \setminus A_n)\right] \quad (2.7.9)$$

$$= \mu(A_1) - \lim_{n \rightarrow \infty} \mu(A_1 \setminus A_n) = \mu(A_1) - \lim_{n \rightarrow \infty} [\mu(A_1) - \mu(A_n)] = \lim_{n \rightarrow \infty} \mu(A_n) \quad (2.7.10)$$

Recall that if (A_1, A_2, \dots) is increasing, $\bigcup_{i=1}^{\infty} A_i$ is denoted $\lim_{n \rightarrow \infty} A_n$, and if (A_1, A_2, \dots) is decreasing, $\bigcap_{i=1}^{\infty} A_i$ is denoted $\lim_{n \rightarrow \infty} A_n$. In both cases, the continuity theorem has the form $\mu(\lim_{n \rightarrow \infty} A_n) = \lim_{n \rightarrow \infty} \mu(A_n)$. The continuity theorem for decreasing events fails without the additional assumption of finite measure. A simple counterexample is given [below](#).

The following corollary of the inclusion-exclusion law gives a condition for countable additivity that does not require that the sets be disjoint, but only that the intersections have measure 0. The result is used below in the [theorem on completion](#).

Suppose that $A_i \in \mathcal{S}$ for each i in a countable index set I and that $\mu(A_i) < \infty$ for $i \in I$ and $\mu(A_i \cap A_j) = 0$ for distinct $i, j \in I$. Then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i) \quad (2.7.11)$$

Proof

We will assume that $I = \mathbb{N}_+$. For $n \in \mathbb{N}_+$,

$$\mu\left(\bigcup_{i=1}^n A_i\right) = \sum_{i=1}^n \mu(A_i) \quad (2.7.12)$$

as an immediate consequence of the inclusion-exclusion law, under the assumption that $\mu(A_i \cap A_j) = 0$ for distinct $i, j \in \{1, 2, \dots, n\}$. Next $\bigcup_{i=1}^n A_i \uparrow \bigcup_{i=1}^{\infty} A_i$ as $n \rightarrow \infty$, and hence by the continuity theorem for increasing events, $\mu(\bigcup_{i=1}^n A_i) \rightarrow \mu(\bigcup_{i=1}^{\infty} A_i)$ as $n \rightarrow \infty$. On the other hand, $\sum_{i=1}^n \mu(A_i) \rightarrow \sum_{i=1}^{\infty} \mu(A_i)$ as $n \rightarrow \infty$ by the definition of an infinite series of nonnegative terms.

More Definitions

If a positive measure is not finite, then the following definition gives the next best thing.

The measure space (S, \mathcal{S}, μ) is σ -finite if there exists a countable collection $\{A_i : i \in I\} \subseteq \mathcal{S}$ with $\bigcup_{i \in I} A_i = S$ and $\mu(A_i) < \infty$ for each $i \in I$.

So of course, if μ is a finite measure on (S, \mathcal{S}) then μ is σ -finite, but not conversely in general. On the other hand, for $i \in I$, let $\mathcal{S}_i = \{A \in \mathcal{S} : A \subseteq A_i\}$. Then \mathcal{S}_i is a σ -algebra of subsets of A_i and μ restricted to \mathcal{S}_i is a finite measure. The point of this (and the reason for the definition) is that often nice properties of finite measures can be extended to σ -finite measures. In particular, σ -finite measure spaces play a crucial role in the construction of product measure spaces, and for the [completion](#) of a measure space considered below.

Suppose that (S, \mathcal{S}, μ) is a σ -finite measure space.

1. There exists an increasing sequence satisfying the σ -finite definition
2. There exists a disjoint sequence satisfying the σ -finite definition.

Proof

Without loss of generality, we can take \mathbb{N}_+ as the index set in the definition. So there exists $A_n \in \mathcal{S}$ for $n \in \mathbb{N}_+$ such that $\mu(A_n) < \infty$ for each $n \in \mathbb{N}_+$ and $S = \bigcup_{n=1}^{\infty} A_n$. The proof uses some of the same tricks that we have seen before.

1. Let $B_n = \bigcup_{i=1}^n A_i$. Then $B_n \in \mathcal{S}$ for $n \in \mathbb{N}_+$ and this sequence is increasing. Moreover, $\mu(B_n) \leq \sum_{i=1}^n \mu(A_i) < \infty$ for $n \in \mathbb{N}_+$ and $\bigcup_{n=1}^{\infty} B_n = \bigcup_{n=1}^{\infty} A_n = S$.
2. Let $C_1 = A_1$ and let $C_n = A_n \setminus \bigcup_{i=1}^{n-1} A_i$ for $n \in \{2, 3, \dots\}$. Then $C_n \in \mathcal{S}$ for each $n \in \mathbb{N}_+$ and this sequence is disjoint. Moreover, $C_n \subseteq A_n$ so $\mu(C_n) \leq \mu(A_n) < \infty$ and $\bigcup_{n=1}^{\infty} C_n = \bigcup_{n=1}^{\infty} A_n = S$.

Our next definition concerns sets where a measure is concentrated, in a certain sense.

Suppose that (S, \mathcal{S}, μ) is a measure space. An *atom* of the space is a set $A \in \mathcal{S}$ with the following properties:

1. $\mu(A) > 0$
2. If $B \in \mathcal{S}$ and $B \subseteq A$ then either $\mu(B) = \mu(A)$ or $\mu(B) = 0$.

A measure space that has no atoms is called *non-atomic* or *diffuse*.

In probability theory, we are often particularly interested in atoms that are singleton sets. Note that $\{x\} \in \mathcal{S}$ is an atom if and only if $\mu(\{x\}) > 0$, since the only subsets of $\{x\}$ are $\{x\}$ itself and \emptyset .

Constructions

There are several simple ways to construct new positive measures from existing ones. As usual, we start with a measurable space (S, \mathcal{S}) .

Suppose that (R, \mathcal{R}) is a measurable subspace of (S, \mathcal{S}) . If μ is a positive measure on (S, \mathcal{S}) then μ restricted to \mathcal{R} is a positive measure on (R, \mathcal{R}) . If μ is a finite measure on (S, \mathcal{S}) then μ is a finite measure on (R, \mathcal{R}) .

Proof

The assumption is that \mathcal{R} is a σ -algebra of subsets of R and $\mathcal{R} \subseteq \mathcal{S}$. In particular $R \in \mathcal{S}$. Since the additivity property of μ holds for a countable, disjoint collection of events in \mathcal{S} , it trivially holds for a countable, disjoint collection of events in \mathcal{R} . Finally, by the increasing property, $\mu(R) \leq \mu(S)$ so if $\mu(S) < \infty$ then $\mu(R) < \infty$.

However, if μ is σ -finite on (S, \mathcal{S}) , it is *not* necessarily true that μ is σ -finite on (R, \mathcal{R}) . A counterexample is given [below](#). The previous theorem would apply, in particular, when $R = S$ so that \mathcal{R} is a sub σ -algebra of \mathcal{S} . Next, a positive multiple of a positive measure gives another positive measure.

If μ is a positive measure on (S, \mathcal{S}) and $c \in (0, \infty)$, then $c\mu$ is also a positive measure on (S, \mathcal{S}) . If μ is finite (σ -finite) then $c\mu$ is finite (σ -finite) respectively.

Proof

Clearly $c\mu : \mathcal{S} \rightarrow [0, \infty]$. Also $(c\mu)(\emptyset) = c\mu(\emptyset) = 0$. Next if $\{A_i : i \in I\}$ is a countable, disjoint collection of events in \mathcal{S} then

$$(c\mu)\left(\bigcup_{i \in I} A_i\right) = c\mu\left(\bigcup_{i \in I} A_i\right) = c \sum_{i \in I} \mu(A_i) = \sum_{i \in I} c\mu(A_i) \quad (2.7.13)$$

Finally, since $\mu(A) < \infty$ if and only if $(c\mu)(A) < \infty$ for $A \in \mathcal{S}$, the finiteness and σ -finiteness properties are trivially preserved.

A nontrivial finite positive measure μ is practically just like a probability measure, and in fact can be re-scaled into a probability measure \mathbb{P} , as was done in the section on Probability Measures:

Suppose that μ is a positive measure on (S, \mathcal{S}) with $0 < \mu(S) < \infty$. Then \mathbb{P} defined by $\mathbb{P}(A) = \mu(A)/\mu(S)$ for $A \in \mathcal{S}$ is a probability measure on (S, \mathcal{S}) .

Proof

\mathbb{P} is a measure by the previous result, and trivially $\mathbb{P}(S) = 1$.

Sums of positive measures are also positive measures.

If μ_i is a positive measure on (S, \mathcal{S}) for each i in a countable index set I then $\mu = \sum_{i \in I} \mu_i$ is also a positive measure on (S, \mathcal{S}) .

1. If I is finite and μ_i is finite for each $i \in I$ then μ is finite.
2. If I is finite and μ_i is σ -finite for each $i \in I$ then μ is σ -finite.

Proof

Clearly $\mu : \mathcal{S} \rightarrow [0, \infty]$. First $\mu(\emptyset) = \sum_{i \in I} \mu_i(\emptyset) = 0$. Next if $\{A_j : j \in J\}$ is a countable, disjoint collection of events in \mathcal{S} then

$$\mu\left(\bigcup_{j \in J} A_j\right) = \sum_{i \in I} \mu_i\left(\bigcup_{j \in J} A_j\right) = \sum_{i \in I} \sum_{j \in J} \mu_i(A_j) = \sum_{j \in J} \sum_{i \in I} \mu_i(A_j) = \sum_{j \in J} \mu(A_j) \quad (2.7.14)$$

The interchange of sums is permissible since the terms are nonnegative. Suppose now that I is finite.

1. If μ_i is finite for each $i \in I$ then $\mu(S) = \sum_{i \in I} \mu_i(S) < \infty$ so μ is finite.
2. Suppose that μ_i is σ -finite for each $i \in I$. Then for each $i \in I$ there exists a collection $\mathcal{A}_i = \{A_{ij} : j \in \mathbb{N}\} \subseteq \mathcal{S}$ such that $\bigcup_{j=1}^{\infty} A_{ij} = S$ and $\mu_i(A_{ij}) < \infty$ for each $j \in \mathbb{N}$. For $j \in \mathbb{N}$, let $B_j = \bigcap_{i \in I} A_{ij}$. Then $B_j \in \mathcal{S}$ for each $j \in \mathbb{N}$ and

$$\bigcup_{j=1}^{\infty} B_j = \bigcup_{j=1}^{\infty} \bigcap_{i \in I} A_{ij} = \bigcap_{i \in I} \bigcup_{j=1}^{\infty} A_{ij} = \bigcap_{i \in I} S = S \quad (2.7.15)$$

Moreover,

$$\mu(B_j) = \sum_{i \in I} \mu_i(B_j) \leq \sum_{i \in I} \mu_i(A_{ij}) < \infty \quad (2.7.16)$$

so μ is σ -finite.

In the context of the last result, if I is countably infinite and μ_i is finite for each $i \in I$, then μ is not necessarily σ -finite. A counterexample is given below. In this case, μ is said to be *s-finite*, but we've had enough definitions, so we won't pursue this one. From [scaling](#) and [sum](#) properties, note that a positive linear combination of positive measures is a positive measure. The next method is sometimes referred to as a *change of variables*.

Suppose that (S, \mathcal{S}, μ) is a measure space. Suppose also that (T, \mathcal{T}) is another measurable space and that $f : S \rightarrow T$ is measurable. Then ν defined as follows is a positive measure on (T, \mathcal{T})

$$\nu(B) = \mu[f^{-1}(B)], \quad B \in \mathcal{T} \quad (2.7.17)$$

If μ is finite then ν is finite.

Proof

Clearly $\nu : \mathcal{T} \rightarrow [0, \infty]$. The proof is easy since inverse images preserve all set operations. First $f^{-1}(\emptyset) = \emptyset$ so $\nu(\emptyset) = 0$. Next, if $\{B_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{T} , then $\{f^{-1}(B_i) : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} , and $f^{-1}(\bigcup_{i \in I} B_i) = \bigcup_{i \in I} f^{-1}(B_i)$. Hence

$$\nu\left(\bigcup_{i \in I} B_i\right) = \mu\left[f^{-1}\left(\bigcup_{i \in I} B_i\right)\right] = \mu\left[\bigcup_{i \in I} f^{-1}(B_i)\right] = \sum_{i \in I} \mu[f^{-1}(B_i)] = \sum_{i \in I} \nu(B_i) \quad (2.7.18)$$

Finally, if μ is finite then $\nu(T) = \mu[f^{-1}(T)] = \mu(S) < \infty$ so ν is finite.

In the context of the last result, if μ is σ -finite on (S, \mathcal{S}) , it is not necessarily true that ν is σ -finite on (T, \mathcal{T}) , even if f is one-to-one. A counterexample is given below. The takeaway is that σ -finiteness of ν depends very much on the nature of the σ -algebra \mathcal{T} . Our next result shows that it's easy to explicitly construct a positive measure on a *countably generated σ -algebra*, that is, a σ -algebra generated by a countable partition. Such σ -algebras are important for counterexamples and to gain insight, and also because many σ -algebras that occur in applications can be constructed from them.

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of S into nonempty sets, and that $\mathcal{S} = \sigma(\mathcal{A})$, the σ -algebra generated by the partition. For $i \in I$, define $\mu(A_i) \in [0, \infty]$ arbitrarily. For $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$, define

$$\mu(A) = \sum_{j \in J} \mu(A_j) \quad (2.7.19)$$

Then μ is a positive measure on (S, \mathcal{S}) .

1. The atoms of the measure are the sets of the form $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$ and where $\mu(A_j) > 0$ for one and only one $j \in J$.
2. If $\mu(A_i) < \infty$ for $i \in I$ and I is finite then μ is finite.
3. If $\mu(A_i) < \infty$ for $i \in I$ and I is countably infinite then μ is σ -finite.

Proof

Recall that every $A \in \mathcal{S}$ has a unique representation of the form $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$. In particular, $J = \emptyset$ in this representation gives $A = \emptyset$. The sum over an empty index set is 0, so $\mu(\emptyset) = 0$. Next suppose that $\{B_k : k \in K\}$ is a countable, disjoint collection of sets in \mathcal{S} . Then there exists a disjoint collection $\{J_k : k \in K\}$ of subsets of I such that $B_k = \bigcup_{j \in J_k} A_j$. Hence

$$\mu\left(\bigcup_{k \in K} B_k\right) = \mu\left(\bigcup_{k \in K} \bigcup_{j \in J_k} A_j\right) = \sum_{k \in K} \sum_{j \in J_k} \mu(A_j) = \sum_{k \in K} \mu(B_k) \quad (2.7.20)$$

The fact that the terms are all nonnegative means that we do not have to worry about the order of summation.

1. Again, every $A \in \mathcal{S}$ has the unique representation $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$. The subsets of A that are in \mathcal{S} are $\bigcup_{k \in K} A_k$ where $K \subseteq J$. Hence A is an atom if and only if $\mu(A_j) > 0$ for one and only one $j \in J$.
2. If I is finite and $\mu(A_i) < \infty$ then $\mu(S) = \sum_{i \in I} \mu(A_i) < \infty$, so μ is finite.
3. If I is countably infinite and $\mu(A_i) < \infty$ for $i \in I$ then \mathcal{A} satisfies the condition for μ to be σ -finite.

One of the most general ways to construct new measures from old ones is via the theory of integration with respect to a positive measure, which is explored in the chapter on Distributions. The construction of positive measures more or less “from scratch” is considered in the next section on Existence and Uniqueness. We close this discussion with a simple result that is useful for counterexamples.

Suppose that the measure space (S, \mathcal{S}, μ) has an atom $A \in \mathcal{S}$ with $\mu(A) = \infty$. Then the space is not σ -finite.

Proof

Let $\{A_i : i \in I\}$ be a countable disjoint collection of sets in \mathcal{S} that partitions S . Then $\{A \cap A_i : i \in I\}$ partitions A . Since $\mu(A) = \sum_{i \in I} \mu(A \cap A_i)$, we must have $\mu(A \cap A_i) > 0$ for some $i \in I$. Since A is an atom and $A \cap A_i \subseteq A$ it follows that $\mu(A \cap A_i) = \infty$. Hence also therefore $\mu(A_i) = \infty$.

Measure and Topology

Often the spaces that occur in probability and stochastic processes are topological spaces. Recall that a topological space (S, \mathcal{T}) consists of a set S and a topology \mathcal{T} on S (the collection of open sets). The topology as well as the measure theory plays an important role, so it's natural to want these two types of structures to be compatible. We have already seen the most important step in this direction: Recall that $\mathcal{S} = \sigma(\mathcal{T})$, the σ -algebra generated by the topology, is the *Borel σ -algebra* on S , named for Émile Borel.

Since the complement of an open set is a closed set, \mathcal{S} is also the σ -algebra generated by the collection of closed sets. Moreover, \mathcal{S} contains countable intersections of open sets (called G_δ sets) and countable unions of closed sets (called F_σ sets).

Suppose that (S, \mathcal{T}) is a topological space and let $\mathcal{S} = \sigma(\mathcal{T})$ be the Borel σ -algebra. A positive measure μ on (S, \mathcal{S}) is a *Borel measure* and then (S, \mathcal{S}, μ) is a *Borel measure space*.

The next definition concerns the subset on which a Borel measure is concentrated, in a certain sense.

Suppose that (S, \mathcal{S}, μ) is a Borel measure space. The *support* of μ is

$$\text{supp}(\mu) = \{x \in S : \mu(U) > 0 \text{ for every open neighborhood } U \text{ of } x\} \quad (2.7.21)$$

The set $\text{supp}(\mu)$ is closed.

Proof

Let $A = \text{supp}(\mu)^c$. For $x \in A^c$, there exists an open neighborhood V_x of x such that $\mu(V_x) = 0$. If $y \in V_x$, then V_x is also an open neighborhood of y , so $y \in A^c$. Hence $V_x \subseteq A^c$ for every $x \in A^c$ and so A^c is open.

The term *Borel measure* has different definitions in the literature. Often the topological space is required to be locally compact, Hausdorff, and with a countable base (*LCCB*). Then a Borel measure μ is required to have the additional condition that $\mu(C) < \infty$ if $C \subseteq S$ is compact. In this text, we use the term Borel measures in this more restricted sense.

Suppose that (S, \mathcal{S}, μ) is a Borel measure space corresponding to an LCCB topology. Then the space is σ -finite.

Proof

Since the topological space is locally compact and has a countable base, $S = \bigcup_{i \in I} C_i$ where $\{C_i : i \in I\}$ is a countable collection of compact sets. Since μ is a Borel measure, $\mu(C_i) < \infty$ and hence μ is σ -finite.

Here are a couple of other definitions that are important for Borel measures, again linking topology and measure in natural ways.

Suppose again that (S, \mathcal{S}, μ) is a Borel measure space.

1. μ is *inner regular* if $\mu(A) = \sup\{\mu(C) : C \text{ is compact and } C \subseteq A\}$ for $A \in \mathcal{S}$.
2. μ is *outer regular* if $\mu(A) = \inf\{\mu(U) : U \text{ is open and } A \subseteq U\}$ for $A \in \mathcal{S}$.
3. μ is *regular* if it is both inner regular and outer regular.

The measure spaces that occur in probability and stochastic processes are usually regular Borel spaces associated with LCCB topologies.

Null Sets and Equivalence

Sets of measure 0 in a measure space turn out to be very important precisely because we can often ignore the differences between mathematical objects on such sets. In this discussion, we assume that we have a fixed measure space (S, \mathcal{S}, μ) .

A set $A \in \mathcal{S}$ is *null* if $\mu(A) = 0$.

Consider a measurable “statement” with $x \in S$ as a free variable. (Technically, such a statement is a *predicate* on S .) If the statement is true for all $x \in S$ except for x in a null set, we say that the statement holds *almost everywhere* on S . This terminology is used often in measure theory and captures the importance of the definition.

Let $\mathcal{D} = \{A \in \mathcal{S} : \mu(A) = 0 \text{ or } \mu(A^c) = 0\}$, the collection of *null and co-null* sets. Then \mathcal{D} is a sub σ -algebra of \mathcal{S} .

Proof

Trivially $S \in \mathcal{D}$ since $S^c = \emptyset$ and $\mu(\emptyset) = 0$. Next if $A \in \mathcal{D}$ then $A^c \in \mathcal{D}$ by the symmetry of the definition. Finally, suppose that $A_i \in \mathcal{D}$ for $i \in I$ where I is a countable index set. If $\mu(A_i) = 0$ for every $i \in I$ then $\mu(\bigcup_{i \in I} A_i) \leq \sum_{i \in I} \mu(A_i) = 0$ by the [subadditive property](#). On the other hand, if $\mu(A_j^c) = 0$ for some $j \in J$ then $\mu[(\bigcup_{i \in I} A_i)^c] = \mu(\bigcap_{i \in I} A_i^c) \leq \mu(A_j^c) = 0$. In either case, $\bigcup_{i \in I} A_i \in \mathcal{D}$.

Of course μ restricted to \mathcal{S} is not very interesting since $\mu(A) = 0$ or $\mu(A) = \mu(S)$ for every $A \in \mathcal{S}$. Our next definition is a type of equivalence between sets in \mathcal{S} . To make this precise, recall first that the *symmetric difference* between subsets A and B of S is $A \triangle B = (A \setminus B) \cup (B \setminus A)$. This is the set that consists of points in one of the two sets, but not both, and corresponds to *exclusive or*.

Sets $A, B \in \mathcal{S}$ are *equivalent* if $\mu(A \triangle B) = 0$, and we denote this by $A \equiv B$.

Thus $A \equiv B$ if and only if $\mu(A \triangle B) = \mu(A \setminus B) + \mu(B \setminus A) = 0$ if and only if $\mu(A \setminus B) = \mu(B \setminus A) = 0$. In the predicate terminology mentioned above, the statement

$$x \in A \text{ if and only if } x \in B \quad (2.7.22)$$

is true for almost every $x \in S$. As the name suggests, the relation \equiv really is an equivalence relation on \mathcal{S} and hence \mathcal{S} is partitioned into disjoint classes of mutually equivalent sets. Two sets in the same equivalence class differ by a set of measure 0.

The relation \equiv is an equivalence relation on \mathcal{S} . That is, for $A, B, C \in \mathcal{S}$,

1. $A \equiv A$ (the *reflexive property*).
2. If $A \equiv B$ then $B \equiv A$ (the *symmetric property*).
3. If $A \equiv B$ and $B \equiv C$ then $A \equiv C$ (the *transitive property*).

Proof

1. The reflexive property is trivial since $A \triangle A = \emptyset$.
2. The symmetric property is also trivial since $A \triangle B = B \triangle A$.
3. For the transitive property, suppose that $A \equiv B$ and $B \equiv C$. Note that $A \setminus C \subseteq (A \setminus B) \cup (B \setminus C)$, and hence $\mathbb{P}(A \setminus C) = 0$. By a symmetric argument, $\mathbb{P}(C \setminus A) = 0$.

Equivalence is preserved under the standard set operations.

If $A, B \in \mathcal{S}$ and $A \equiv B$ then $A^c \equiv B^c$.

Proof

Note that $A^c \setminus B^c = B \setminus A$ and $B^c \setminus A^c = A \setminus B$, so $A^c \triangle B^c = A \triangle B$.

Suppose that $A_i, B_i \in \mathcal{S}$ and that $A_i \equiv B_i$ for i in a countable index set I . Then

1. $\bigcup_{i \in I} A_i \equiv \bigcup_{i \in I} B_i$
2. $\bigcap_{i \in I} A_i \equiv \bigcap_{i \in I} B_i$

Proof

1. Note that

$$\left(\bigcup_{i \in I} A_i \right) \triangle \left(\bigcup_{i \in I} B_i \right) \subseteq \bigcup_{i \in I} (A_i \triangle B_i) \quad (2.7.23)$$

To see this, note that if x is in the set on the left then either $x \in A_j$ for some $j \in I$ and $x \notin B_i$ for every $i \in I$, or $x \notin A_i$ for every $i \in I$ and $x \in B_j$ for some $j \in I$. In either case, $x \in A_j \triangle B_j$ for some $j \in I$.

2. Similarly

$$\left(\bigcap_{i \in I} A_i \right) \triangle \left(\bigcap_{i \in I} B_i \right) \subseteq \bigcup_{i \in I} (A_i \triangle B_i) \quad (2.7.24)$$

If x is in the set on the left then $x \in A_i$ for every $i \in I$ and $x \notin B_j$ for some $j \in I$, or $x \in B_i$ for every $i \in I$ and $x \notin A_j$ for some $j \in I$. In either case, $x \in A_j \triangle B_j$ for some $j \in I$.

In both parts, the proof is completed by noting that the common set on the right in the displayed equations is null:

$$\mu \left[\bigcup_{i \in I} (A_i \triangle B_i) \right] \leq \sum_{i \in I} \mu(A_i \triangle B_i) = 0 \quad (2.7.25)$$

Equivalent sets have the same measure.

If $A, B \in \mathcal{S}$ and $A \equiv B$ then $\mu(A) = \mu(B)$.

Proof

Note again that $A = (A \cap B) \cup (A \setminus B)$. If $A \equiv B$ then $\mu(A) = \mu(A \cap B)$. By a symmetric argument, $\mu(B) = \mu(A \cap B)$.

The converse trivially fails, and a counterexample is given [below](#). However, the collection of null sets and the collection of co-null sets do form equivalence classes.

Suppose that $A \in \mathcal{S}$.

1. If $\mu(A) = 0$ then $A \equiv B$ if and only if $\mu(B) = 0$.
2. If $\mu(A^c) = 0$ then $A \equiv B$ if and only if $\mu(B^c) = 0$.

Proof

1. Suppose that $\mu(A) = 0$ and $A \equiv B$. Then $\mu(B) = 0$ by the result [above](#). Conversely, note that $A \setminus B \subseteq A$ and $B \setminus A \subseteq B$ so if $\mu(A) = \mu(B) = 0$ then $\mu(A \triangle B) = 0$ so $A \equiv B$.
2. Part (b) follows from part (a) and the result above on [complements](#).

We can extend the notion of equivalence to measurable functions with a common range space. Thus suppose that (T, \mathcal{T}) is another measurable space. If $f, g: S \rightarrow T$ are measurable, then $(f, g): S \rightarrow T \times T$ is measurable with respect to the usual product σ -algebra $\mathcal{T} \otimes \mathcal{T}$. We assume that the diagonal set $D = \{(y, y) : y \in T\} \in \mathcal{T} \otimes \mathcal{T}$, which is almost always true in applications.

Measurable functions $f, g: S \rightarrow T$ are equivalent if $\mu\{x \in S : f(x) \neq g(x)\} = 0$. Again we write $f \equiv g$.

Details

Note that $\{x \in S : f(x) \neq g(x)\} = \{x \in S : (f(x), g(x)) \in D^c\} \in \mathcal{S}$ by our assumption, so the definition makes sense.

In the terminology discussed earlier, $f \equiv g$ means that $f(x) = g(x)$ almost everywhere on S . As with measurable sets, the relation \equiv really does define an equivalence relation on the collection of measurable functions from S to T . Thus, the collection of such functions is partitioned into disjoint classes of mutually equivalent variables.

The relation \equiv is an equivalence relation on the collection of measurable functions from S to T . That is, for measurable $f, g, h: S \rightarrow T$,

1. $f \equiv f$ (the reflexive property).
2. If $f \equiv g$ then $g \equiv f$ (the symmetric property).
3. If $f \equiv g$ and $g \equiv h$ then $f \equiv h$ (the transitive property).

Proof

Parts (a) and (b) are trivially. For (c) note that $f(x) = g(x)$ and $g(x) = h(x)$ implies $f(x) = h(x)$ for $x \in S$. Negating this statement gives $f(x) \neq h(x)$ implies $f(x) \neq g(x)$ or $g(x) \neq h(x)$. So

$$\{x \in S : f(x) \neq h(x)\} \subseteq \{x \in S : f(x) \neq g(x)\} \cup \{x \in S : g(x) \neq h(x)\} \quad (2.7.26)$$

Since $f \equiv g$ and $g \equiv h$, the two sets on the right have measure 0. Hence, so does the set on the left.

Suppose again that $f, g: S \rightarrow T$ are measurable and that $f \equiv g$. Then for every $B \in \mathcal{T}$, the sets $f^{-1}(B) \equiv g^{-1}(B)$.

Proof

Note that $f^{-1}(B) \triangle g^{-1}(B) \subseteq \{x \in S : f(x) \neq g(x)\}$.

Thus if $f, g: S \rightarrow T$ are measurable and $f \equiv g$, then by the previous result, $\nu_f = \nu_g$ where ν_f, ν_g are the measures on (T, \mathcal{T}) associated with f and g , as [above](#). Again, the converse fails with a passion.

It often happens that a definition for functions subsumes the corresponding definition for sets, by considering the indicator functions of the sets. So it is with equivalence. In the following result, we can take $T = \{0, 1\}$ with \mathcal{T} the collection of all subsets.

Suppose that $A, B \in \mathcal{S}$. Then $A \equiv B$ if and only if $\mathbf{1}_A \equiv \mathbf{1}_B$.

Proof

Note that $\{x \in S : 1_A(x) \neq 1_B(x)\} = A \triangle B$.

Equivalence is preserved under composition. For the next result, suppose that (U, \mathcal{U}) is yet another measurable space.

Suppose that $f, g : S \rightarrow T$ are measurable and that $h : T \rightarrow U$ is measurable. If $f \equiv g$ then $h \circ f \equiv h \circ g$.

Proof

Note that $\{x \in S : h[f(x)] \neq h[g(x)]\} \subseteq \{x \in S : f(x) \neq g(x)\}$.

Suppose again that (S, \mathcal{S}, μ) is a measure space. Let \mathcal{V} denote the collection of all measurable real-valued random functions from S into \mathbb{R} . (As usual, \mathbb{R} is given the Borel σ -algebra.) From our previous discussion of measure theory, we know that with the usual definitions of addition and scalar multiplication, $(\mathcal{V}, +, \cdot)$ is a vector space. However, in measure theory, we often do not want to distinguish between functions that are equivalent, so it's nice to know that the vector space structure is preserved when we *identify* equivalent functions. Formally, let $[f]$ denote the equivalence class generated by $f \in \mathcal{V}$, and let \mathcal{W} denote the collection of all such equivalence classes. In modular notation, \mathcal{W} is \mathcal{V}/\equiv . We define addition and scalar multiplication on \mathcal{W} by

$$[f] + [g] = [f + g], \quad c[f] = [cf]; \quad f, g \in \mathcal{V}, \quad c \in \mathbb{R} \quad (2.7.27)$$

$(\mathcal{W}, +, \cdot)$ is a vector space.

Proof

All that we have to show is that addition and scalar multiplication are well defined. That is, we must show that the definitions do not depend on the particular representative of the equivalence class. Then the other properties that define a vector space are inherited from $(\mathcal{V}, +, \cdot)$. Thus we must show that if $f_1 \equiv f_2$ and $g_1 \equiv g_2$, and if $c \in \mathbb{R}$, then $f_1 + g_1 \equiv f_2 + g_2$ and $cf_1 \equiv cf_2$. For the first problem, note that (f_1, g_1) and (f_2, g_2) are measurable functions from S to \mathbb{R}^2 . (\mathbb{R}^2 is given the product σ -algebra which also happens to be the Borel σ -algebra corresponding to the standard Euclidean topology). Moreover, $(f_1, g_1) \equiv (f_2, g_2)$ since

$$\{x \in S : (f_1(x), g_1(x)) \neq (f_2(x), g_2(x))\} = \{x \in S : f_1(x) \neq f_2(x)\} \cup \{x \in S : g_1(x) \neq g_2(x)\} \quad (2.7.28)$$

But the function $(a, b) \mapsto a + b$ from \mathbb{R}^2 into \mathbb{R} is measurable and hence from [composition property](#), it follows that $f_1 + g_1 \equiv f_2 + g_2$. The second problem is easier. The function $a \mapsto ca$ from \mathbb{R} into \mathbb{R} is measurable so again it follows from [composition property](#) that $cf_1 \equiv cf_2$.

Often we don't bother to use the special notation for the equivalence class associated with a function. Rather, it's understood that equivalent functions represent the same object. Spaces of functions in a measure space are studied further in the chapter on Distributions.

Completion

Suppose that (S, \mathcal{S}, μ) is a measure space and let $\mathcal{N} = \{A \in \mathcal{S} : \mu(A) = 0\}$ denote the collection of null sets of the space. If $A \in \mathcal{N}$ and $B \in \mathcal{S}$ is a subset of A , then we know that $\mu(B) = 0$ so $B \in \mathcal{N}$ also. However, in general there might be subsets of A that are not in \mathcal{S} . This leads naturally to the following definition.

The measure space (S, \mathcal{S}, μ) is *complete* if $A \in \mathcal{N}$ and $B \subseteq A$ imply $B \in \mathcal{S}$ (and hence $B \in \mathcal{N}$).

Our goal in this discussion is to show that if (S, \mathcal{S}, μ) is a σ -finite measure that is not complete, then it can be *completed*. That is μ can be extended to σ -algebra that includes all of the sets in \mathcal{S} and all subsets of null sets. The first step is to extend the equivalence relation defined in our previous discussion to $\mathcal{P}(S)$.

For $A, B \subseteq S$, define $A \equiv B$ if and only if there exists $N \in \mathcal{N}$ such that $A \triangle B \subseteq N$. The relation \equiv is an equivalence relation on $\mathcal{P}(S)$: For $A, B, C \subseteq S$,

1. $A \equiv A$ (the *reflexive property*).
2. If $A \equiv B$ then $B \equiv A$ (the *symmetric property*).
3. If $A \equiv B$ and $B \equiv C$ then $A \equiv C$ (the *transitive property*).

Proof

1. Note that $A \triangle A = \emptyset$ and $\emptyset \in \mathcal{N}$.
2. Suppose that $A \triangle B \subseteq N$ where $N \in \mathcal{N}$. Then $B \triangle A = A \triangle B \subseteq N$.
3. Suppose that $A \triangle B \subseteq N_1$ and $B \triangle C \subseteq N_2$ where $N_1, N_2 \in \mathcal{N}$. Then $A \triangle C \subseteq (A \triangle B) \cup (B \triangle C) \subseteq N_1 \cup N_2$, and $N_1 \cup N_2 \in \mathcal{N}$.

So the equivalence relation \equiv partitions $\mathcal{P}(S)$ into mutually disjoint equivalence classes. Two sets in an equivalence class differ by a subset of a null set. In particular, $A \equiv \emptyset$ if and only if $A \subseteq N$ for some $N \in \mathcal{N}$. The extended relation \equiv is preserved under the set operations, just as before. Our next step is to enlarge the σ -algebra \mathcal{S} by adding any set that is equivalent to a set in \mathcal{S} .

Let $\mathcal{S}_0 = \{A \subseteq S : A \equiv B \text{ for some } B \in \mathcal{S}\}$. Then \mathcal{S}_0 is a σ -algebra of subsets of S , and in fact is the σ -algebra generated by $\mathcal{S} \cup \{A \subseteq S : A \equiv \emptyset\}$.

Proof

Note that if $A \in \mathcal{S}$ then $A \equiv A$ so $A \in \mathcal{S}_0$. In particular, $S \in \mathcal{S}_0$. Also, $\emptyset \in \mathcal{S}$ so if $A \equiv \emptyset$ then $A \in \mathcal{S}_0$. Suppose that $A \in \mathcal{S}_0$ so that $A \equiv B$ for some $B \in \mathcal{S}$. Then $B^c \in \mathcal{S}$ and $A^c \equiv B^c$ so $A^c \in \mathcal{S}_0$. Next suppose that $A_i \in \mathcal{S}_0$ for i in a countable index set I . Then for each $i \in I$ there exists $B_i \in \mathcal{S}$ such that $A_i \equiv B_i$. But then $\bigcup_{i \in I} B_i \in \mathcal{S}$ and $\bigcup_{i \in I} A_i \equiv \bigcup_{i \in I} B_i$, so $\bigcup_{i \in I} A_i \in \mathcal{S}_0$. Therefore \mathcal{S}_0 is a σ -algebra of subsets of S . Finally, suppose that \mathcal{T} is a σ -algebra of subsets of S and that $\mathcal{S} \cup \{A \subseteq S : A \equiv \emptyset\} \subseteq \mathcal{T}$. We need to show that $\mathcal{S}_0 \subseteq \mathcal{T}$. Thus, suppose that $A \in \mathcal{S}_0$. Then there exists $B \in \mathcal{S}$ such that $A \equiv B$. But $B \in \mathcal{T}$ and $A \triangle B \in \mathcal{T}$ so $A \cap B = B \setminus (A \triangle B) \in \mathcal{T}$. Also $A \setminus B \in \mathcal{T}$, so $A = (A \cap B) \cup (A \setminus B) \in \mathcal{T}$.

Our last step is to extend μ to a positive measure on the enlarged σ -algebra \mathcal{S}_0 .

Suppose that $A \in \mathcal{S}_0$ so that $A \equiv B$ for some $B \in \mathcal{S}$. Define $\mu_0(A) = \mu(B)$. Then

1. μ_0 is well defined.
2. $\mu_0(A) = \mu(A)$ for $A \in \mathcal{S}$.
3. μ_0 is a positive measure on \mathcal{S}_0 .

The measure space $(S, \mathcal{S}_0, \mu_0)$ is complete and is known as the *completion* of (S, \mathcal{S}, μ) .

Proof

1. Suppose that $A \in \mathcal{S}_0$ and that $A \equiv B_1$ and $A \equiv B_2$ where $B_1, B_2 \in \mathcal{S}$. Then $B_1 \equiv B_2$ so by the result [above](#) $\mu(B_1) = \mu(B_2)$. Thus, μ_0 is well-defined.
2. Next, if $A \in \mathcal{S}$ then of course $A \equiv A$ so $\mu_0(A) = \mu(A)$.
3. Trivially $\mu_0(A) \geq 0$ for $A \in \mathcal{S}_0$. Thus we just need to show the countable additivity property. To understand the proof you need to keep several facts in mind: the functions μ and μ_0 agree on \mathcal{S} (property (b)); equivalence is preserved under set operations; equivalent sets have the same value under μ_0 (property (a)). Since the measure space (S, \mathcal{S}, μ) is σ -finite, there exists a countable disjoint collection $\{C_i : i \in I\}$ of sets in \mathcal{S} such that $S = \bigcup_{i \in I} C_i$ and $\mu(C_i) < \infty$ for each $i \in I$. Suppose first that $A \in \mathcal{S}_0$, so that there exists $B \in \mathcal{S}$ with $A \equiv B$. Then

$$\mu_0(A) = \mu_0\left[\bigcup_{i \in I} (A \cap C_i)\right] = \mu\left[\bigcup_{i \in I} (B \cap C_i)\right] = \sum_{i \in I} \mu(B \cap C_i) = \sum_{i \in I} \mu_0(A \cap C_i) \quad (2.7.29)$$

Suppose next that (A_1, A_2, \dots) is a sequence of pairwise disjoint sets in \mathcal{S}_0 so that there exists a sequence (B_1, B_2, \dots) of sets in \mathcal{S} such that $A_i \equiv B_i$ for each $i \in \mathbb{N}_+$. For fixed $i \in I$,

$$\mu_0\left[\bigcup_{n=1}^{\infty} (A_n \cap C_i)\right] = \mu_0\left[\bigcup_{n=1}^{\infty} (B_n \cap C_i)\right] = \mu\left[\bigcup_{n=1}^{\infty} (B_n \cap C_i)\right] = \sum_{n=1}^{\infty} \mu(B_n \cap C_i) = \sum_{n=1}^{\infty} \mu_0(A_n \cap C_i) \quad (2.7.30)$$

The next-to-the-last equality use the inclusion-exclusion law, since we don't know (and it's probably not true) that the sequence (B_1, B_2, \dots) is disjoint. The use of inclusion-exclusion is why we need (S, \mathcal{S}, μ) to be σ -finite. Finally, using the previous displayed equations,

$$\begin{aligned}\mu_0\left(\bigcup_{n=1}^{\infty} A_n\right) &= \sum_{i \in I} \mu_0\left[\left(\bigcup_{n=1}^{\infty} A_n\right) \cap C_i\right] = \sum_{i \in I} \mu_0\left(\bigcup_{n=1}^{\infty} A_n \cap C_i\right) \\ &= \sum_{i \in I} \sum_{n=1}^{\infty} \mu_0(A_n \cap C_i) = \sum_{n=1}^{\infty} \sum_{i \in I} \mu_0(A_n \cap C_i) = \sum_{n=1}^{\infty} \mu_0(A_n)\end{aligned}$$

Examples and Exercises

As always, be sure to try the computational exercises and proofs yourself before reading the answers and proofs in the text. Recall that a *discrete measure space* consists of a countable set, with the σ -algebra of all subsets, and with counting measure $\#$.

Counterexamples

The continuity theorem for decreasing events can fail if the events do not have finite measure.

Consider \mathbb{Z} with counting measure $\#$ on the σ -algebra of all subsets. Let $A_n = \{z \in \mathbb{Z} : z \leq -n\}$ for $n \in \mathbb{N}_+$. The continuity theorem fails for (A_1, A_2, \dots) .

Proof

The sequence is decreasing and $\#(A_n) = \infty$ for each n , but $\#(\bigcap_{i=1}^{\infty} A_i) = \#(\emptyset) = 0$.

Equal measure certainly does not imply equivalent sets.

Suppose that (S, \mathcal{S}, μ) is a measure space with the property that there exist disjoint sets $A, B \in \mathcal{S}$ such that $\mu(A) = \mu(B) > 0$. Then A and B are not equivalent.

Proof

Note that $A \triangle B = A \cup B$ and $\mu(A \cup B) > 0$.

For a concrete example, we could take $S = \{0, 1\}$ with counting measure $\#$ on σ -algebra of all subsets, and $A = \{0\}$, $B = \{1\}$.

The σ -finite property is not necessarily inherited by a sub-measure space. To set the stage for the counterexample, let \mathcal{B} denote the Borel σ -algebra of \mathbb{R} , that is, the σ -algebra generated by the standard Euclidean topology. There exists a positive measure λ on $(\mathbb{R}, \mathcal{B})$ that generalizes length. The measure λ , known as *Lebesgue measure*, is constructed in the section on Existence. Next let \mathcal{C} denote the σ -algebra of *countable and co-countable* sets:

$$\mathcal{C} = \{A \subseteq \mathbb{R} : A \text{ is countable or } A^c \text{ is countable}\} \quad (2.7.31)$$

That \mathcal{C} is a σ -algebra was shown in the section on measure theory in the chapter on foundations.

$(\mathbb{R}, \mathcal{C})$ is a subspace of $(\mathbb{R}, \mathcal{B})$. Moreover, $(\mathbb{R}, \mathcal{B}, \lambda)$ is σ -finite but $(\mathbb{R}, \mathcal{C}, \lambda)$ is not.

Proof

If $x \in \mathbb{R}$, then the singleton $\{x\}$ is closed and hence is in \mathcal{B} . A countable set is a countable union of singletons, so if A is countable then $A \in \mathcal{B}$. It follows that $\mathcal{C} \subset \mathcal{B}$. Next, let I_n denote the interval $[n, n+1)$ for $n \in \mathbb{Z}$. Then $\lambda(I_n) = 1$ for $n \in \mathbb{Z}$ and $\mathbb{R} = \bigcup_{n \in \mathbb{Z}} I_n$, so $(\mathbb{R}, \mathcal{B}, \lambda)$ is σ -finite. On the other hand, $\lambda\{x\} = 0$ for $x \in \mathbb{R}$ (since the set is an interval of length 0). Therefore $\lambda(A) = 0$ if A is countable and $\lambda(A) = \infty$ if A^c is countable. It follows that \mathbb{R} cannot be written as a countable union of sets in \mathcal{C} , each with finite measure.

A sum of finite measures may not be σ -finite.

Let S be a nonempty, finite set with the σ -algebra \mathcal{S} of all subsets. Let $\mu_n = \#$ be counting measure on (S, \mathcal{S}) for $n \in \mathbb{N}_+$. Then μ_n is a finite measure for each $n \in \mathbb{N}_+$, but $\mu = \sum_{n \in \mathbb{N}_+} \mu_n$ is not σ -finite.

Proof

Note that μ is the trivial measure on (S, \mathcal{S}) given by $\mu(A) = \infty$ if $A \neq \emptyset$ (and of course $\mu(\emptyset) = 0$).

Basic Properties

In the following problems, μ is a positive measure on the measurable space (S, \mathcal{S}) .

Suppose that $\mu(S) = 20$ and that $A, B \in \mathcal{S}$ with $\mu(A) = 5$, $\mu(B) = 6$, $\mu(A \cap B) = 2$. Find the measure of each of the following sets:

1. $A \setminus B$
2. $A \cup B$
3. $A^c \cup B^c$
4. $A^c \cap B^c$
5. $A \cup B^c$

Answer

1. 3
2. 9
3. 18
4. 11
5. 16

Suppose that $\mu(S) = \infty$ and that $A, B \in \mathcal{S}$ with $\mu(A \setminus B) = 2$, $\mu(B \setminus A) = 3$, and $\mu(A \cap B) = 4$. Find the measure of each of the following sets:

1. A
2. B
3. $A \cup B$
4. $A^c \cap B^c$
5. $A^c \cup B^c$

Answer

1. 6
2. 7
3. 9
4. ∞
5. ∞

Suppose that $\mu(S) = 10$ and that $A, B \in \mathcal{S}$ with $\mu(A) = 3$, $\mu(A \cup B) = 7$, and $\mu(A \cap B) = 2$. Find the measure of each of the following events:

1. B
2. $A \setminus B$
3. $B \setminus A$
4. $A^c \cup B^c$
5. $A^c \cap B^c$

Answer

1. 6
2. 1
3. 4
4. 8
5. 3

Suppose that $A, B, C \in \mathcal{S}$ with $\mu(A) = 10$, $\mu(B) = 12$, $\mu(C) = 15$, $\mu(A \cap B) = 3$, $\mu(A \cap C) = 4$, $\mu(B \cap C) = 5$, and $\mu(A \cap B \cap C) = 1$. Find the probabilities of the various unions:

1. $A \cup B$
2. $A \cup C$
3. $B \cup C$

4. $A \cup B \cup C$

Answer

1. 21
2. 23
3. 22
4. 28

This page titled [2.7: Measure Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.8: Existence and Uniqueness

Suppose that S is a set and \mathcal{S} a σ -algebra of subsets of S , so that (S, \mathcal{S}) is a measurable space. In many cases, it is impossible to define a positive measure μ on \mathcal{S} explicitly, by giving a “formula” for computing $\mu(A)$ for each $A \in \mathcal{S}$. Rather, we often know how the measure μ should work on some class of sets \mathcal{B} that generates \mathcal{S} . We would then like to know that μ can be *extended* to a positive measure on \mathcal{S} , and that this extension is *unique*. The purpose of this section is to discuss the basic results on this topic. To understand this section you will need to review the sections on Measure Theory and Special Set Structures in the chapter on Foundations, and the section on Measure Spaces in this chapter. If you are not interested in questions of existence and uniqueness of positive measures, you can safely skip this section.

Basic Theory

Positive Measures on Algebras

Suppose first that \mathcal{A} is an algebra of subsets of S . Recall that this means that \mathcal{A} is a collection of subsets that contains S and is closed under complements and finite unions (and hence also finite intersections). Here is our first definition:

A *positive measure* on \mathcal{A} is a function $\mu : \mathcal{A} \rightarrow [0, \infty]$ that satisfies the following properties:

1. $\mu(\emptyset) = 0$
2. If $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{A} and if $\bigcup_{i \in I} A_i \in \mathcal{A}$ then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i) \quad (2.8.1)$$

Clearly the definition of a positive measure on an algebra is very similar to the definition for a σ -algebra. If the collection of sets in (b) is finite, then $\bigcup_{i \in I} A_i$ must be in the algebra \mathcal{A} . Thus, μ is finitely additive. If the collection is countably infinite, then there is no guarantee that the union is in \mathcal{A} . If it is however, then μ must be additive over this collection. Given the similarity, it is not surprising that μ shares many of the basic properties of a positive measure on a σ -algebra, with proofs that are almost identical.

If $A, B \in \mathcal{A}$, then $\mu(B) = \mu(A \cap B) + \mu(B \setminus A)$.

Proof

Note that $B = (A \cap B) \cup (B \setminus A)$, and the sets in the union are in the algebra \mathcal{A} and are disjoint.

If $A, B \in \mathcal{A}$ and $A \subseteq B$ then

1. $\mu(B) = \mu(A) + \mu(B \setminus A)$
2. $\mu(A) \leq \mu(B)$

Proof

Part (a) follows from the previous theorem, since $A \cap B = A$. Part (b) follows from part (a).

Thus μ is increasing, relative to the subset partial order \subseteq on \mathcal{A} and the ordinary order \leq on $[0, \infty]$. Note also that if $A, B \in \mathcal{A}$ and $\mu(B) < \infty$ then $\mu(B \setminus A) = \mu(B) - \mu(A \cap B)$. In the special case that $A \subseteq B$, this becomes $\mu(B \setminus A) = \mu(B) - \mu(A)$. If $\mu(S) < \infty$ then $\mu(A^c) = \mu(S) - \mu(A)$. These are the familiar *difference* and *complement* rules.

The following result is the *subadditive property* for a positive measure μ on an algebra \mathcal{A} .

Suppose that $\{A_i : i \in I\}$ is a countable collection of sets in \mathcal{A} and that $\bigcup_{i \in I} A_i \in \mathcal{A}$. Then

$$\mu\left(\bigcup_{i \in I} A_i\right) \leq \sum_{i \in I} \mu(A_i) \quad (2.8.2)$$

Proof

The proof is just like before. Assume that $I = \mathbb{N}_+$. Let $B_1 = A_1$ and $B_i = A_i \setminus (A_1 \cup \dots \cup A_{i-1})$ for $i \in \{2, 3, \dots\}$. Then $\{B_i : i \in I\}$ is a disjoint collection of sets in \mathcal{A} with the same union as $\{A_i : i \in I\}$. Also $B_i \subseteq A_i$ for each i so $\mu(B_i) \leq \mu(A_i)$. Hence if the union is in \mathcal{A} then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \mu\left(\bigcup_{i \in I} B_i\right) = \sum_{i \in I} \mu(B_i) \leq \sum_{i \in I} \mu(A_i) \quad (2.8.3)$$

For a finite union of sets with finite measure, the inclusion-exclusion formula holds, and the proof is just like the one for a probability measure.

Suppose that $\{A_i : i \in I\}$ is a finite collection of sets in \mathcal{A} where $\#(I) = n \in \mathbb{N}_+$, and that $\mu(A_i) < \infty$ for $i \in I$. Then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mu\left(\bigcap_{j \in J} A_j\right) \quad (2.8.4)$$

The continuity theorems hold for a positive measure μ on an algebra \mathcal{A} , just as for a positive measure on a σ -algebra, assuming that the appropriate union and intersection are in the algebra. The proofs are just as before.

Suppose that (A_1, A_2, \dots) is a sequence of sets in \mathcal{A} .

1. If the sequence is increasing, so that $A_n \subseteq A_{n+1}$ for each $n \in \mathbb{N}_+$, and $\bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$, then $\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \lim_{n \rightarrow \infty} \mu(A_n)$.
2. If the sequence is decreasing, so that $A_{n+1} \subseteq A_n$ for each $n \in \mathbb{N}_+$, and $\mu(A_1) < \infty$ and $\bigcap_{i=1}^{\infty} A_i \in \mathcal{A}$, then $\mu\left(\bigcap_{i=1}^{\infty} A_i\right) = \lim_{n \rightarrow \infty} \mu(A_n)$.

Proof

1. Note that if $\mu(A_k) = \infty$ for some k then $\mu(A_n) = \infty$ for $n \geq k$ and $\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \infty$ if this union is in \mathcal{A} . Thus, suppose that $\mu(A_i) < \infty$ for each i . Let $B_1 = A_1$ and $B_i = A_i \setminus A_{i-1}$ for $i \in \{2, 3, \dots\}$. Then (B_1, B_2, \dots) is a disjoint sequence in \mathcal{A} with the same union as (A_1, A_2, \dots) . Also, $\mu(B_1) = \mu(A_1)$ and $\mu(B_i) = \mu(A_i) - \mu(A_{i-1})$ for $i \in \{2, 3, \dots\}$. Hence if the union is in \mathcal{A} ,

$$\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \mu\left(\bigcup_{i=1}^{\infty} B_i\right) = \sum_{i=1}^{\infty} \mu(B_i) = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mu(B_i) \quad (2.8.5)$$

But $\sum_{i=1}^n \mu(B_i) = \mu(A_1) + \sum_{i=2}^n [\mu(A_i) - \mu(A_{i-1})] = \mu(A_n)$.

2. Note that $A_1 \setminus A_n \in \mathcal{A}$ and this sequence is increasing. Moreover, $\bigcup_{n=1}^{\infty} (A_1 \setminus A_n) = \left(\bigcap_{n=1}^{\infty} A_n\right)^c \cap A_1$. Hence if $\bigcap_{n=1}^{\infty} A_n \in \mathcal{A}$ then $\bigcup_{n=1}^{\infty} (A_1 \setminus A_n) \in \mathcal{A}$. Thus using the continuity result for increasing sets,

$$\mu\left(\bigcap_{i=1}^{\infty} A_i\right) = \mu\left[A_1 \setminus \bigcup_{i=1}^{\infty} (A_1 \setminus A_i)\right] = \mu(A_1) - \mu\left[\bigcup_{i=1}^{\infty} (A_1 \setminus A_n)\right] \quad (2.8.6)$$

$$= \mu(A_1) - \lim_{n \rightarrow \infty} \mu(A_1 \setminus A_n) = \mu(A_1) - \lim_{n \rightarrow \infty} [\mu(A_1) - \mu(A_n)] = \lim_{n \rightarrow \infty} \mu(A_n) \quad (2.8.7)$$

Recall that if the sequence (A_1, A_2, \dots) is increasing, then we define $\lim_{n \rightarrow \infty} A_n = \bigcup_{n=1}^{\infty} A_n$, and if the sequence is decreasing then we define $\lim_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} A_n$. Thus the conclusion of both parts of the continuity theorem is

$$\mathbb{P}\left(\lim_{n \rightarrow \infty} A_n\right) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n) \quad (2.8.8)$$

Finite additivity and continuity for increasing events imply countable additivity:

If $\mu : \mathcal{A} \rightarrow [0, \infty]$ satisfies the properties below then μ is a positive measure on \mathcal{A} .

1. $\mu(\emptyset) = 0$
2. $\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i)$ if $\{A_i : i \in I\}$ is a finite disjoint collection of sets in \mathcal{A}
3. $\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \lim_{n \rightarrow \infty} \mu(A_n)$ if (A_1, A_2, \dots) is an increasing sequence of events in \mathcal{A} and $\bigcup_{i=1}^{\infty} A_i \in \mathcal{A}$.

Proof

All that is left to prove is additivity over a countably infinite collection of sets in \mathcal{A} when the union is also in \mathcal{A} . Thus suppose that $\{A_n : n \in \mathbb{N}\}$ is a disjoint collection of sets in \mathcal{A} with $\bigcup_{n=1}^{\infty} A_n \in \mathcal{A}$. Let $B_n = \bigcup_{i=1}^n A_i$ for $n \in \mathbb{N}_+$. Then $B_n \in \mathcal{A}$ and $\bigcup_{n=1}^{\infty} B_n = \bigcup_{n=1}^{\infty} A_n$. Hence using the finite additivity and the continuity property we have

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \mathbb{P}\left(\bigcup_{n=1}^{\infty} B_n\right) = \lim_{n \rightarrow \infty} \mathbb{P}(B_n) = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{P}(A_i) = \sum_{i=1}^{\infty} \mathbb{P}(A_i) \quad (2.8.9)$$

Many of the basic theorems in measure theory require that the measure not be too far removed from being finite. This leads to the following definition, which is just like the one for a positive measure on a σ -algebra.

A measure μ on an algebra \mathcal{A} of subsets of S is σ -finite if there exists a sequence of sets (A_1, A_2, \dots) in \mathcal{A} such that $\bigcup_{n=1}^{\infty} A_n = S$ and $\mu(A_n) < \infty$ for each $n \in \mathbb{N}_+$. The sequence is called a σ -finite sequence for μ .

Suppose that μ is a σ -finite measure on an algebra \mathcal{A} of subsets of S .

1. There exists an increasing σ -finite sequence.
2. There exists a disjoint σ -finite sequence.

Proof

We use the same tricks that we have used before. Suppose that (A_1, A_2, \dots) is a σ -finite sequence for μ .

1. Let $B_n = \bigcup_{i=1}^n A_i$. Then $B_n \in \mathcal{A}$ for $n \in \mathbb{N}_+$ and this sequence is increasing. Moreover, $\mu(B_n) \leq \sum_{i=1}^n \mu(A_i) < \infty$ for $n \in \mathbb{N}_+$ and $\bigcup_{n=1}^{\infty} B_n = \bigcup_{n=1}^{\infty} A_n = S$.
2. Let $C_1 = A_1$ and let $C_n = A_n \setminus \bigcup_{i=1}^{n-1} A_i$ for $n \in \{2, 3, \dots\}$. Then $C_n \in \mathcal{A}$ for each $n \in \mathbb{N}_+$ and this sequence is disjoint. Moreover, $C_n \subseteq A_n$ so $\mu(C_n) \leq \mu(A_n) < \infty$ and $\bigcup_{n=1}^{\infty} C_n = \bigcup_{n=1}^{\infty} A_n = S$.

Extension and Uniqueness Theorems

The fundamental theorem on measures states that a positive, σ -finite measure μ on an algebra \mathcal{A} can be uniquely extended to $\sigma(\mathcal{A})$. The extension part is sometimes referred to as the *Carathéodory extension theorem*, and is named for the Greek mathematician Constantin Carathéodory.

If μ is a positive, σ -finite measure on an algebra \mathcal{A} , then μ can be extended to a positive measure on $\mathcal{S} = \sigma(\mathcal{A})$.

Proof

The proof is complicated, but here is a broad outline. First, for $A \subseteq S$, we define a *cover* of A to be a countable collection $\{A_i : i \in I\}$ of sets in \mathcal{A} such that $A \subseteq \bigcup_{i \in I} A_i$. Next, we define a new set function μ^* , the *outer measure*, on all subsets of S :

$$\mu^*(A) = \inf \left\{ \sum_{i \in I} \mu(A_i) : \{A_i : i \in I\} \text{ is a cover of } A \right\}, \quad A \subseteq S \quad (2.8.10)$$

Outer measure satisfies the following properties.

1. $\mu^*(A) \geq 0$ for $A \subseteq S$, so μ^* is nonnegative.
2. $\mu^*(A) = \mu(A)$ for $A \in \mathcal{A}$, so μ^* extends μ .
3. If $A \subseteq B$ then $\mu^*(A) \leq \mu^*(B)$, so μ^* is increasing.
4. If $A_i \subseteq S$ for each i in a countable index set I then $\mu^*\left(\bigcup_{i \in I} A_i\right) \leq \sum_{i \in I} \mu^*(A_i)$, so μ^* is countably subadditive.

Next, $A \subseteq S$ is said to be *measurable* if

$$\mu^*(B) = \mu^*(B \cap A) + \mu^*(B \setminus A), \quad B \subseteq S \quad (2.8.11)$$

Thus, A is measurable if μ^* is additive with respect to the partition of B induced by $\{A, A^c\}$, for every $B \subseteq S$. We let \mathcal{M} denote the collection of measurable subsets of S . The proof is finished by showing that $\mathcal{A} \subseteq \mathcal{M}$, \mathcal{M} is a σ -algebra of subsets of S , and μ^* is a positive measure on \mathcal{M} . It follows that $\sigma(\mathcal{A}) = \mathcal{S} \subseteq \mathcal{M}$ and hence μ^* is a measure on \mathcal{S} that extends μ .

Our next goal is the basic uniqueness result, which serves as the complement to the basic extension result. But first we need another variation of the term σ -finite.

Suppose that μ is a measure on a σ -algebra \mathcal{S} of subsets of S and $\mathcal{B} \subseteq \mathcal{S}$. Then μ is σ -finite on \mathcal{B} if there exists a countable collection $\{B_i : i \in I\} \subseteq \mathcal{B}$ such that $\mu(B_i) < \infty$ for $i \in I$ and $\bigcup_{i \in I} B_i = S$.

The next result is the uniqueness theorem. The proof, like others that we have seen, uses Dynkin's π - λ theorem, named for Eugene Dynkin.

Suppose that \mathcal{B} is a π -system and that $\mathcal{S} = \sigma(\mathcal{B})$. If μ_1 and μ_2 are positive measures on \mathcal{S} and are σ -finite on \mathcal{B} , and if $\mu_1(A) = \mu_2(A)$ for all $A \in \mathcal{B}$, then $\mu_1(A) = \mu_2(A)$ for all $A \in \mathcal{S}$.

Proof

Suppose that $B \in \mathcal{B}$ and that $\mu_1(B) = \mu_2(B) < \infty$. Let $\mathcal{L}_B = \{A \in \mathcal{S} : \mu_1(A \cap B) = \mu_2(A \cap B)\}$. Then $S \in \mathcal{L}_B$ since $\mu_1(B) = \mu_2(B)$. If $A \in \mathcal{L}_B$ then $\mu_1(A \cap B) = \mu_2(A \cap B)$ so $\mu_1(A^c \cap B) = \mu_1(B) - \mu_1(A \cap B) = \mu_2(B) - \mu_2(A \cap B) = \mu_2(A^c \cap B)$ and hence $A^c \in \mathcal{L}_B$. Finally, suppose that $\{A_j : j \in J\}$ is a countable, disjoint collection of events in \mathcal{L}_B . Then $\mu_1(A_j \cap B) = \mu_2(A_j \cap B)$ for each $j \in J$ and hence

$$\mu_1 \left[\left(\bigcup_{j \in J} A_j \right) \cap B \right] = \mu_1 \left(\bigcup_{j \in J} (A_j \cap B) \right) = \sum_{j \in J} \mu_1(A_j \cap B) \quad (2.8.12)$$

$$= \sum_{j \in J} \mu_2(A_j \cap B) = \mu_2 \left(\bigcup_{j \in J} (A_j \cap B) \right) = \mu_2 \left[\left(\bigcup_{j \in J} A_j \right) \cap B \right] \quad (2.8.13)$$

Therefore $\bigcup_{j \in J} A_j \in \mathcal{L}_B$, and so \mathcal{L}_B is a λ -system. By assumption, $\mathcal{B} \subseteq \mathcal{L}_B$ and therefore by the π - λ theorem, $\mathcal{S} = \sigma(\mathcal{B}) \subseteq \mathcal{L}_B$.

Next, by assumption there exists $B_i \in \mathcal{B}$ with $\mu_1(B_i) = \mu_2(B_i) < \infty$ for each $i \in \mathbb{N}_+$ and $S = \bigcup_{i=1}^{\infty} B_i$. If $A \in \mathcal{S}$ then the inclusion-exclusion rule can be applied to

$$\mu_k \left[\left(\bigcup_{i=1}^n B_i \right) \cap A \right] = \mu_k \left[\bigcup_{i=1}^n (A \cap B_i) \right] \quad (2.8.14)$$

where $k \in \{1, 2\}$ and $n \in \mathbb{N}_+$. But the inclusion-exclusion formula only has terms of the form $\mu_k \left[\bigcap_{j \in J} (A \cap B_j) \right] = \mu_k \left[A \cap \left(\bigcap_{j \in J} B_j \right) \right]$ where $J \subseteq \{1, 2, \dots, n\}$. But $\bigcap_{j \in J} B_j \in \mathcal{B}$ since \mathcal{B} is a π -system, so by the previous paragraph, $\mu_1 \left[\bigcap_{j \in J} (A \cap B_j) \right] = \mu_2 \left[\bigcap_{j \in J} (A \cap B_j) \right]$. It then follows that for each $n \in \mathbb{N}_+$

$$\mu_1 \left[\left(\bigcup_{i=1}^n B_i \right) \cap A \right] = \mu_2 \left[\left(\bigcup_{i=1}^n B_i \right) \cap A \right] \quad (2.8.15)$$

Finally, letting $n \rightarrow \infty$ and using the continuity theorem for increasing sets gives $\mu_1(A) = \mu_2(A)$.

An algebra \mathcal{A} of subsets of S is trivially a π -system. Hence, if μ_1 and μ_2 are positive measures on $\mathcal{S} = \sigma(\mathcal{A})$ and are σ -finite on \mathcal{A} , and if $\mu_1(A) = \mu_2(A)$ for $A \in \mathcal{A}$, then $\mu_1(A) = \mu_2(A)$ for $A \in \mathcal{S}$. This completes the second part of the fundamental theorem.

Of course, the results of this subsection hold for probability measures. Formally, a probability measure \mathbb{P} on an algebra \mathcal{A} of subsets of S is a positive measure on \mathcal{A} with the additional requirement that $\mathbb{P}(S) = 1$. Probability measures are trivially σ -finite, so a probability measure \mathbb{P} on an algebra \mathcal{A} can be uniquely extended to $\mathcal{S} = \sigma(\mathcal{A})$.

However, usually we start with a collection that is more primitive than an algebra. The next result combines the definition with the main theorem associated with the definition. For a proof see the section on Special Set Structures in the chapter on Foundations.

Suppose that \mathcal{B} is a nonempty collection of subsets of S and let

$$\mathcal{A} = \left\{ \bigcup_{i \in I} B_i : \{B_i : i \in I\} \text{ is a finite, disjoint collection of sets in } \mathcal{B} \right\} \quad (2.8.16)$$

If the following conditions are satisfied, then \mathcal{B} is a *semi-algebra* of subsets of S , and then \mathcal{A} is the algebra generated by \mathcal{B} .

1. If $B_1, B_2 \in \mathcal{B}$ then $B_1 \cap B_2 \in \mathcal{B}$.
2. If $B \in \mathcal{B}$ then $B^c \in \mathcal{A}$.

Suppose now that we know how a measure μ should work on a semi-algebra \mathcal{B} that generates an algebra \mathcal{A} and then a σ -algebra $\mathcal{S} = \sigma(\mathcal{A}) = \sigma(\mathcal{B})$. That is, we know $\mu(B) \in [0, \infty]$ for each $B \in \mathcal{B}$. Because of the additivity property, there is no question as to how we should extend μ to \mathcal{A} . We must have

$$\mu(A) = \sum_{i \in I} \mu(B_i) \quad (2.8.17)$$

if $A = \bigcup_{i \in I} B_i$ for some finite, disjoint collection $\{B_i : i \in I\}$ of sets in \mathcal{B} (so that $A \in \mathcal{A}$). However, we cannot assign the values $\mu(B)$ for $B \in \mathcal{B}$ arbitrarily. The following extension theorem states that, subject just to some essential consistency conditions, the extension of μ from the semi-algebra \mathcal{B} to the algebra \mathcal{A} does in fact produce a measure on \mathcal{A} . The consistency conditions are that μ be *finitely additive* and *countably subadditive* on \mathcal{B} .

Suppose that \mathcal{B} is a semi-algebra of subsets of S and that \mathcal{A} is the algebra of subsets of S generated by \mathcal{B} . A function $\mu : \mathcal{B} \rightarrow [0, \infty]$ can be uniquely extended to a measure on \mathcal{A} if and only if μ satisfies the following properties:

1. If $\emptyset \in \mathcal{B}$ then $\mu(\emptyset) = 0$.
2. If $\{B_i : i \in I\}$ is a finite, disjoint collection of sets in \mathcal{B} and $B = \bigcup_{i \in I} B_i \in \mathcal{B}$ then $\mu(B) = \sum_{i \in I} \mu(B_i)$.
3. If $B \in \mathcal{B}$ and $B \subseteq \bigcup_{i \in I} B_i$ where $\{B_i : i \in I\}$ is a countable collection of sets in \mathcal{B} then $\mu(B) \leq \sum_{i \in I} \mu(B_i)$.

If the measure μ on the algebra \mathcal{A} is σ -finite, then the [extension theorem](#) and the [uniqueness theorem](#) apply, so μ can be extended uniquely to a measure on the σ -algebra $\mathcal{S} = \sigma(\mathcal{A}) = \sigma(\mathcal{B})$. This chain of extensions, starting with a semi-algebra \mathcal{B} , is often how measures are constructed.

Examples and Applications

Product Spaces

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. For the Cartesian product set $S \times T$, recall that the *product σ -algebra* is

$$\mathcal{S} \otimes \mathcal{T} = \sigma\{A \times B : A \in \mathcal{S}, B \in \mathcal{T}\} \quad (2.8.18)$$

the σ -algebra generated by the Cartesian products of measurable sets, sometimes referred to as *measurable rectangles*.

Suppose that (S, \mathcal{S}, μ) and (T, \mathcal{T}, ν) are σ -finite measure spaces. Then there exists a unique σ -finite measure $\mu \otimes \nu$ on $(S \times T, \mathcal{S} \otimes \mathcal{T})$ such that

$$(\mu \otimes \nu)(A \times B) = \mu(A)\nu(B); \quad A \in \mathcal{S}, B \in \mathcal{T} \quad (2.8.19)$$

The measure space $(S \times T, \mathcal{S} \otimes \mathcal{T}, \mu \otimes \nu)$ is the *product measure space* associated with (S, \mathcal{S}, μ) and (T, \mathcal{T}, ν) .

Proof

Recall that the collection $\mathcal{B} = \{A \times B : A \in \mathcal{S}, B \in \mathcal{T}\}$ is a semi-algebra: the intersection of two product sets is another product set, and the complement of a product set is the union of two disjoint product sets. We define $\rho : \mathcal{B} \rightarrow [0, \infty]$ by $\rho(A \times B) = \mu(A)\nu(B)$. The [consistency conditions](#) hold, so ρ can be extended to a measure on the algebra \mathcal{A} generated by \mathcal{B} . The algebra \mathcal{A} is the collection of all finite, disjoint unions of products of measurable sets. We will now show that the extended measure ρ is σ -finite on \mathcal{A} . Since μ is σ -finite, [there exists](#), an increasing sequence (A_1, A_2, \dots) of sets in \mathcal{S} with $\mu(A_i) < \infty$ and $\bigcup_{i=1}^{\infty} A_i = S$. Similarly, there exists an increasing sequence (B_1, B_2, \dots) of sets in \mathcal{T} with $\nu(B_j) < \infty$ and $\bigcup_{j=1}^{\infty} B_j = T$. Then $\rho(A_i \times B_j) = \mu(A_i)\nu(B_j) < \infty$, and since the sets are increasing, $\bigcup_{(i,j) \in \mathbb{N}_+ \times \mathbb{N}_+} A_i \times B_j = S \times T$. The standard [extension theorem](#) and uniqueness theorem [uniqueness theorem](#) now apply, so ρ can be extended uniquely to a measure on $\sigma(\mathcal{A}) = \mathcal{S} \otimes \mathcal{T}$.

Recall that for $C \subseteq S \times T$, the cross section of C in the first coordinate at $x \in S$ is $C_x = \{y \in T : (x, y) \in C\}$. Similarly, the cross section of C in the second coordinate at $y \in T$ is $C^y = \{x \in S : (x, y) \in C\}$. We know that the cross sections of a measurable set are measurable. The following result shows that the measures of the cross sections of a measurable set form measurable functions.

Suppose again that (S, \mathcal{S}, μ) and (T, \mathcal{T}, ν) are σ -finite measure spaces. If $C \in \mathcal{S} \otimes \mathcal{T}$ then

1. $x \mapsto \nu(C_x)$ is a measurable function from S to $[0, \infty]$.
2. $y \mapsto \mu(C^y)$ is a measurable function from T to $[0, \infty]$.

Proof

We prove part (a), since of course the proof for part (b) is symmetric. Suppose first that the measure spaces are finite. Let $\mathcal{R} = \{A \times B : A \in \mathcal{S}, B \in \mathcal{T}\}$ denote the set of measurable rectangles. Let $\mathcal{C} = \{C \in \mathcal{S} \otimes \mathcal{T} : x \mapsto \nu(C_x) \text{ is measurable}\}$. If $A \times B \in \mathcal{R}$, then $A \times B \in \mathcal{C}$, since $\nu[(A \times B)_x] = \nu(B)\mathbf{1}_A(x)$. Next, suppose $C \in \mathcal{C}$. Then $(C^c)_x = (C_x)^c$, so $\nu[(C^c)_x] = \nu(T) - \nu(C_x)$ and this is a measurable function of $x \in S$. Hence $C^c \in \mathcal{C}$. Next, suppose that $\{C_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{C} and let $C = \bigcup_{i \in I} C_i$. Then $\{(C_i)_x : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{T} , and $C_x = \bigcup_{i \in I} (C_i)_x$. Hence $\nu(C_x) = \sum_{i \in I} \nu[(C_i)_x]$, and this is a measurable function of $x \in S$. Hence $C \in \mathcal{C}$. It follows that \mathcal{C} is a λ -system that contains \mathcal{R} , which in turn is a π -system. It follows from Dynkin's π - λ theorem, that $\mathcal{S} \otimes \mathcal{T} = \sigma(\mathcal{R}) \subseteq \mathcal{C}$. Thus $\mathcal{C} = \mathcal{S} \otimes \mathcal{T}$.

Consider now the general case where the measure spaces are σ -finite. There exists a countable, increasing sequence of sets $C_n \in \mathcal{S} \otimes \mathcal{T}$ for $n \in \mathbb{N}_+$ with $(\mu \otimes \nu)(C_n) < \infty$ for $n \in \mathbb{N}_+$. If $C \in \mathcal{S} \otimes \mathcal{T}$, then $C \cap C_n$ is increasing in $n \in \mathbb{N}_+$, and $C = \bigcup_{n=1}^{\infty} (C \cap C_n)$. Hence, for $x \in S$, $(C \cap C_n)_x$ is increasing in $n \in \mathbb{N}_+$ and $C_x = \bigcup_{n=1}^{\infty} (C \cap C_n)_x$. Therefore $\nu(C_x) = \lim_{n \rightarrow \infty} \nu[(C \cap C_n)_x]$. But $x \mapsto \nu[(C \cap C_n)_x]$ is a measurable function of $x \in S$ for each $n \in \mathbb{N}_+$ by the previous argument, so $x \mapsto \nu(C_x)$ is a measurable function of $x \in S$.

In the next chapter, where we study integration with respect to a measure, we will see that for $C \in \mathcal{S} \otimes \mathcal{T}$, the product measure $(\mu \otimes \nu)(C)$ can be computed by integrating $\nu(C_x)$ over $x \in S$ with respect to μ or by integrating $\mu(C^y)$ over $y \in T$ with respect to ν . These results, generalizing the definition of the [product measure](#), are special cases of Fubini's theorem, named for the Italian mathematician Guido Fubini.

Except for more complicated notation, these results extend in a perfectly straightforward way to the product of a finite number of σ -finite measure spaces.

Suppose that $n \in \mathbb{N}_+$ and that $(S_i, \mathcal{S}_i, \mu_i)$ is a σ -finite measure space for $i \in \{1, 2, \dots, n\}$. Let $S = \prod_{i=1}^n S_i$ and let \mathcal{S} denote the corresponding product σ -algebra. There exists a unique σ -finite measure μ on (S, \mathcal{S}) satisfying

$$\mu \left(\prod_{i=1}^n A_i \right) = \prod_{i=1}^n \mu_i(A_i), \quad A_i \in \mathcal{S}_i \text{ for } i \in \{1, 2, \dots, n\} \quad (2.8.20)$$

The measure space (S, \mathcal{S}, μ) is the *product measure space* associated with the given measure spaces.

Lebesgue Measure

The next discussion concerns our most important and essential application. Recall that the *Borel* σ -algebra on \mathbb{R} , named for Émile Borel, is the σ -algebra \mathcal{B} generated by the standard Euclidean topology on \mathbb{R} . Equivalently, $\mathcal{B} = \sigma(\mathcal{I})$ where \mathcal{I} is the collection of intervals of \mathbb{R} (of all types—bounded and unbounded, with any type of closure, and including single points and the empty set). Next recall how the *length* of an interval is defined. For $a, b \in \mathbb{R}$ with $a \leq b$, each of the intervals (a, b) , $[a, b)$, $(a, b]$, and $[a, b]$ has length $b - a$. For $a \in \mathbb{R}$, each of the intervals (a, ∞) , $[a, \infty)$, $(-\infty, a)$, $(-\infty, a]$ has length ∞ , as does \mathbb{R} itself. The standard measure on \mathcal{B} generalizes the length measurement for intervals.

There exists a unique measure λ on \mathcal{B} such that $\lambda(I) = \text{length}(I)$ for $I \in \mathcal{I}$. The measure λ is *Lebesgue measure* on $(\mathbb{R}, \mathcal{B})$.

Proof

Recall that \mathcal{I} is a semi-algebra: The intersection of two intervals is another interval, and the complement of an interval is either another interval or the union of two disjoint intervals. Define λ on \mathcal{I} by $\lambda(I) = \text{length}(I)$ for $I \in \mathcal{I}$. Then λ satisfies

the [consistency condition](#) and hence λ can be extended to a measure on the algebra \mathcal{J} generated by \mathcal{I} , namely the collection of finite, disjoint unions of intervals. The measure λ on \mathcal{J} is clearly σ -finite, since \mathbb{R} can be written as a countably infinite union of bounded intervals. Hence the standard [extension theorem](#) and [uniqueness theorem](#) apply, so λ can be extended to a measure on $\mathcal{R} = \sigma(\mathcal{J})$.

The is name in honor of Henri Lebesgue, of course. Since λ is σ -finite, the σ -algebra of Borel sets \mathcal{R} can be completed with respect to λ .

The completion of the Borel σ -algebra \mathcal{R} with respect to λ is the *Lebesgue σ -algebra* \mathcal{R}^* .

Recall that *completed* means that if $A \in \mathcal{R}^*$, $\lambda(A) = 0$ and $B \subseteq A$, then $B \in \mathcal{R}^*$ (and then $\lambda(B) = 0$). The Lebesgue measure λ on \mathbb{R} , with either the Borel σ -algebra \mathcal{R} , or its completion \mathcal{R}^* is the standard measure that is used for the real numbers. Other properties of the measure space $(\mathbb{R}, \mathcal{R}, \lambda)$ are given below, in the discussion of Lebesgue measure on \mathbb{R}^n .

For $n \in \mathbb{N}_+$, let \mathcal{R}_n denote the Borel σ -algebra corresponding to the the standard Euclidean topology on \mathbb{R}^n , so that $(\mathbb{R}^n, \mathcal{R}_n)$ is the n -dimensional Euclidean measurable space. The σ -algebra, \mathcal{R}_n is also the n -fold power of \mathcal{R} , the Borel σ -algebra of \mathbb{R} . That is, $\mathcal{R}_n = \mathcal{R} \otimes \mathcal{R} \otimes \cdots \otimes \mathcal{R}$ (n times). It is also the σ -algebra generated by the products of intervals:

$$\mathcal{R}_n = \sigma \{ I_1 \times I_2 \times \cdots \times I_n : I_j \in \mathcal{I} \text{ for } j \in \{1, 2, \dots, n\} \} \quad (2.8.21)$$

As above, let λ denote Lebesgue measure on $(\mathbb{R}, \mathcal{R})$.

For $n \in \mathbb{N}_+$ the n -fold power of λ , denoted λ_n is *Lebesgue measure* on $(\mathbb{R}^n, \mathcal{R}_n)$. In particular,

$$\lambda_n(A_1 \times A_2 \times \cdots \times A_n) = \lambda(A_1)\lambda(A_2) \cdots \lambda(A_n); \quad A_1, \dots, A_n \in \mathcal{R} \quad (2.8.22)$$

Specializing further, if $I_j \in \mathcal{I}$ is an interval for $j \in \{1, 2, \dots, n\}$ then

$$\lambda_n(I_1 \times I_2 \times \cdots \times I_n) = \text{length}(I_1)\text{length}(I_2) \cdots \text{length}(I_n) \quad (2.8.23)$$

In particular, λ_2 extends the *area* measure on \mathcal{R}_2 and λ_3 extends the *volume* measure on \mathcal{R}_3 . In general, $\lambda_n(A)$ is sometimes referred to as *n -dimensional volume* of $A \in \mathcal{R}_n$. As in the one-dimensional case, \mathcal{R}_n can be completed with respect to λ_n , essentially adding all subsets of sets of measure 0 to \mathcal{R}_n . The completed σ -algebra is the σ -algebra of *Lebesgue measurable* sets. Since $\lambda_n(U) > 0$ if $U \subseteq \mathbb{R}^n$ is open, the support of λ_n is all of \mathbb{R}^n . In addition, Lebesgue measure has the regularity properties that are concerned with approximating the measure of a set, from below with the measure of a compact set, and from above with the measure of an open set.

The measure space $(\mathbb{R}^n, \mathcal{R}_n, \lambda_n)$ is *regular*. That is, for $A \in \mathcal{R}_n$,

1. $\lambda_n(A) = \sup\{\lambda_n(C) : C \text{ is compact and } C \subseteq A\}$, (*inner regularity*)
2. $\lambda_n(A) = \inf\{\lambda_n(U) : U \text{ is open and } A \subseteq U\}$ (*outer regularity*).

The following theorem describes how the measure of a set is changed under certain basic transformations. These are essential properties of Lebesgue measure. To setup the notation, suppose that $n \in \mathbb{N}_+$, $A \subseteq \mathbb{R}^n$, $x \in \mathbb{R}^n$, $c \in (0, \infty)$ and that T is an $n \times n$ matrix. Define

$$A + x = \{a + x : a \in A\}, \quad cA = \{ca : a \in A\}, \quad TA = \{Ta : a \in A\} \quad (2.8.24)$$

Suppose that $A \in \mathcal{R}_n$.

1. If $x \in \mathbb{R}^n$ then $\lambda_n(A + x) = \lambda_n(A)$ (*translation invariance*)
2. If $c \in (0, \infty)$ then $\lambda_n(cA) = c^n \lambda_n(A)$ (*dilation property*)
3. If T is an $n \times n$ matrix then $\lambda_n(TA) = |\det(T)| \lambda_n(A)$ (*the scaling property*)

Lebesgue-Stieltjes Measures on \mathbb{R}

The construction of Lebesgue measure on \mathbb{R} can be generalized. Here is the definition that we will need.

A function $F : \mathbb{R} \rightarrow \mathbb{R}$ that satisfies the following properties is a *distribution function* on \mathbb{R}

1. F is increasing: if $x \leq y$ then $F(x) \leq F(y)$.
2. F is continuous from the right: $\lim_{t \downarrow x} F(t) = F(x)$ for all $x \in \mathbb{R}$.

Since F is increasing, the limit from the left at $x \in \mathbb{R}$ exists in \mathbb{R} and is denoted $F(x^-) = \lim_{t \uparrow x} F(t)$. Similarly $F(\infty) = \lim_{x \rightarrow \infty} F(x)$ exists, as a real number or ∞ , and $F(-\infty) = \lim_{x \rightarrow -\infty} F(x)$ exists, as a real number or $-\infty$.

If F is a distribution function on \mathbb{R} , then there exists a unique measure μ on \mathcal{R} that satisfies

$$\mu(a, b] = F(b) - F(a), \quad -\infty \leq a \leq b \leq \infty \quad (2.8.25)$$

The measure μ is called the *Lebesgue-Stieltjes measure* associated with F , named for Henri Lebesgue and Thomas Joannes Stieltjes. Distribution functions and the measures associated with them are studied in more detail in the chapter on Distributions. When the function F takes values in $[0, 1]$, the associated measure \mathbb{P} is a probability measure, and the function F is the *probability distribution function* of \mathbb{P} . Probability distribution functions are also studied in much more detail (but with less technicality) in the chapter on Distributions.

Note that the identity function $x \mapsto x$ for $x \in \mathbb{R}$ is a distribution function, and the measure associated with this function is ordinary Lebesgue measure on \mathbb{R} constructed in (15).

This page titled [2.8: Existence and Uniqueness](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.9: Probability Spaces Revisited

In this section we discuss probability spaces from the more advanced point of view of measure theory. The previous two sections on positive measures and existence and uniqueness are prerequisites. The discussion is divided into two parts: first those concepts that are shared rather equally between probability theory and general measure theory, and second those concepts that are for the most part unique to probability theory. In particular, it's a mistake to think of probability theory as a mere branch of measure theory. Probability has its own notation, terminology, point of view, and applications that makes it an incredibly rich subject on its own.

Basic Concepts

Our first discussion concerns topics that were discussed in the section on positive measures. So no proofs are necessary, but you will notice that the notation, and in some cases the terminology, is very different.

Definitions

We can now give a precise definition of the probability space, the mathematical model of a random experiment.

A *probability space* $(S, \mathcal{S}, \mathbb{P})$, consists of three essential parts:

1. A set of outcomes S .
2. A σ -algebra of events \mathcal{S} .
3. A probability measure \mathbb{P} on the *sample space* (S, \mathcal{S}) .

Often the special notation $(\Omega, \mathcal{F}, \mathbb{P})$ is used for a probability space in the literature—the symbol Ω for the set of outcomes is intended to remind us that these are *all* possible outcomes. However in this text, we don't insist on the special notation, and use whatever notation seems most appropriate in a given context.

In probability, σ -algebras are not just important for theoretical and foundational purposes, but are important for practical purposes as well. A σ -algebra can be used to specify *partial information* about an experiment—a concept of fundamental importance. Specifically, suppose that \mathcal{A} is a collection of events in the experiment, and that we know whether or not A occurred for each $A \in \mathcal{A}$. Then in fact, we can determine whether or not A occurred for each $A \in \sigma(\mathcal{A})$, the σ -algebra generated by \mathcal{A} .

Technically, a random variable for our experiment is a *measurable* function from the sample space into another measurable space.

Suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space and that (T, \mathcal{T}) is another measurable space. A *random variable* X with values in T is a measurable function from S into T .

1. The *probability distribution* of X is the mapping on \mathcal{T} given by $B \mapsto \mathbb{P}(X \in B)$.
2. The collection of events $\{X \in B : B \in \mathcal{T}\}$ is a sub σ -algebra of \mathcal{S} , and is the σ -algebra generated by X , denoted $\sigma(X)$.

Details




Figure 2.9.1: The event $\{X \in B\}$ associated with $B \in \mathcal{T}$

If we observe the value of X , then we know whether or not each event in $\sigma(X)$ has occurred. More generally, we can construct the σ -algebra associated with any collection of random variables.

suppose that (T_i, \mathcal{T}_i) is a measurable space for each i in an index set I , and that X_i is a random variable taking values in T_i for each $i \in I$. The σ -algebra generated by $\{X_i : i \in I\}$ is

$$\sigma\{X_i : i \in I\} = \sigma\{X \in B_i : B_i \in \mathcal{T}_i, i \in I\} \quad (2.9.1)$$

If we observe the value of X_i for each $i \in I$ then we know whether or not each event in $\sigma\{X_i : i \in I\}$ has occurred. This idea is very important in the study of stochastic processes.

Null Events, Almost Sure Events, and Equivalence

Suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space.

Define the following collections of events:

1. $\mathcal{N} = \{A \in \mathcal{S} : \mathbb{P}(A) = 0\}$, the collection of *null* events
2. $\mathcal{M} = \{A \in \mathcal{S} : \mathbb{P}(A) = 1\}$, the collection of *almost sure* events
3. $\mathcal{D} = \mathcal{N} \cup \mathcal{M} = \{A \in \mathcal{S} : \mathbb{P}(A) = 0 \text{ or } \mathbb{P}(A) = 1\}$, the collection of *essentially deterministic* events

The collection of essentially deterministic events \mathcal{D} is a sub σ -algebra of \mathcal{S} .

In the section on independence, we showed that \mathcal{D} is also a collection of independent events.

Intuitively, equivalent events or random variables are those that are indistinguishable from a probabilistic point of view. Recall first that the *symmetric difference* between events A and B is $A \triangle B = (A \setminus B) \cup (B \setminus A)$; it is the event that occurs if and only if one of the events occurs, but not the other, and corresponds to *exclusive or*. Here is the definition for events:

Events A and B are *equivalent* if $A \triangle B \in \mathcal{N}$, and we denote this by $A \equiv B$. The relation \equiv is an equivalence relation on \mathcal{S} . That is, for $A, B, C \in \mathcal{S}$,

1. $A \equiv A$ (the *reflexive property*).
2. If $A \equiv B$ then $B \equiv A$ (the *symmetric property*).
3. If $A \equiv B$ and $B \equiv C$ then $A \equiv C$ (the *transitive property*).

Thus $A \equiv B$ if and only if $\mathbb{P}(A \triangle B) = \mathbb{P}(A \setminus B) + \mathbb{P}(B \setminus A) = 0$ if and only if $\mathbb{P}(A \setminus B) = \mathbb{P}(B \setminus A) = 0$. The equivalence relation \equiv partitions \mathcal{S} into disjoint classes of mutually equivalent events. Equivalence is preserved under the set operations.

Suppose that $A, B \in \mathcal{S}$. If $A \equiv B$ then $A^c \equiv B^c$.

Suppose that $A_i, B_i \in \mathcal{S}$ for i in a countable index set I . If $A_i \equiv B_i$ for $i \in I$ then

1. $\bigcup_{i \in I} A_i \equiv \bigcup_{i \in I} B_i$
2. $\bigcap_{i \in I} A_i \equiv \bigcap_{i \in I} B_i$

Equivalent events have the same probability.

If $A, B \in \mathcal{S}$ and $A \equiv B$ then $\mathbb{P}(A) = \mathbb{P}(B)$.

The converse trivially fails, and a [counterexample](#) is given below. However, the null and almost sure events do form equivalence classes.

Suppose that $A \in \mathcal{S}$.

1. If $A \in \mathcal{N}$ then $A \equiv B$ if and only if $B \in \mathcal{N}$.
2. If $A \in \mathcal{M}$ then $A \equiv B$ if and only if $B \in \mathcal{M}$.

We can extend the notion of equivalence to random variables taking values in the same space. Thus suppose that (T, \mathcal{T}) is another measurable space. If X and Y are random variables with values in T , then (X, Y) is a random variable with values in $T \times T$, which is given the usual product σ -algebra $\mathcal{T} \otimes \mathcal{T}$. We assume that the diagonal set $D = \{(x, x) : x \in T\} \in \mathcal{T} \otimes \mathcal{T}$, which is almost always true in applications.

Random variables X and Y taking values in T are *equivalent* if $\mathbb{P}(X = Y) = 1$. Again we write $X \equiv Y$. The relation \equiv is an equivalence relation on the collection of random variables that take values in T . That is, for random variables X, Y , and Z with values in T ,

1. $X \equiv X$ (the *reflexive property*).
2. If $X \equiv Y$ then $Y \equiv X$ (the *symmetric property*).
3. If $X \equiv Y$ and $Y \equiv Z$ then $X \equiv Z$ (the *transitive property*).

So the collection of random variables with values in T is partitioned into disjoint classes of mutually equivalent variables.

Suppose that X and Y are random variables taking values in T and that $X \equiv Y$. Then

1. $\{X \in B\} \equiv \{Y \in B\}$ for every $B \in \mathcal{T}$.
2. X and Y have the same probability distribution on (T, \mathcal{T}) .

Again, the converse to part (b) fails with a passion, and a [counterexample](#) is given below. It often happens that a definition for random variables subsumes the corresponding definition for events, by considering the indicator variables of the events. So it is with equivalence.

Suppose that $A, B \in \mathcal{S}$. Then $A \equiv B$ if and only if $\mathbf{1}_A \equiv \mathbf{1}_B$.

Equivalence is preserved under a deterministic transformation of the variables. For the next result, suppose that (U, \mathcal{U}) is yet another measurable space, along with (T, \mathcal{T}) .

Suppose X, Y are random variables with values in T and that $g: T \rightarrow U$ is measurable. If $X \equiv Y$ then $g(X) \equiv g(Y)$.

Suppose again that $(S, \mathcal{S}, \mathbb{P})$ is a probability space corresponding to a random experiment. Let \mathcal{V} denote the collection of all real-valued random variables for the experiment, that is, all measurable functions from S into \mathbb{R} . With the usual definitions of addition and scalar multiplication, $(\mathcal{V}, +, \cdot)$ is a vector space. However, in probability theory, we often do not want to distinguish between random variables that are equivalent, so it's nice to know that the vector space structure is preserved when we *identify* equivalent random variables. Formally, let $[X]$ denote the equivalence class generated by a real-valued random variable $X \in \mathcal{V}$, and let \mathcal{W} denote the collection of all such equivalence classes. In modular notation, \mathcal{W} is the set \mathcal{V} / \equiv . We define addition and scalar multiplication on \mathcal{W} by

$$[X] + [Y] = [X + Y], \quad c[X] = [cX]; \quad [X], [Y] \in \mathcal{W}, \quad c \in \mathbb{R} \quad (2.9.2)$$

$(\mathcal{W}, +, \cdot)$ is a vector space.

Often we don't bother to use the special notation for the equivalence class associated with a random variable. Rather, it's understood that equivalent random variables represent the same object. Spaces of functions in a general measure space are studied in the chapter on Distributions, and spaces of random variables are studied in more detail in the chapter on Expected Value.

Completion

Suppose again that $(S, \mathcal{S}, \mathbb{P})$ is a probability space, and that \mathcal{N} denotes the collection of null events, as above. Suppose that $A \in \mathcal{N}$ so that $\mathbb{P}(A) = 0$. If $B \subseteq A$ and $B \in \mathcal{S}$, then we know that $\mathbb{P}(B) = 0$ so $B \in \mathcal{N}$ also. However, in general there might be subsets of A that are not in \mathcal{S} . This leads naturally to the following definition.

The probability space $(S, \mathcal{S}, \mathbb{P})$ is *complete* if $A \in \mathcal{N}$ and $B \subseteq A$ imply $B \in \mathcal{S}$ (and hence $B \in \mathcal{N}$).

So the probability space is complete if every subset of an event with probability 0 is also an event (and hence also has probability 0). We know from our work on positive measures that every σ -finite measure space that is not complete can be completed. So in particular a probability space that is not complete can be completed. To review the construction, recall that the equivalence relation \equiv that we used above on \mathcal{S} is extended to $\mathcal{P}(S)$ (the power set of S).

For $A, B \subseteq S$, define $A \equiv B$ if and only if there exists $N \in \mathcal{N}$ such that $A \triangle B \subseteq N$. The relation \equiv is an equivalence relation on $\mathcal{P}(S)$.

Here is how the probability space is completed:

Let $\mathcal{S}_0 = \{A \subseteq S : A \equiv B \text{ for some } B \in \mathcal{S}\}$. For $A \in \mathcal{S}_0$, define $\mathbb{P}_0(A) = \mathbb{P}(B)$ where $B \in \mathcal{S}$ and $A \equiv B$. Then

1. \mathcal{S}_0 is a σ -algebra of subsets of S and $\mathcal{S} \subseteq \mathcal{S}_0$.
2. \mathbb{P}_0 is a probability measure on (S, \mathcal{S}_0) .
3. $(S, \mathcal{S}_0, \mathbb{P}_0)$ is complete, and is the *completion* of $(S, \mathcal{S}, \mathbb{P})$.

Product Spaces

Our next discussion concerns the construction of probability spaces that correspond to specified distributions. To set the stage, suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space. If we let X denote the identity function on S , so that $X(x) = x$ for $x \in S$, then $\{X \in A\} = A$ for $A \in \mathcal{S}$ and hence $\mathbb{P}(X \in A) = \mathbb{P}(A)$. That is, \mathbb{P} is the probability distribution of X . We have seen this before—every probability measure can be thought of as the distribution of a random variable. The next result shows how to construct a probability space that corresponds to a sequence of independent random variables with specified distributions.

Suppose $n \in \mathbb{N}_+$ and that $(S_i, \mathcal{S}_i, \mathbb{P}_i)$ is a probability space for $i \in \{1, 2, \dots, n\}$. The corresponding product measure space $(S, \mathcal{S}, \mathbb{P})$ is a probability space. If $X_i: S \rightarrow S_i$ is the i th coordinate function on S so that $X_i(\mathbf{x}) = x_i$ for $\mathbf{x} = (x_1, x_2, \dots, x_n) \in S$ then (X_1, X_2, \dots, X_n) is a sequence of independent random variables on $(S, \mathcal{S}, \mathbb{P})$, and X_i has distribution \mathbb{P}_i on (S_i, \mathcal{S}_i) for each $i \in \{1, 2, \dots, n\}$.

Proof

Of course, the existence of the product space $(S, \mathcal{S}, \mathbb{P})$ follows immediately from the more general result for products of positive measure spaces. Recall that $S = \prod_{i=1}^n S_i$ and that \mathcal{S} is the σ -algebra generated by sets of the form $\prod_{i=1}^n A_i$ where $A_i \in \mathcal{S}_i$ for each $i \in \{1, 2, \dots, n\}$. Finally, \mathbb{P} is the unique positive measure on (S, \mathcal{S}) satisfying

$$\mathbb{P}\left(\prod_{i=1}^n A_i\right) = \prod_{i=1}^n \mathbb{P}_i(A_i) \quad (2.9.3)$$

where again, $A_i \in \mathcal{S}_i$ for each $i \in \{1, 2, \dots, n\}$. Clearly \mathbb{P} is a probability measure since $\mathbb{P}(S) = \prod_{i=1}^n \mathbb{P}_i(S_i) = 1$. Suppose that $A_i \in \mathcal{S}_i$ for $i \in \{1, 2, \dots, n\}$. Then $\{X_1 \in A_1, X_2 \in A_2, \dots, X_n \in A_n\} = \prod_{i=1}^n A_i \in \mathcal{S}$. Hence

$$\mathbb{P}(X_1 \in A_1, X_2 \in A_2, \dots, X_n \in A_n) = \prod_{i=1}^n \mathbb{P}_i(A_i) \quad (2.9.4)$$

If we fix $i \in \{1, 2, \dots, n\}$ and let $A_j = S_j$ for $j \neq i$, then the displayed equation give $\mathbb{P}(X_i \in A_i) = \mathbb{P}_i(A_i)$, so X_i has distribution \mathbb{P}_i on (S_i, \mathcal{S}_i) . Returning to the displayed equation we have

$$\mathbb{P}(X_1 \in A_1, X_2 \in A_2, \dots, X_n \in A_n) = \prod_{i=1}^n \mathbb{P}(X_i \in A_i) \quad (2.9.5)$$

so (X_1, X_2, \dots, X_n) are independent.

Intuitively, the given probability spaces correspond to n random experiments. The product space then is the probability space that corresponds to the experiments performed independently. When modeling a random experiment, if we say that we have a finite sequence of independent random variables with specified distributions, we can rest assured that there actually is a probability space that supports this statement

We can extend the last result to an infinite sequence of probability spaces. Suppose that (S_i, \mathcal{S}_i) is a measurable space for each $i \in \mathbb{N}_+$. Recall that the product space $\prod_{i=1}^{\infty} S_i$ consists of all sequences $\mathbf{x} = (x_1, x_2, \dots)$ such that $x_i \in S_i$ for each $i \in \mathbb{N}_+$. The corresponding product σ -algebra \mathcal{S} is generated by the collection of *cylinder sets*. That is, $\mathcal{S} = \sigma(\mathcal{B})$ where

$$\mathcal{B} = \left\{ \prod_{i=1}^{\infty} A_i : A_i \in \mathcal{S}_i \text{ for each } i \in \mathbb{N}_+ \text{ and } A_i = S_i \text{ for all but finitely many } i \in \mathbb{N}_+ \right\} \quad (2.9.6)$$

Suppose that $(S_i, \mathcal{S}_i, \mathbb{P}_i)$ is a probability space for $i \in \mathbb{N}_+$. Let (S, \mathcal{S}) denote the product measurable space so that $\mathcal{S} = \sigma(\mathcal{B})$ where \mathcal{B} is the collection of cylinder sets. Then there exists a unique probability measure \mathbb{P} on (S, \mathcal{S}) that satisfies

$$\mathbb{P}\left(\prod_{i=1}^{\infty} A_i\right) = \prod_{i=1}^{\infty} \mathbb{P}_i(A_i), \quad \prod_{i=1}^{\infty} A_i \in \mathcal{B} \quad (2.9.7)$$

If $X_i : S \rightarrow S_i$ is the i th coordinate function on S for $i \in \mathbb{N}_+$, so that $X_i(\mathbf{x}) = x_i$ for $\mathbf{x} = (x_1, x_2, \dots) \in S$, then (X_1, X_2, \dots) is a sequence of independent random variables on $(S, \mathcal{S}, \mathbb{P})$, and X_i has distribution \mathbb{P}_i on (S_i, \mathcal{S}_i) for each $i \in \mathbb{N}_+$.

Proof

The proof is similar to the one in for positive measure spaces in the section on existence and uniqueness. First recall that the collection of cylinder sets \mathcal{B} is a semi-algebra. We define $\mathbb{P} : \mathcal{B} \rightarrow [0, 1]$ as in the statement of the theorem. Note that all but finitely many factors are 1. The consistency conditions are satisfied, so \mathbb{P} can be extended to a probability measure on the algebra \mathcal{A} generated by \mathcal{B} . That is, \mathcal{A} is the collection of all finite, disjoint unions of cylinder sets. The standard extension theorem and uniqueness theorem now apply, so \mathbb{P} can be extended uniquely to a measure on $\mathcal{S} = \sigma(\mathcal{A})$. The proof that (X_1, X_2, \dots) are independent and that X_i has distribution \mathbb{P}_i for each $i \in \mathbb{N}_+$ is just as in the previous theorem.

Once again, if we model a random process by starting with an infinite sequence of independent random variables, we can be sure that there exists a probability space that supports this sequence. The particular probability space constructed in the last theorem is called the *canonical probability space* associated with the sequence of random variables. Note also that it was important that we had *probability measures* rather than just general positive measures in the construction, since the infinite product $\prod_{i=1}^{\infty} \mathbb{P}_i(A_i)$ is always well defined. The next section on Stochastic Processes continues the discussion of how to construct probability spaces that correspond to a collection of random variables with specified distributional properties.

Probability Concepts

Our next discussion concerns topics that are unique to probability theory and do not have simple analogies in general measure theory.

Independence

As usual, suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space. We have already studied the independence of collections of events and the independence of collections of random variables. A more complete and general treatment results if we define the independence of collections of *collections* of events, and most importantly, the independence of collections of σ -algebras. This extension actually occurred already, when we went from independence of a collection of events to independence of a collection of random variables, but we did not note it at the time. In spite of the layers of set theory, the basic idea is the same.

Suppose that \mathcal{A}_i is a collection of events for each i in an index set I . Then $\mathcal{A} = \{\mathcal{A}_i : i \in I\}$ is independent if and only if for every choice of $A_i \in \mathcal{A}_i$ for $i \in I$, the collection of events $\{A_i : i \in I\}$ is independent. That is, for every finite $J \subseteq I$,

$$\mathbb{P}\left(\bigcap_{j \in J} A_j\right) = \prod_{j \in J} \mathbb{P}(A_j) \quad (2.9.8)$$

As noted above, independence of random variables, as we defined previously, is a special case of our new definition.

Suppose that (T_i, \mathcal{T}_i) is a measurable space for each i in an index set I , and that X_i is a random variable taking values in a set T_i for each $i \in I$. The independence of $\{X_i : i \in I\}$ is equivalent to the independence of $\{\sigma(X_i) : i \in I\}$.

Independence of events is also a special case of the new definition, and thus our new definition really does subsume our old one.

Suppose that A_i is an event for each $i \in I$. The independence of $\{A_i : i \in I\}$ is equivalent to the independence of $\{\mathcal{A}_i : i \in I\}$ where $\mathcal{A}_i = \sigma\{A_i\} = \{S, \emptyset, A_i, A_i^c\}$ for each $i \in I$.

For every collection of objects that we have considered (collections of events, collections of random variables, collections of collections of events), the notion of independence has the basic inheritance property.

Suppose that \mathcal{A} is a collection of collections of events.

1. If \mathcal{A} is independent then \mathcal{B} is independent for every $\mathcal{B} \subseteq \mathcal{A}$.
2. If \mathcal{B} is independent for every finite $\mathcal{B} \subseteq \mathcal{A}$ then \mathcal{A} is independent.

Our most important collections are σ -algebras, and so we are most interested in the independence of a collection of σ -algebras. The next result allows us to go from the independence of certain types of collections to the independence of the σ -algebras generated by these collections. To understand the result, you will need to review the definitions and theorems concerning π -systems and λ -systems. The proof uses Dynkin's π - λ theorem, named for Eugene Dynkin.

Suppose that \mathcal{A}_i is a collection of events for each i in an index set I , and that \mathcal{A}_i is a π -system for each $i \in I$. If $\{\mathcal{A}_i : i \in I\}$ is independent, then $\{\sigma(\mathcal{A}_i) : i \in I\}$ is independent.

Proof

In light of the previous result, it suffices to consider a finite set of collections. Thus, suppose that $\{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n\}$ is independent. Now, fix $A_i \in \mathcal{A}_i$ for $i \in \{2, 3, \dots, n\}$ and let $E = \bigcap_{i=2}^n A_i$. Let $\mathcal{L} = \{B \in \mathcal{S} : \mathbb{P}(B \cap E) = \mathbb{P}(B)\mathbb{P}(E)\}$. Trivially $S \in \mathcal{L}$ since $\mathbb{P}(S \cap E) = \mathbb{P}(E) = \mathbb{P}(S)\mathbb{P}(E)$. Next suppose that $A \in \mathcal{L}$. Then

$$\mathbb{P}(A^c \cap E) = \mathbb{P}(E) - \mathbb{P}(A \cap E) = \mathbb{P}(E) - \mathbb{P}(A)\mathbb{P}(E) = [1 - \mathbb{P}(A)]\mathbb{P}(E) = \mathbb{P}(A^c)\mathbb{P}(E) \quad (2.9.9)$$

Thus $A^c \in \mathcal{L}$. Finally, suppose that $\{A_j : j \in J\}$ is a countable collection of disjoint sets in \mathcal{L} . Then

$$\mathbb{P}\left[\left(\bigcup_{j \in J} A_j\right) \cap E\right] = \mathbb{P}\left[\bigcup_{j \in J} (A_j \cap E)\right] = \sum_{j \in J} \mathbb{P}(A_j \cap E) = \sum_{j \in J} \mathbb{P}(A_j)\mathbb{P}(E) = \mathbb{P}(E) \sum_{j \in J} \mathbb{P}(A_j) = \mathbb{P}(E)\mathbb{P}\left(\bigcup_{j \in J} A_j\right) \quad (2.9.10)$$

Therefore $\bigcup_{j \in J} A_j \in \mathcal{L}$ and so \mathcal{L} is a λ -system. Trivially $\mathcal{A}_1 \subseteq \mathcal{L}$ by the original independence assumption, so by the π - λ theorem, $\sigma(\mathcal{A}_1) \subseteq \mathcal{L}$. Thus, we have that for every $A_1 \in \sigma(\mathcal{A}_1)$ and $A_i \in \mathcal{A}_i$ for $i \in \{2, 3, \dots, n\}$,

$$\mathbb{P}\left(\bigcap_{i=1}^n A_i\right) = \prod_{i=1}^n \mathbb{P}(A_i) \quad (2.9.11)$$

Thus we have shown that $\{\sigma(\mathcal{A}_1), \mathcal{A}_2, \dots, \mathcal{A}_n\}$ is independent. Repeating the argument $n-1$ additional times, we get that $\{\sigma(\mathcal{A}_1), \sigma(\mathcal{A}_2), \dots, \sigma(\mathcal{A}_n)\}$ is independent.

The next result is a rigorous statement of the strong independence that is implied the independence of a collection of events.

Suppose that \mathcal{A} is an independent collection of events, and that $\{\mathcal{B}_j : j \in J\}$ is a partition of \mathcal{A} . That is, $\mathcal{B}_j \cap \mathcal{B}_k = \emptyset$ for $j \neq k$ and $\bigcup_{j \in J} \mathcal{B}_j = \mathcal{A}$. Then $\{\sigma(\mathcal{B}_j) : j \in J\}$ is independent.

Proof

Let \mathcal{B}_j^* denote the set of all finite intersections of sets in \mathcal{B}_j , for each $j \in J$. Then clearly \mathcal{B}_j^* is a π -system for each j , and $\{\mathcal{B}_j^* : j \in J\}$ is independent. By the previous theorem, $\{\sigma(\mathcal{B}_j^*) : j \in J\}$ is independent. But clearly $\sigma(\mathcal{B}_j^*) = \sigma(\mathcal{B}_j)$ for $j \in J$.

Let's bring the result down to earth. Suppose that A, B, C, D are independent events. In our elementary discussion of independence, you were asked to show, for example, that $A \cup B^c$ and $C^c \cup D^c$ are independent. This is a consequence of the much stronger statement that the σ -algebras $\sigma\{A, B\}$ and $\sigma\{C, D\}$ are independent.

Exchangeability

As usual, suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space corresponding to a random experiment. Roughly speaking, a sequence of events or a sequence of random variables is *exchangeable* if the probability law that governs the sequence is unchanged when the order of the events or variables is changed. Exchangeable variables arise naturally in sampling experiments and many other settings, and are a natural generalization of a sequence of independent, identically distributed (IID) variables. Conversely, it turns out that any exchangeable sequence of variables can be constructed from an IID sequence. First we give the definition for events:

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a collection of events, where I is a nonempty index set. Then \mathcal{A} is *exchangeable* if the probability of the intersection of a finite number of the events depends only on the number of events. That is, if J and K are finite subsets of I and $\#(J) = \#(K)$ then

$$\mathbb{P}\left(\bigcap_{j \in J} A_j\right) = \mathbb{P}\left(\bigcap_{k \in K} A_k\right) \quad (2.9.12)$$

Exchangeability has the same basic inheritance property that we have seen before.

Suppose that \mathcal{A} is a collection of events.

1. If \mathcal{A} is exchangeable then \mathcal{B} is exchangeable for every $\mathcal{B} \subseteq \mathcal{A}$.
2. Conversely, if \mathcal{B} is exchangeable for every finite $\mathcal{B} \subseteq \mathcal{A}$ then \mathcal{A} is exchangeable.

For a collection of exchangeable events, the inclusion exclusion law for the probability of a union is much simpler than the general version.

Suppose that $\{A_1, A_2, \dots, A_n\}$ is an exchangeable collection of events. For $J \subseteq \{1, 2, \dots, n\}$ with $\#(J) = k$, let $p_k = \mathbb{P}\left(\bigcap_{j \in J} A_j\right)$. Then

$$\mathbb{P}\left(\bigcup_{i=1}^n A_i\right) = \sum_{k=1}^n (-1)^{k-1} \binom{n}{k} p_k \quad (2.9.13)$$

Proof

The inclusion-exclusion rule gives

$$\mathbb{P}\left(\bigcup_{i \in I} A_i\right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mathbb{P}\left(\bigcap_{j \in J} A_j\right) \quad (2.9.14)$$

But $p_k = \mathbb{P}\left(\bigcap_{j \in J} A_j\right)$ for every $J \subseteq \{1, 2, \dots, n\}$ with $\#(J) = k$, and there are $\binom{n}{k}$ such subsets.

The concept of exchangeability can be extended to random variables in the natural way. Suppose that (T, \mathcal{T}) is a measurable space.

Suppose that \mathcal{A} is a collection of random variables, each taking values in T . The collection \mathcal{A} is *exchangeable* if for any $\{X_1, X_2, \dots, X_n\} \subseteq \mathcal{A}$, the distribution of the random vector (X_1, X_2, \dots, X_n) depends only on n .

Thus, the distribution of the random vector is unchanged if the coordinates are permuted. Once again, exchangeability has the same basic inheritance property as a collection of independent variables.

Suppose that \mathcal{A} is a collection of random variables, each taking values in T .

1. If \mathcal{A} is exchangeable then \mathcal{B} is exchangeable for every $\mathcal{B} \subseteq \mathcal{A}$.
2. Conversely, if \mathcal{B} is exchangeable for every finite $\mathcal{B} \subseteq \mathcal{A}$ then \mathcal{A} is exchangeable.

Suppose that \mathcal{A} is a collection of random variables, each taking values in T , and that \mathcal{A} is exchangeable. Then trivially the variables are identically distributed: if $X, Y \in \mathcal{A}$ and $A \in \mathcal{T}$, then $\mathbb{P}(X \in A) = \mathbb{P}(Y \in A)$. Also, the definition of exchangeable variables subsumes the definition for events:

Suppose that \mathcal{A} is a collection of events, and let $\mathcal{B} = \{\mathbf{1}_A : A \in \mathcal{A}\}$ denote the corresponding collection of indicator random variables. Then \mathcal{A} is exchangeable if and only if \mathcal{B} is exchangeable.

Tail Events and Variables

Suppose again that we have a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$.

Suppose that (X_1, X_2, \dots) be a sequence of random variables. The *tail sigma algebra* of the sequence is

$$\mathcal{T} = \bigcap_{n=1}^{\infty} \sigma\{X_n, X_{n+1}, \dots\} \quad (2.9.15)$$

1. An event $B \in \mathcal{T}$ is a *tail event* for the sequence.
2. A random variable Y that is measurable with respect to \mathcal{T} is a *tail random variable* for the sequence.

Informally, a tail event (random variable) is an event (random variable) that can be defined in terms of $\{X_n, X_{n+1}, \dots\}$ for each $n \in \mathbb{N}_+$. The tail sigma algebra for a sequence of events (A_1, A_2, \dots) is defined analogously (or simply let $X_k = \mathbf{1}(A_k)$, the indicator variable of A , for each k). For the following results, you may need to review some of the definitions in the section on Convergence.

Suppose that (A_1, A_2, \dots) is a sequence of events.

1. If the sequence is increasing then $\lim_{n \rightarrow \infty} A_n = \bigcup_{n=1}^{\infty} A_n$ is a tail event of the sequence.
2. If the sequence is decreasing then $\lim_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} A_n$ is a tail event of the sequence.

Proof

1. If the sequence is increasing then $\bigcup_{n=1}^{\infty} A_n = \bigcup_{n=k}^{\infty} A_n \in \sigma\{A_k, A_{k+1}, \dots\}$ for every $k \in \mathbb{N}_+$.
2. If the sequence is decreasing then $\bigcap_{n=1}^{\infty} A_n = \bigcap_{n=k}^{\infty} A_n \in \sigma\{A_k, A_{k+1}, \dots\}$ for every $k \in \mathbb{N}_+$.

Suppose again that (A_1, A_2, \dots) is a sequence of events. Each of the following is a tail event of the sequence:

1. $\limsup_{n \rightarrow \infty} A_n = \bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} A_i$
2. $\liminf_{n \rightarrow \infty} A_n = \bigcup_{n=1}^{\infty} \bigcap_{i=n}^{\infty} A_i$

Proof

1. The events $\bigcup_{i=n}^{\infty} A_i$ are decreasing in n and hence $\limsup_{n \rightarrow \infty} A_n = \lim_{n \rightarrow \infty} \bigcup_{i=n}^{\infty} A_i \in \mathcal{T}$ by the previous result.
2. The events $\bigcap_{i=n}^{\infty} A_i$ are increasing in n and hence $\liminf_{n \rightarrow \infty} A_n = \lim_{n \rightarrow \infty} \bigcap_{i=n}^{\infty} A_i \in \mathcal{T}$ by the previous result.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of real-valued random variables.

1. $\{X_n \text{ converges as } n \rightarrow \infty\}$ is a tail event for \mathbf{X} .
2. $\liminf_{n \rightarrow \infty} X_n$ is a tail random variable for \mathbf{X} .
3. $\limsup_{n \rightarrow \infty} X_n$ is a tail random variable for \mathbf{X} .

Proof

1. The Cauchy criterion for convergence (named for Augustin Cauchy of course) states that X_n converges as $n \rightarrow \infty$ if and only if for every $\epsilon > 0$ there exists $N \in \mathbb{N}_+$ (depending on ϵ) such that if $m, n \geq N$ then $|X_n - X_m| < \epsilon$. In this criterion, we can without loss of generality take ϵ to be rational, and for a given $k \in \mathbb{N}_+$ we can insist that $m, n \geq k$. With these restrictions, the Cauchy criterion is a countable intersection of events, each of which is in $\sigma\{X_k, X_{k+1}, \dots\}$.
2. Recall that $\liminf_{n \rightarrow \infty} X_n = \lim_{n \rightarrow \infty} \inf\{X_k : k \geq n\}$.
3. Similarly, recall that $\limsup_{n \rightarrow \infty} X_n = \lim_{n \rightarrow \infty} \sup\{X_k : k \geq n\}$.

The random variable in part (b) may take the value $-\infty$, and the random variable in (c) may take the value ∞ . From parts (b) and (c) together, note that if $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ on the sample space \mathcal{S} , then X_∞ is a tail random variable for \mathbf{X} .

There are a number of *zero-one laws* in probability. These are theorems that give conditions under which an event will be essentially deterministic; that is, have probability 0 or probability 1. Interestingly, it can sometimes be difficult to determine which of these extremes is actually the case. The following result is the *Kolmogorov zero-one law*, named for Andrey Kolmogorov. It states that an event in the tail σ -algebra of an *independent* sequence will have probability 0 or 1.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is an independent sequence of random variables

1. If B is a tail event for \mathbf{X} then $\mathbb{P}(B) = 0$ or $\mathbb{P}(B) = 1$.
2. If Y is a real-valued tail random variable for \mathbf{X} then Y is constant with probability 1.

Proof

1. By definition $B \in \sigma\{X_{n+1}, X_{n+2}, \dots\}$ for each $n \in \mathbb{N}_+$, and hence $\{X_1, X_2, \dots, X_n, \mathbf{1}_B\}$ is an independent set of random variables. Thus $\{X_1, X_2, \dots, \mathbf{1}_B\}$ is an independent set of random variables. But $B \in \sigma\{X_1, X_2, \dots\}$, so it follows that the event B is independent of itself. Therefore $\mathbb{P}(B) = 0$ or $\mathbb{P}(B) = 1$.

2. The function $y \mapsto \mathbb{P}(Y \leq y)$ on \mathbb{R} is the (cumulative) distribution function of Y . This function is clearly increasing. Moreover, simple applications of the continuity theorems show that it is right continuous and that $\mathbb{P}(Y \leq y) \rightarrow 0$ as $y \rightarrow -\infty$ and $\mathbb{P}(Y \leq y) \rightarrow 1$ as $y \rightarrow \infty$. (Explicit proofs are given in the section on distribution functions in the chapter on Distributions.) But since Y is a tail random variable, $\{Y \leq y\}$ is a tail event and hence $\mathbb{P}(Y \leq y) \in \{0, 1\}$ for each $y \in \mathbb{R}$. It follows that there exists $c \in \mathbb{R}$ such that $\mathbb{P}(Y \leq y) = 0$ for $y < c$ and $\mathbb{P}(Y \leq y) = 1$ for $y \geq c$. Hence $\mathbb{P}(Y = c) = 1$.

From the Komogorov zero-one law and the [result above](#), note that if (A_1, A_2, \dots) is a sequence of independent events, then $\limsup_{n \rightarrow \infty} A_n$ must have probability 0 or 1. The Borel-Cantelli lemmas give conditions for which of these is correct:

Suppose that (A_1, A_2, \dots) is a sequence of independent events.

1. If $\sum_{i=1}^{\infty} \mathbb{P}(A_i) < \infty$ then $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = 0$.
2. If $\sum_{i=1}^{\infty} \mathbb{P}(A_i) = \infty$ then $\mathbb{P}(\limsup_{n \rightarrow \infty} A_n) = 1$.

Another proof of the Kolmogorov zero-one law will be given using the martingale convergence theorem.

Examples and Exercises

As always, be sure to try the computational exercises and proofs yourself before reading the answers and proofs in the text.

Counterexamples

Equal probability certainly does not imply equivalent events.

Consider the simple experiment of tossing a fair coin. The event that the coin lands heads and the event that the coin lands tails have the same probability, but are not equivalent.

Proof

Let S denote the sample space, and H the event of heads, so that H^c is the event of tails. Since the coin is fair, $\mathbb{P}(H) = \mathbb{P}(H^c) = \frac{1}{2}$. But $H \triangle H^c = S$, so $\mathbb{P}(H \triangle H^c) = 1$, so H and H^c are as far from equivalent as possible.

Similarly, equivalent distributions does not imply equivalent random variables.

Consider the experiment of rolling a standard, fair die. Let X denote the score and $Y = 7 - X$. Then X and Y have the same distribution but are not equivalent.

Proof

Since the die is fair, X is uniformly distributed on $S = \{1, 2, 3, 4, 5, 6\}$. Also $\mathbb{P}(Y = k) = \mathbb{P}(X = 7 - k) = \frac{1}{6}$ for $k \in S$, so Y also has the uniform distribution on S . But $\mathbb{P}(X = Y) = \mathbb{P}\left(X = \frac{7}{2}\right) = 0$, so X and Y are as far from equivalent as possible.

Consider the experiment of rolling two standard, fair dice and recording the sequence of scores (X, Y) . Then X and Y are independent and have the same distribution, but are not equivalent.

Proof

Since the dice are fair, (X, Y) has the uniform distribution on $\{1, 2, 3, 4, 5, 6\}^2$. Equivalently, X and Y are independent, and each has the uniform distribution on $\{1, 2, 3, 4, 5, 6\}$. But $\mathbb{P}(X = Y) = \frac{1}{6}$, so X and Y are not equivalent.

This page titled [2.9: Probability Spaces Revisited](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.10: Stochastic Processes

Introduction

This section requires measure theory, so you may need to review the advanced sections in the chapter on Foundations and in this chapter. In particular, recall that a set E almost always comes with a σ -algebra \mathcal{E} of admissible subsets, so that (E, \mathcal{E}) is a measurable space. Usually in fact, E has a topology and \mathcal{E} is the corresponding Borel σ -algebra, that is, the σ -algebra generated by the topology. If E is countable, we almost always take \mathcal{E} to be the collection of all subsets of E , and in this case (E, \mathcal{E}) is a *discrete space*. The other common case is when E is an uncountable measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}$, in which case \mathcal{E} is the collection of measurable subsets of E . If $(E_1, \mathcal{E}_1), (E_2, \mathcal{E}_2), \dots, (E_n, \mathcal{E}_n)$ are measurable spaces for some $n \in \mathbb{N}_+$, then the Cartesian product $E_1 \times E_2 \times \dots \times E_n$ is given the product σ -algebra $\mathcal{E}_1 \otimes \mathcal{E}_2 \otimes \dots \otimes \mathcal{E}_n$. As a special case, the Cartesian power E^n is given the corresponding power σ -algebra \mathcal{E}^n .

With these preliminary remarks out of the way, suppose that $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, so that Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . Suppose also that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. Here is our main definition:

A *random process* or *stochastic process* on $(\Omega, \mathcal{F}, \mathbb{P})$ with *state space* (S, \mathcal{S}) and *index set* T is a collection of random variables $\mathbf{X} = \{X_t : t \in T\}$ such that X_t takes values in S for each $t \in T$.

Sometimes it's notationally convenient to write $X(t)$ instead of X_t for $t \in T$. Often $T = \mathbb{N}$ or $T = [0, \infty)$ and the elements of T are interpreted as *points in time* (discrete time in the first case and continuous time in the second). So then $X_t \in S$ is the *state* of the random process at *time* $t \in T$, and the index space (T, \mathcal{T}) becomes the *time space*.

Since X_t is itself a function from Ω into S , it follows that ultimately, a stochastic process is a function from $\Omega \times T$ into S . Stated another way, $t \mapsto X_t$ is a *random function* on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. To make this precise, recall that S^T is the notation sometimes used for the collection of functions from T into S . Recall also that a natural σ -algebra used for S^T is the one generated by sets of the form

$$\{f \in S^T : f(t) \in A_t \text{ for all } t \in T\}, \text{ where } A_t \in \mathcal{S} \text{ for every } t \in T \text{ and } A_t = S \text{ for all but finitely many } t \in T \quad (2.10.1)$$

This σ -algebra, denoted \mathcal{S}^T , generalizes the ordinary power σ -algebra \mathcal{S}^n mentioned in the opening paragraph and will be important in the discussion of existence below.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with state space (S, \mathcal{S}) and index set T . Then the mapping that takes ω into the function $t \mapsto X_t(\omega)$ is measurable with respect to (Ω, \mathcal{F}) and (S^T, \mathcal{S}^T) .

Proof

Recall that a mapping with values in S^T is measurable if and only if each of its “coordinate functions” is measurable. In the present context that means that we must show that the function X_t is measurable with respect to (Ω, \mathcal{F}) and (S, \mathcal{S}) for each $t \in T$. But of course, that follows from the very meaning of the term *random variable*.

For $\omega \in \Omega$, the function $t \mapsto X_t(\omega)$ is known as a *sample path* of the process. So S^T , the set of functions from T into S , can be thought of as a *set of outcomes* of the stochastic process \mathbf{X} , a point we will return to in our discussion of existence below.

As noted in the proof of the last theorem, X_t is a measurable function from Ω into S for each $t \in T$, by the very meaning of the term *random variable*. But it does not follow in general that $(\omega, t) \mapsto X_t(\omega)$ is measurable as a function from $\Omega \times T$ into S . In fact, the σ -algebra on T has played no role in our discussion so far. Informally, a statement about X_t for a *fixed* $t \in T$ or even a statement about X_t for countably many $t \in T$ defines an event. But it does not follow that a statement about X_t for uncountably many $t \in T$ defines an event. We often want to make such statements, so the following definition is inevitable:

A stochastic process $\mathbf{X} = \{X_t : t \in T\}$ defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and with index space (T, \mathcal{T}) and state space (S, \mathcal{S}) is *measurable* if $(\omega, t) \mapsto X_t(\omega)$ is a measurable function from $\Omega \times T$ into S .

Every stochastic process indexed by a countable set T is measurable, so the definition is only important when T is uncountable, and in particular for $T = [0, \infty)$.

Equivalent Processes

Our next goal is to study different ways that two stochastic processes, with the same state and index spaces, can be “equivalent”. We will assume that the diagonal $D = \{(x, x) : x \in S\} \in \mathcal{S}^2$, an assumption that almost always holds in applications, and in particular for the discrete and Euclidean spaces that are most important to us. Sufficient conditions are that \mathcal{S} have a sub σ -algebra that is countably generated and contains all of the singleton sets, properties that hold for the Borel σ -algebra when the topology on S is locally compact, Hausdorff, and has a countable base.

First, we often feel that we understand a random process $\mathbf{X} = \{X_t : t \in T\}$ well if we know the *finite dimensional distributions*, that is, if we know the distribution of $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ for every choice of $n \in \mathbb{N}_+$ and $(t_1, t_2, \dots, t_n) \in T^n$. Thus, we can compute $\mathbb{P}[(X_{t_1}, X_{t_2}, \dots, X_{t_n}) \in A]$ for

every $n \in \mathbb{N}_+$, $(t_1, t_2, \dots, t_n) \in T^n$, and $A \in \mathcal{S}^n$. Using various rules of probability, we can compute the probabilities of many events involving infinitely many values of the index parameter t as well. With this idea in mind, we have the following definition:

Random processes $\mathbf{X} = \{X_t : t \in T\}$ and $\mathbf{Y} = \{Y_t : t \in T\}$ with state space (S, \mathcal{S}) and index set T are *equivalent in distribution* if they have the same finite dimensional distributions. This defines an equivalence relation on the collection of stochastic processes with this state space and index set. That is, if \mathbf{X} , \mathbf{Y} , and \mathbf{Z} are such processes then

1. \mathbf{X} is equivalent in distribution to \mathbf{X} (the *reflexive property*)
2. If \mathbf{X} is equivalent in distribution to \mathbf{Y} then \mathbf{Y} is equivalent in distribution to \mathbf{X} (the *symmetric property*)
3. If \mathbf{X} is equivalent in distribution to \mathbf{Y} and \mathbf{Y} is equivalent in distribution to \mathbf{Z} then \mathbf{X} is equivalent in distribution to \mathbf{Z} (the *transitive property*)

Note that since only the finite-dimensional distributions of the processes \mathbf{X} and \mathbf{Y} are involved in the definition, the processes need not be defined on the same probability space. Thus, *equivalence in distribution* partitions the collection of all random processes with a given state space and index set into mutually disjoint equivalence classes. But of course, we already know that two random variables can have the same *distribution* but be very different as variables (functions on the sample space). Clearly, the same statement applies to random processes.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter $p = \frac{1}{2}$. Let $Y_n = 1 - X_n$ for $n \in \mathbb{N}_+$. Then $\mathbf{Y} = (Y_1, Y_2, \dots)$ is equivalent in distribution to \mathbf{X} but

$$\mathbb{P}(X_n \neq Y_n \text{ for every } n \in \mathbb{N}_+) = 1 \quad (2.10.2)$$

Proof

By the meaning of *Bernoulli trials*, \mathbf{X} is a sequence of independent indicator random variables with $\mathbb{P}(X_n = 1) = \frac{1}{2}$ for each $n \in \mathbb{N}_+$. It follows that \mathbf{Y} is also a Bernoulli trials sequence with success parameter $\frac{1}{2}$, so \mathbf{X} and \mathbf{Y} are equivalent in distribution. Also, of course, the state set is $\{0, 1\}$ and $Y_n = 1$ if and only if $X_n = 0$.

Motivated by this example, let's look at another, stronger way that random processes can be equivalent. First recall that random variables X and Y on $(\Omega, \mathcal{F}, \mathbb{P})$, with values in S , are *equivalent* if $\mathbb{P}(X = Y) = 1$.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ and $\mathbf{Y} = \{Y_t : t \in T\}$ are stochastic processes defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and both with state space (S, \mathcal{S}) and index set T . Then \mathbf{Y} is a *version* of \mathbf{X} if Y_t is equivalent to X_t (so that $\mathbb{P}(X_t = Y_t) = 1$) for every $t \in T$. This defines an equivalence relation on the collection of stochastic processes on the same probability space and with the same state space and index set. That is, if \mathbf{X} , \mathbf{Y} , and \mathbf{Z} are such processes then

1. \mathbf{X} is a version of \mathbf{X} (the *reflexive property*)
2. If \mathbf{X} is a version of \mathbf{Y} then \mathbf{Y} is a version of \mathbf{X} (the *symmetric property*)
3. If \mathbf{X} is a version of \mathbf{Y} and \mathbf{Y} is of \mathbf{Z} then \mathbf{X} is a version of \mathbf{Z} (the *transitive property*)

Proof

Note that (X_t, Y_t) is a random variable with values in S^2 (and so the function $\omega \mapsto (X_t(\omega), Y_t(\omega))$ is measurable). The event $\{X_t = Y_t\}$ is the inverse image of the diagonal $D \in \mathcal{S}^2$ under this mapping, and so the definition makes sense.

So the *version of* relation partitions the collection of stochastic processes on a given probability space and with a given state space and index set into mutually disjoint equivalence classes.

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ and $\mathbf{Y} = \{Y_t : t \in T\}$ are random processes on $(\Omega, \mathcal{F}, \mathbb{P})$ with state space (S, \mathcal{S}) and index set T . If \mathbf{Y} is a version of \mathbf{X} then \mathbf{Y} and \mathbf{X} are equivalent in distribution.

Proof

Suppose that $(t_1, t_2, \dots, t_n) \in T^n$ and that $A \in \mathcal{S}^n$. Recall that the intersection of a finite (or even countably infinite) collection of events with probability 1 still has probability 1. Hence

$$\mathbb{P}[(X_{t_1}, X_{t_2}, \dots, X_{t_n}) \in A] = \mathbb{P}[(X_{t_1}, X_{t_2}, \dots, X_{t_n}) \in A, X_{t_1} = Y_{t_1}, X_{t_2} = Y_{t_2}, \dots, X_{t_n} = Y_{t_n}] \quad (2.10.3)$$

$$= \mathbb{P}[(Y_{t_1}, Y_{t_2}, \dots, Y_{t_n}) \in A, X_{t_1} = Y_{t_1}, X_{t_2} = Y_{t_2}, \dots, X_{t_n} = Y_{t_n}] = \mathbb{P}[(Y_{t_1}, Y_{t_2}, \dots, Y_{t_n}) \in A] \quad (2.10.4)$$

As noted in the proof, a countable intersection of events with probability 1 still has probability 1. Hence if T is countable and random processes \mathbf{X} is a version of \mathbf{Y} then

$$\mathbb{P}(X_t = Y_t \text{ for all } t \in T) = 1 \quad (2.10.5)$$

so \mathbf{X} and \mathbf{Y} really are essentially the same random process. But when T is uncountable the result in the displayed equation may not be true, and \mathbf{X} and \mathbf{Y} may be very different as random functions on T . Here is a simple example:

Suppose that $\Omega = T = [0, \infty)$, $\mathcal{F} = \mathcal{S}$ is the σ -algebra of Borel measurable subsets of $[0, \infty)$, and \mathbb{P} is any continuous probability measure on (Ω, \mathcal{F}) . Let $S = \{0, 1\}$ (with all subsets measurable, of course). For $t \in T$ and $\omega \in \Omega$, define $X_t(\omega) = \mathbf{1}_t(\omega)$ and $Y_t(\omega) = 0$. Then $\mathbf{X} = \{X_t : t \in T\}$ is a version of $\mathbf{Y} = \{Y_t : t \in T\}$, but $\mathbb{P}(X_t = Y_t \text{ for all } t \in T) = 0$.

Proof

For $t \in [0, \infty)$, $\mathbb{P}(X_t \neq Y_t) = \mathbb{P}\{t\} = 0$ since P is a continuous measure. But $\{\omega \in \Omega : X_t(\omega) = Y_t(\omega) \text{ for all } t \in T\} = \emptyset$.

Motivated by this example, we have our strongest form of equivalence:

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ and $\mathbf{Y} = \{Y_t : t \in T\}$ are measurable random processes on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and with state space (S, \mathcal{S}) and index space (T, \mathcal{T}) . Then \mathbf{X} is *indistinguishable from* \mathbf{Y} if $\mathbb{P}(X_t = Y_t \text{ for all } t \in T) = 1$. This defines an equivalence relation on the collection of measurable stochastic processes defined on the same probability space and with the same state and index spaces. That is, if \mathbf{X} , \mathbf{Y} , and \mathbf{Z} are such processes then

1. \mathbf{X} is indistinguishable from \mathbf{X} (the *reflexive property*)
2. If \mathbf{X} is indistinguishable from \mathbf{Y} then \mathbf{Y} is indistinguishable from \mathbf{X} (the *symmetric property*)
3. If \mathbf{X} is indistinguishable from \mathbf{Y} and \mathbf{Y} is indistinguishable from \mathbf{Z} then \mathbf{X} is indistinguishable from \mathbf{Z} (the *transitive property*)

Details

The measurability requirement for the stochastic processes is needed to ensure that $\{X_t = Y_t \text{ for all } t \in T\}$ is a valid event. To see this, note that $(\omega, t) \mapsto (X_t(\omega), Y_t(\omega))$ is measurable, as a function from $\Omega \times T$ into S^2 . As before, let $D = \{(x, x) : x \in S\}$ denote the diagonal. Then $D^c \in \mathcal{S}^2$ and the inverse image of D^c under our mapping is

$$\{(\omega, t) \in \Omega \times T : X_t(\omega) \neq Y_t(\omega)\} \in \mathcal{F} \otimes \mathcal{T} \quad (2.10.6)$$

The projection of this set onto Ω

$$\{\omega \in \Omega : X_t(\omega) \neq Y_t(\omega) \text{ for some } t \in T\} \in \mathcal{F} \quad (2.10.7)$$

since the projection of a measurable set in the product space is also measurable. Hence the complementary event

$$\{\omega \in \Omega : X_t(\omega) = Y_t(\omega) \text{ for all } t \in T\} \in \mathcal{F} \quad (2.10.8)$$

So the *indistinguishable from* relation partitions the collection of measurable stochastic processes on a given probability space and with given state space and index space into mutually disjoint equivalence classes. Trivially, if \mathbf{X} is indistinguishable from \mathbf{Y} , then \mathbf{X} is a version of \mathbf{Y} . As noted above, when T is countable, the converse is also true, but not, as our previous example shows, when T is uncountable. So to summarize, *indistinguishable from* implies *version of* implies *equivalent in distribution*, but none of the converse implications hold in general.

The Kolmogorov Construction

In applications, a stochastic process is often modeled by giving various distributional properties that the process should satisfy. So the basic existence problem is to construct a process that has these properties. More specifically, how can we construct random processes with specified finite dimensional distributions? Let's start with the simplest case, one that we have seen several times before, and build up from there. Our simplest case is to construct a single random variable with a specified distribution.

Suppose that (S, \mathcal{S}, P) is a probability space. Then there exists a random variable X on probability space $(\Omega, \mathcal{F}, \mathbb{P})$ such that X takes values in S and has distribution P .

Proof

The proof is utterly trivial. Let $(\Omega, \mathcal{F}, \mathbb{P}) = (S, \mathcal{S}, P)$ and define $X : \Omega \rightarrow S$ by $X(\omega) = \omega$, so that X is the identity function. Then $\{X \in A\} = A$ and so $\mathbb{P}(X \in A) = P(A)$ for $A \in \mathcal{S}$.

In spite of its triviality the last result contains the seeds of everything else we will do in this discussion. Next, let's see how to construct a sequence of independent random variables with specified distributions.

Suppose that P_i is a probability measure on the measurable space (S, \mathcal{S}) for $i \in \mathbb{N}_+$. Then there exists an independent sequence of random variables (X_1, X_2, \dots) on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ such that X_i takes values in S and has distribution P_i for $i \in \mathbb{N}_+$.

Proof

Let $\Omega = S^\infty = S \times S \times \dots$. Next let $\mathcal{F} = \mathcal{S}^\infty$, the corresponding product σ -algebra. Recall that this is the σ -algebra generated by sets of the form

$$A_1 \times A_2 \times \dots \text{ where } A_i \in \mathcal{S} \text{ for each } i \in I \text{ and } A_i = S \text{ for all but finitely many } i \in I \quad (2.10.9)$$

Finally, let $\mathbb{P} = P_1 \otimes P_2 \otimes \dots$, the corresponding product measure on (Ω, \mathcal{F}) . Recall that this is the unique probability measure that satisfies

$$\mathbb{P}(A_1 \times A_2 \times \dots) = P_1(A_1)P_2(A_2) \dots \quad (2.10.10)$$

where $A_1 \times A_2 \times \dots$ is a set of the type in the first displayed equation. Now define X_i on Ω by $X_i(\omega_1, \omega_2, \dots) = \omega_i$, for $i \in \mathbb{N}_+$, so that X_i is simply the coordinate function for index i . If $A_1 \times A_2 \times \dots$ is a set of the type in the first displayed equation then

$$\{X_1 \in A_1, X_2 \in A_2, \dots\} = A_1 \times A_2 \times \dots \quad (2.10.11)$$

and so by the definition of the product measure,

$$\mathbb{P}(X_1 \in A_1, X_2 \in A_2, \dots) = P_1(A_1)P_2(A_2) \dots \quad (2.10.12)$$

It follows that (X_1, X_2, \dots) is a sequence of independent variables and that X_i has distribution P_i for $i \in \mathbb{N}$.

If you looked at the proof of the last two results you might notice that the last result can be viewed as a special case of the one before, since $\mathbf{X} = (X_1, X_2, \dots)$ is simply the identity function on $\Omega = S^\infty$. The important step is the existence of the product measure \mathbb{P} on (Ω, \mathcal{F}) .

The full generalization of these results is known as the *Kolmogorov existence theorem* (named for Andrei Kolmogorov). We start with the state space (S, \mathcal{S}) and the index set T . The theorem states that if we specify the finite dimensional distributions in a consistent way, then there exists a stochastic process defined on a suitable probability space that has the given finite dimensional distributions. The consistency condition is a bit clunky to state in full generality, but the basic idea is very easy to understand. Suppose that s and t are distinct elements in T and that we specify the distribution (probability measure) P_s of X_s , P_t of X_t , $P_{s,t}$ of (X_s, X_t) , and $P_{t,s}$ of (X_t, X_s) . Then clearly we must specify these so that

$$P_s(A) = P_{s,t}(A \times S), \quad P_t(B) = P_{t,s}(S \times B) \quad (2.10.13)$$

For all $A, B \in \mathcal{S}$. Clearly we also must have $P_{s,t}(C) = P_{t,s}(C')$ for all measurable $C \in \mathcal{S}^2$, where $C' = \{(y, x) : (x, y) \in C\}$.

To state the consistency conditions in general, we need some notation. For $n \in \mathbb{N}_+$, let $T^{(n)} \subset T^n$ denote the set of n -tuples of distinct elements of T , and let $\mathbf{T} = \bigcup_{n=1}^{\infty} T^{(n)}$ denote the set of all finite sequences of distinct elements of T . If $n \in \mathbb{N}_+$, $\mathbf{t} = (t_1, t_2, \dots, t_n) \in T^{(n)}$ and π is a permutation of $\{1, 2, \dots, n\}$, let $\mathbf{t}\pi$ denote the element of $T^{(n)}$ with coordinates $(\mathbf{t}\pi)_i = t_{\pi(i)}$. That is, we permute the coordinates of \mathbf{t} according to π . If $C \in \mathcal{S}^n$, let

$$\pi C = \{(x_1, x_2, \dots, x_n) \in S^n : (x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(n)}) \in C\} \in \mathcal{S}^n \quad (2.10.14)$$

finally, if $n > 1$, let \mathbf{t}_- denote the vector $(t_1, t_2, \dots, t_{n-1}) \in T^{(n-1)}$

Now suppose that $P_{\mathbf{t}}$ is a probability measure on (S^n, \mathcal{S}^n) for each $n \in \mathbb{N}_+$ and $\mathbf{t} \in T^{(n)}$. The idea, of course, is that we want the collection $\mathcal{P} = \{P_{\mathbf{t}} : \mathbf{t} \in \mathbf{T}\}$ to be the finite dimensional distributions of a random process with index set T and state space (S, \mathcal{S}) . Here is the critical definition:

The collection of probability distributions \mathcal{P} relative to T and (S, \mathcal{S}) is *consistent* if

1. $P_{\mathbf{t}\pi}(C) = P_{\mathbf{t}}(\pi C)$ for every $n \in \mathbb{N}_+$, $\mathbf{t} \in T^{(n)}$, permutation π of $\{1, 2, \dots, n\}$, and measurable $C \subseteq S^n$.
2. $P_{\mathbf{t}_-}(C) = P_{\mathbf{t}}(C \times S)$ for every $n > 1$, $\mathbf{t} \in T^{(n)}$, and measurable $C \subseteq S^{n-1}$

With the proper definition of consistence, we can state the fundamental theorem.

Kolmogorov Existence Theorem. If \mathcal{P} is a consistent collection of probability distributions relative to the index set T and the state space (S, \mathcal{S}) , then there exists a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on this probability space such that \mathcal{P} is the collection of finite dimensional distribution of \mathbf{X} .

Proof sketch

Let $\Omega = S^T$, the set of functions from T to S . Such functions are the *outcomes* of the stochastic process. Let $\mathcal{F} = \mathcal{S}^T$, the product σ -algebra, generated by sets of the form

$$B = \{\omega \in \Omega : \omega(t) \in A_t \text{ for all } t \in T\} \quad (2.10.15)$$

where $A_t \in \mathcal{S}$ for all $t \in T$ and $A_t = S$ for all but finitely many $t \in T$. We know how our desired probability measure \mathbb{P} should work on the sets that generate \mathcal{F} . Specifically, suppose that B is a set of the type in the displayed equation, and $A_t = S$ except for $\mathbf{t} = (t_1, t_2, \dots, t_n) \in T^{(n)}$. Then we want

$$\mathbb{P}(B) = P_{\mathbf{t}}(A_{t_1} \times A_{t_2} \times \dots \times A_{t_n}) \quad (2.10.16)$$

Basic existence and uniqueness theorems in measure theory that we discussed earlier, and the consistency of \mathcal{P} , guarantee that \mathbb{P} can be extended to a probability measure on all of \mathcal{F} . Finally, for $t \in T$ we define $X_t : \Omega \rightarrow S$ by $X_t(\omega) = \omega(t)$ for $\omega \in \Omega$, so that X_t is simply the coordinate function of index t . Thus, we have a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ with state space (S, \mathcal{S}) , defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with \mathcal{P} as the collection of finite dimensional distributions.

Note that except for the more complicated notation, the construction is very similar to the one for a [sequence of independent variables](#). Again, \mathbf{X} is essentially the identity function on $\Omega = S^T$. The important and more difficult part is the construction of the probability measure \mathbb{P} on (Ω, \mathcal{F}) .

Applications

Our last discussion is a summary of the stochastic processes that are studied in this text. All are classics and are immensely important in applications.

Random processes are associated with Bernoulli trials include

1. the Bernoulli trials sequence itself
2. the sequence of binomial variables

3. the sequence of geometric variables
4. the sequence of negative binomial variables
5. the simple random walk

Construction

The Bernoulli trials sequence in (a) is a sequence of independent, identically distributed indicator random variables, and so can be constructed as in (). The random processes in (b)–(e) are constructed from the Bernoulli trials sequence.

Random process associated with the Poisson model include

1. the sequence of inter-arrival times
2. the sequence of arrival times
3. the counting process on $[0, \infty)$, both in the homogeneous and non-homogeneous cases.
4. A compound Poisson process.
5. the counting process on a general measure space

Constructions

The random process in (a) is a sequence of independent random variable with a common exponential distribution, and so can be constructed as in (). The processes in (b) and (c) can be constructed from the sequence in (a).

Random processes associated with renewal theory include

1. the sequence of inter-arrival times
2. the sequence of arrival times
3. the counting process on $[0, \infty)$

Markov chains form a very important family of random processes as do Brownian motion and related processes. We will study these in subsequent chapters.

This page titled [2.10: Stochastic Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

2.11: Filtrations and Stopping Times

Introduction

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process with state space (S, \mathcal{S}) defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$. To review, Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on (S, \mathcal{S}) . Also S is the set of states, and \mathcal{S} the σ -algebra of admissible subsets of S . Usually, S is a topological space and \mathcal{S} the Borel σ -algebra generated by the open subsets of S . A standard set of assumptions is that the topology is locally compact, Hausdorff, and has a countable base, which we will abbreviate by *LCCB*. For the index set, we assume that either $T = \mathbb{N}$ or that $T = [0, \infty)$ and as usual in these cases, we interpret the elements of T as points of time. The set T is also given a topology, the discrete topology in the first case and the standard Euclidean topology in the second case, and then the Borel σ -algebra \mathcal{T} . So in discrete time with $T = \mathbb{N}$, $\mathcal{T} = \mathcal{P}(T)$, the power set of T , so every subset of T is measurable, as is every function from T into a another measurable space. Finally, X_t is a random variable and so by definition is measurable with respect to \mathcal{F} and \mathcal{S} for each $t \in T$. We interpret X_t is the state of some random system at time $t \in T$. Many important concepts involving \mathbf{X} are based on how the *future* behavior of the process depends on the *past* behavior, relative to a given *current* time.

For $t \in T$, let $\mathcal{F}_t = \sigma\{X_s : s \in T, s \leq t\}$, the σ -algebra of events that can be defined in terms of the process up to time t . Roughly speaking, for a given $A \in \mathcal{F}_t$, we can tell whether or not A has occurred if we are allowed to observe the process up to time t . The family of σ -algebras $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ has two critical properties: the family is increasing in $t \in T$, relative to the subset partial order, and all of the σ -algebras are sub σ -algebras of \mathcal{F} . That is for $s, t \in T$ with $s \leq t$, we have $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$.

Filtrations

Basic Definitions

Sometimes we need σ -algebras that are a bit larger than the ones in the last paragraph. For example, there may be other random variables that we get to observe, as time goes by, besides the variables in \mathbf{X} . Sometimes, particularly in continuous time, there are technical reasons for somewhat different σ -algebras. Finally, we may want to describe how our information grows, as a family of σ -algebras, without reference to a random process. For the remainder of this section, we have a fixed measurable space (Ω, \mathcal{F}) which we again think of as a sample space, and the time space (T, \mathcal{T}) as described above.

A family of σ -algebras $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a *filtration* on (Ω, \mathcal{F}) if $s, t \in T$ and $s \leq t$ imply $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$. The object $(\Omega, \mathcal{F}, \mathfrak{F})$ is a *filtered sample space*. If \mathbb{P} is a probability measure on (Ω, \mathcal{F}) , then $(\Omega, \mathcal{F}, \mathfrak{F}, \mathbb{P})$ is a *filtered probability space*.

So a filtration is simply an increasing family of sub- σ -algebras of \mathcal{F} , indexed by T . We think of \mathcal{F}_t as the σ -algebra of events up to time $t \in T$. The larger the σ -algebras in a filtration, the more events that are available, so the following relation on filtrations is natural.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ and $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ are filtrations on (Ω, \mathcal{F}) . We say that \mathfrak{F} is *coarser than* \mathfrak{G} and \mathfrak{G} is *finer than* \mathfrak{F} , and we write $\mathfrak{F} \preceq \mathfrak{G}$, if $\mathcal{F}_t \subseteq \mathcal{G}_t$ for all $t \in T$. The relation \preceq is a partial order on the collection of filtrations on (Ω, \mathcal{F}) . That is, if $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$, and $\mathfrak{H} = \{\mathcal{H}_t : t \in T\}$ are filtrations then

1. $\mathfrak{F} \preceq \mathfrak{F}$, the *reflexive property*.
2. If $\mathfrak{F} \preceq \mathfrak{G}$ and $\mathfrak{G} \preceq \mathfrak{F}$ then $\mathfrak{F} = \mathfrak{G}$, the *antisymmetric property*.
3. If $\mathfrak{F} \preceq \mathfrak{G}$ and $\mathfrak{G} \preceq \mathfrak{H}$ then $\mathfrak{F} \preceq \mathfrak{H}$, the *transitive property*.

Proof

The proof is a simple consequence of the fact that the subset relation defines a partial order.

1. $\mathcal{F}_t \subseteq \mathcal{F}_t$ for each $t \in T$ so $\mathfrak{F} \preceq \mathfrak{F}$.
2. If $\mathfrak{F} \preceq \mathfrak{G}$ and $\mathfrak{G} \preceq \mathfrak{F}$ then $\mathcal{F}_t \subseteq \mathcal{G}_t$ and $\mathcal{G}_t \subseteq \mathcal{F}_t$ for each $t \in T$. Hence $\mathcal{F}_t = \mathcal{G}_t$ for each $t \in T$ and so $\mathfrak{F} = \mathfrak{G}$.
3. If $\mathfrak{F} \preceq \mathfrak{G}$ and $\mathfrak{G} \preceq \mathfrak{H}$ then $\mathcal{F}_t \subseteq \mathcal{G}_t$ and $\mathcal{G}_t \subseteq \mathcal{H}_t$ for each $t \in T$. Hence $\mathcal{F}_t \subseteq \mathcal{H}_t$ for each $t \in T$ and so $\mathfrak{F} \preceq \mathfrak{H}$.

So the coarsest filtration on (Ω, \mathcal{F}) is the one where $\mathcal{F}_t = \{\Omega, \emptyset\}$ for every $t \in T$ while the finest filtration is the one where $\mathcal{F}_t = \mathcal{F}$ for every $t \in T$. In the first case, we gain no information as time evolves, and in the second case, we have complete information from the beginning of time. Usually neither of these is realistic.

It's also natural to consider the σ -algebra that encodes our information over *all* time.

For a filtration $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ on (Ω, \mathcal{F}) , define $\mathcal{F}_\infty = \sigma(\bigcup\{\mathcal{F}_t : t \in T\})$. Then

1. $\mathcal{F}_\infty = \sigma(\bigcup\{\mathcal{F}_t : t \in T, t \geq s\})$ for $s \in T$.
2. $\mathcal{F}_t \subseteq \mathcal{F}_\infty$ for $t \in T$.

Proof

These results follow since the σ -algebras in a filtration are increasing in time.

Of course, it may be the case that $\mathcal{F}_\infty = \mathcal{F}$, but not necessarily. Recall that the intersection of a collection of σ -algebras on (Ω, \mathcal{F}) is another σ -algebra. We can use this to create new filtrations from a collection of given filtrations.

Suppose that $\mathcal{F}_i = \{\mathcal{F}_t^i : t \in T\}$ is a filtration on (Ω, \mathcal{F}) for each i in a nonempty index set I . Then $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ where $\mathcal{F}_t = \bigcap_{i \in I} \mathcal{F}_t^i$ for $t \in T$ is also a filtration on (Ω, \mathcal{F}) . This filtration is sometimes denoted $\mathcal{F} = \bigwedge_{i \in I} \mathcal{F}_i$, and is the finest filtration that is coarser than \mathcal{F}_i for every $i \in I$.

Proof

Suppose $s, t \in T$ with $s \leq t$. Then $\mathcal{F}_s^i \subseteq \mathcal{F}_t^i \subseteq \mathcal{F}$ for each $i \in I$ so it follows that $\bigcap_{i \in I} \mathcal{F}_s^i \subseteq \bigcap_{i \in I} \mathcal{F}_t^i \subseteq \mathcal{F}$.

Unions of σ -algebras are not in general σ -algebras, but we can construct a new filtration from a given collection of filtrations using unions in a natural way.

Suppose again that $\mathcal{F}_i = \{\mathcal{F}_t^i : t \in T\}$ is a filtration on (Ω, \mathcal{F}) for each i in a nonempty index set I . Then $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ where $\mathcal{F}_t = \sigma(\bigcup_{i \in I} \mathcal{F}_t^i)$ for $t \in T$ is also a filtration on (Ω, \mathcal{F}) . This filtration is sometimes denoted $\mathcal{F} = \bigvee_{i \in I} \mathcal{F}_i$, and is the coarsest filtration that is finer than \mathcal{F}_i for every $i \in I$.

Proof

Suppose $s, t \in T$ with $s \leq t$. Then $\mathcal{F}_s^i \subseteq \mathcal{F}_t^i \subseteq \mathcal{F}$ for each $i \in I$ so it follows that $\bigcup_{i \in I} \mathcal{F}_s^i \subseteq \bigcup_{i \in I} \mathcal{F}_t^i \subseteq \mathcal{F}$, and hence $\sigma(\bigcup_{i \in I} \mathcal{F}_s^i) \subseteq \sigma(\bigcup_{i \in I} \mathcal{F}_t^i) \subseteq \mathcal{F}$.

Stochastic Processes

Note again that we can have a filtration without an underlying stochastic process in the background. However, we usually *do* have a stochastic process $\mathbf{X} = \{X_t : t \in T\}$, and in this case the filtration $\mathcal{F}^0 = \{\mathcal{F}_t^0 : t \in T\}$ where $\mathcal{F}_t^0 = \sigma\{X_s : s \in T, s \leq t\}$ is the *natural filtration* associated with \mathbf{X} . More generally, the following definition is appropriate.

A stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on (Ω, \mathcal{F}) is *adapted* to a filtration $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ on (Ω, \mathcal{F}) if X_t is measurable with respect to \mathcal{F}_t for each $t \in T$.

Equivalently, \mathbf{X} is adapted to \mathcal{F} if \mathcal{F} is finer than \mathcal{F}^0 , the natural filtration associated with \mathbf{X} . That is, $\sigma\{X_s : s \in T, s \leq t\} \subseteq \mathcal{F}_t$ for each $t \in T$. So clearly, if \mathbf{X} is adapted to a filtration, then it is adapted to any finer filtration, and \mathcal{F}^0 is the coarsest filtration to which \mathbf{X} is adapted. The basic idea behind the definition is that if the filtration \mathcal{F} encodes our information as time goes by, then the process \mathbf{X} is observable. In discrete time, there is a related definition.

Suppose that $T = \mathbb{N}$. A stochastic process $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is *predictable* by the filtration $\mathcal{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ if X_{n+1} is measurable with respect to \mathcal{F}_n for all $n \in \mathbb{N}$.

Clearly if \mathbf{X} is predictable by \mathcal{F} then \mathbf{X} is adapted to \mathcal{F} . But predictable is better than adapted, in the sense that if \mathcal{F} encodes our information as time goes by, then we can look one step into the future in terms of \mathbf{X} : at time n we can determine X_{n+1} . The concept of predictability can be extended to continuous time, but the definition is much more complicated.

Note that ultimately, a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ with sample space (Ω, \mathcal{F}) and state space (S, \mathcal{S}) can be viewed as a function from $\Omega \times T$ into S , so $X_t(\omega) \in S$ is the state at time $t \in T$ corresponding to the outcome $\omega \in \Omega$. By definition, $\omega \mapsto X_t(\omega)$ is measurable for each $t \in T$, but it is often necessary for the process to be jointly measurable in ω and t .

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process with sample space (Ω, \mathcal{F}) and state space (S, \mathcal{S}) . Then \mathbf{X} is *measurable* if $\mathbf{X} : \Omega \times T \rightarrow S$ is measurable with respect to $\mathcal{F} \otimes \mathcal{T}$ and \mathcal{S} .

When we have a filtration, as we usually do, there is a stronger condition that is natural. Let $T_t = \{s \in T : s \leq t\}$ for $t \in T$, and let $\mathcal{F}_t = \{A \cap T_t : A \in \mathcal{F}\}$ be the corresponding induced σ -algebra.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process with sample space (Ω, \mathcal{F}) and state space (S, \mathcal{S}) , and that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration. Then \mathbf{X} is *progressively measurable* relative to \mathfrak{F} if $\mathbf{X} : \Omega \times T_t \rightarrow S$ is measurable with respect to $\mathcal{F}_t \otimes \mathcal{T}_t$ and \mathcal{S} for each $t \in T$.

Clearly if \mathbf{X} is progressively measurable with respect to a filtration, then it is progressively measurable with respect to any finer filtration. Of course when T is discrete, then any process \mathbf{X} is measurable, and any process \mathbf{X} adapted to \mathfrak{F} progressively measurable, so these definitions are only of interest in the case of continuous time.

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process with sample space (Ω, \mathcal{F}) and state space (S, \mathcal{S}) , and that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration. If \mathbf{X} is progressively measurable relative to \mathfrak{F} then

1. \mathbf{X} is measurable.
2. \mathbf{X} is adapted to \mathfrak{F} .

Proof

Suppose that \mathbf{X} is progressively measurable relative to \mathfrak{F} .

1. If $A \in \mathcal{S}$ then

$$\mathbf{X}^{-1}(A) = \{(\omega, t) \in \Omega \times T : X_t(\omega) \in A\} = \bigcup_{n=1}^{\infty} \{(\omega, t) \in \Omega \times T_n : X_t(\omega) \in A\} \quad (2.11.1)$$

By assumption, the n th term in the union is in $\mathcal{F} \otimes \mathcal{T}_n \subseteq \mathcal{F} \otimes \mathcal{T}$, so the union is in $\mathcal{F} \otimes \mathcal{T}$.

2. Suppose that $t \in T$. Then $\mathbf{X} : \Omega \times T_t \rightarrow S$ is measurable with respect to $\mathcal{F}_t \otimes \mathcal{T}_t$ and \mathcal{S} . But $X_t : \Omega \rightarrow S$ is just the cross section of this function at t and hence is measurable with respect to \mathcal{F}_t and \mathcal{S} .

When the state space is a topological space (which is usually the case), then as you might guess, there is a natural link between continuity of the sample paths and progressive measurability.

Suppose that S has an LCCB topology and that \mathcal{S} is the σ -algebra of Borel sets. Suppose also that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is right continuous. Then \mathbf{X} is progressively measurable relative to the natural filtration \mathfrak{F}^0 .

So if \mathbf{X} is right continuous, then \mathbf{X} is progressively measurable with respect to any filtration to which \mathbf{X} is adapted. Recall that in the previous section, we studied different ways that two stochastic processes can be equivalent. The following example illustrates some of the subtleties of processes in continuous time.

Suppose that $\Omega = T = [0, \infty)$, $\mathcal{F} = \mathcal{T}$ is the σ -algebra of Borel measurable subsets of $[0, \infty)$, and \mathbb{P} is any continuous probability measure on (Ω, \mathcal{F}) . Let $S = \{0, 1\}$ and $\mathcal{S} = \mathcal{P}(S) = \{\emptyset, \{0\}, \{1\}, \{0, 1\}\}$. For $t \in T$ and $\omega \in \Omega$, define $X_t(\omega) = \mathbf{1}_t(\omega)$ and $Y_t(\omega) = 0$. Then

1. $\mathbf{X} = \{X_t : t \in T\}$ is a version of $\mathbf{Y} = \{Y_t : t \in T\}$
2. \mathbf{X} is not adapted to the natural filtration of \mathbf{Y} .

Proof

1. This was shown in the previous section, but here it is again: For $t \in T$, $\mathbb{P}(X_t \neq Y_t) = \mathbb{P}(\{t\}) = 0$.
2. Trivially, $\sigma(Y_t) = \{\emptyset, \Omega\}$ for every $t \in T$, so $\sigma\{Y_s : 0 \leq s \leq t\} = \{\emptyset, \Omega\}$. But $\sigma(X_t) = \{\emptyset, \{t\}, \Omega \setminus \{t\}, \Omega\}$

Completion

Suppose now that P is a probability measure on (Ω, \mathcal{F}) . Recall that \mathcal{F} is complete with respect to P if $A \in \mathcal{F}$, $B \subseteq A$, and $P(A) = 0$ imply $B \in \mathcal{F}$ (and hence $P(B) = 0$). That is, if A is an event with probability 0 and $B \subseteq A$, then B is also an event (and also has probability 0). For a filtration, the following definition is appropriate.

The filtration $\{\mathcal{F}_t : t \in T\}$ is *complete* with respect to a probability measure P on (Ω, \mathcal{F}) if

1. \mathcal{F} is complete with respect to P
2. If $A \in \mathcal{F}$ and $P(A) = 0$ then $A \in \mathcal{F}_0$.

Suppose P is a probability measure on (Ω, \mathcal{F}) and that the filtration $\{\mathcal{F}_t : t \in T\}$ is complete with respect to P . If $A \in \mathcal{F}$ is a null event ($P(A) = 0$) or an almost certain event ($P(A) = 1$) then $A \in \mathcal{F}_t$ for every $t \in T$.

Proof

This follows since almost certain events are complements of null events and since the σ -algebras are increasing in $t \in T$.

Recall that if P is a probability measure on (Ω, \mathcal{F}) , but \mathcal{F} is not complete with respect to P , then \mathcal{F} can always be *completed*. Here's a review of how this is done: Let

$$\mathcal{N} = \{A \subseteq \Omega : \text{there exists } N \in \mathcal{F} \text{ with } P(N) = 0 \text{ and } A \subseteq N\} \quad (2.11.2)$$

So \mathcal{N} is the collection of null sets. Then we let $\mathcal{F}^P = \sigma(\mathcal{F} \cup \mathcal{N})$ and extend P to \mathcal{F}^P in the natural way: if $A \in \mathcal{F}^P$ and A differs from $B \in \mathcal{F}$ by a null set, then $P(A) = P(B)$. Filtrations can also be completed.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) and that P is a probability measure on (Ω, \mathcal{F}) . As above, let \mathcal{N} denote the collection of null subsets of Ω , and for $t \in T$, let $\mathcal{F}_t^P = \sigma(\mathcal{F}_t \cup \mathcal{N})$. Then $\mathfrak{F}^P = \{\mathcal{F}_t^P : t \in T\}$ is a filtration on (Ω, \mathcal{F}^P) that is finer than \mathfrak{F} and is complete relative to P .

Proof

If $s, t \in T$ with $s \leq t$ then $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ and hence

$$\sigma(\mathcal{F}_s \cup \mathcal{N}) \subseteq \sigma(\mathcal{F}_t \cup \mathcal{N}) \subseteq \sigma(\mathcal{F} \cup \mathcal{N}) \quad (2.11.3)$$

and so $\mathcal{F}_s^P \subseteq \mathcal{F}_t^P \subseteq \mathcal{F}^P$. The probability measure P can be extended to \mathcal{F}^P as described above, and hence is defined on \mathcal{F}_t^P for each $t \in T$. By construction, if $A \in \mathcal{F}^P$ and $P(A) = 0$ then $A \in \mathcal{N}$ so \mathfrak{F}^P is complete with respect to P .

Naturally, \mathfrak{F}^P is the *completion* of \mathfrak{F} with respect to P . Sometimes we need to consider *all* probability measures on (Ω, \mathcal{F}) .

Let \mathcal{P} denote the collection of probability measures on (Ω, \mathcal{F}) , and suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) . Let $\mathcal{F}^* = \bigcap \{\mathcal{F}^P : P \in \mathcal{P}\}$, and let $\mathfrak{F}^* = \bigwedge \{\mathfrak{F}^P : P \in \mathcal{P}\}$. Then \mathfrak{F}^* is a filtration on (Ω, \mathcal{F}^*) , known as the *universal completion* of \mathfrak{F} .

Proof

Note that \mathfrak{F}^P is a filtration on (Ω, \mathcal{F}^P) for each $P \in \mathcal{P}$, so \mathfrak{F}^* is a filtration on (Ω, \mathcal{F}^*) .

The last definition must seem awfully obscure, but it does have a place. In the theory of Markov processes, we usually allow arbitrary initial distributions, which in turn produces a large collection of distributions on the sample space.

Right Continuity

In continuous time, we sometimes need to refine a given filtration somewhat.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is a filtration on (Ω, \mathcal{F}) . For $t \in [0, \infty)$, define $\mathcal{F}_{t+} = \bigcap \{\mathcal{F}_s : s \in (t, \infty)\}$. Then $\mathfrak{F}_+ = \{\mathcal{F}_{t+} : t \in T\}$ is also a filtration on (Ω, \mathcal{F}) and is finer than \mathfrak{F} .

Proof

For $t \in [0, \infty)$ note that \mathcal{F}_{t+} is a σ -algebra since it is the intersection of σ -algebras, and clearly $\mathcal{F}_{t+} \subseteq \mathcal{F}$. Next, if $s, t \in [0, \infty)$ with $s \leq t$, then $\{\mathcal{F}_r : r \in (t, \infty)\} \subseteq \{\mathcal{F}_r : r \in (s, \infty)\}$, so it follows that

$$\mathcal{F}_{s+} = \bigcap \{\mathcal{F}_r : r \in (s, \infty)\} \subseteq \bigcap \{\mathcal{F}_r : r \in (t, \infty)\} = \mathcal{F}_{t+} \quad (2.11.4)$$

Finally, for $t \in [0, \infty)$, $\mathcal{F}_t \subseteq \mathcal{F}_s$ for every $s \in (t, \infty)$ so $\mathcal{F}_t \subseteq \bigcap \{\mathcal{F}_s : s \in (t, \infty)\} = \mathcal{F}_{t+}$.

Since the σ -algebras in a filtration are increasing, it follows that for $t \in [0, \infty)$, $\mathcal{F}_{t+} = \bigcap \{\mathcal{F}_s : s \in (t, t+\epsilon)\}$ for every $\epsilon \in (0, \infty)$. So if the filtration \mathfrak{F} encodes the information available as time goes by, then the filtration \mathfrak{F}_+ allows an “infinitesimal peak into the future” at each $t \in [0, \infty)$. In light of the previous result, the next definition is natural.

A filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is *right continuous* if $\mathfrak{F}_+ = \mathfrak{F}$, so that $\mathcal{F}_{t+} = \mathcal{F}_t$ for every $t \in [0, \infty)$.

Right continuous filtrations have some nice properties, as we will see later. If the original filtration is not right continuous, the slightly refined filtration is:

Suppose again that $\mathcal{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is a filtration. Then \mathcal{F}_+ is a right continuous filtration.

Proof

For $t \in T$

$$\mathcal{F}_{t++} = \bigcap \{\mathcal{F}_{s+} : s \in (t, \infty)\} = \bigcap \left\{ \bigcap \{\mathcal{F}_r : r \in (s, \infty)\} : s \in (t, \infty) \right\} = \bigcap \{\mathcal{F}_u : u \in (t, \infty)\} = \mathcal{F}_{t+} \quad (2.11.5)$$

For a stochastic process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ in continuous time, often the filtration \mathcal{F} that is most useful is the right-continuous refinement of the natural filtration. That is, $\mathcal{F} = \mathcal{F}_+$, so that $\mathcal{F}_t = \sigma\{X_s : s \in [0, t]\}_+$ for $t \in [0, \infty)$.

Stopping Times

Basic Properties

Suppose again that we have a fixed sample space (Ω, \mathcal{F}) . Random variables taking values in the time set T are important, but often as we will see, it's necessary to allow such variables to take the value ∞ as well as finite times. So let $T_\infty = T \cup \{\infty\}$. We extend order to T_∞ by the obvious rule that $t < \infty$ for every $t \in T$. We also extend the topology on T to T_∞ by the rule that for each $s \in T$, the set $\{t \in T_\infty : t > s\}$ is an open neighborhood of ∞ . That is, T_∞ is the *one-point compactification* of T . The reason for this is to preserve the meaning of *time converging to infinity*. That is, if (t_1, t_2, \dots) is a sequence in T_∞ then $t_n \rightarrow \infty$ as $n \rightarrow \infty$ if and only if, for every $t \in T$ there exists $m \in \mathbb{N}_+$ such that $t_n > t$ for $n > m$. We then give T_∞ the Borel σ -algebra \mathcal{T}_∞ as before. In discrete time, this is once again the discrete σ -algebra, so that all subsets are measurable. In both cases, we now have an enhanced time space is $(T_\infty, \mathcal{T}_\infty)$. A random variable τ taking values in T_∞ is called a *random time*.

Suppose that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) . A random time τ is a *stopping time* relative to \mathcal{F} if $\{\tau \leq t\} \in \mathcal{F}_t$ for each $t \in T$.

In a sense, a stopping time is a random time that does not require that we see into the future. That is, we can tell whether or not $\tau \leq t$ from our information at time t . The term *stopping time* comes from gambling. Consider a gambler betting on games of chance. The gambler's decision to stop gambling at some point in time and accept his fortune must define a stopping time. That is, the gambler can base his decision to stop gambling on all of the information that he has at that point in time, but not on what will happen in the future. The terms *Markov time* and *optional time* are sometimes used instead of stopping time. If τ is a stopping time relative to a filtration, then it is also a stopping time relative to any finer filtration:

Suppose that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ and $\mathcal{G} = \{\mathcal{G}_t : t \in T\}$ are filtrations on (Ω, \mathcal{F}) , and that \mathcal{G} is finer than \mathcal{F} . If a random time τ is a stopping time relative to \mathcal{F} then τ is a stopping time relative to \mathcal{G} .

Proof

This is very simple. If $t \in T$ then $\{\tau \leq t\} \in \mathcal{F}_t$ and hence $\{\tau \leq t\} \in \mathcal{G}_t$ since $\mathcal{F}_t \subseteq \mathcal{G}_t$.

So, the finer the filtration, the larger the collection of stopping times. In fact, *every* random time is a stopping time relative to the finest filtration \mathcal{F} where $\mathcal{F}_t = \mathcal{F}$ for every $t \in T$. But this filtration corresponds to having complete information from the beginning of time, which of course is usually not sensible. At the other extreme, for the coarsest filtration \mathcal{F} where $\mathcal{F}_t = \{\Omega, \emptyset\}$ for every $t \in T$, the only stopping times are constants. That is, random times of the form $\tau(\omega) = t$ for every $\omega \in \Omega$, for some $t \in T_\infty$.

Suppose again that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) . A random time τ is a stopping time relative to \mathcal{F} if and only if $\{\tau > t\} \in \mathcal{F}_t$ for each $t \in T$.

Proof

This result is trivial since $\{\tau > t\} = \{\tau \leq t\}^c$ for $t \in T$.

Suppose again that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) , and that τ is a stopping time relative to \mathcal{F} . Then

1. $\{\tau < t\} \in \mathcal{F}_t$ for every $t \in T$.
2. $\{\tau \geq t\} \in \mathcal{F}_t$ for every $t \in T$.
3. $\{\tau = t\} \in \mathcal{F}_t$ for every $t \in T$.

Proof

1. Suppose first that $T = \mathbb{N}$. Then $\{\tau < t\} = \{\tau \leq t-1\} \in \mathcal{F}_{t-1} \subseteq \mathcal{F}_t$ for $t \in \mathbb{N}$. Next suppose that $T = [0, \infty)$. Fix $t \in (0, \infty)$ and let (s_1, s_2, \dots) be a strictly increasing sequence in $[0, \infty)$ with $s_n \uparrow t$ as $n \rightarrow \infty$. Then $\{\tau < t\} = \bigcup_{n=1}^{\infty} \{\tau \leq s_n\}$. But $\{\tau \leq s_n\} \in \mathcal{F}_{s_n} \subseteq \mathcal{F}_t$ for each n , so $\{\tau < t\} \in \mathcal{F}_t$.
2. This follows from (a) since $\{\tau \geq t\} = \{\tau < t\}^c$ for $t \in T$.
3. For $t \in T$ note that $\{\tau = t\} = \{\tau \leq t\} \setminus \{\tau < t\}$. Both events in the set difference are in \mathcal{F}_t .

Note that when $T = \mathbb{N}$, we actually showed that $\{\tau < t\} \in \mathcal{F}_{t-1}$ and $\{\tau \geq t\} \in \mathcal{F}_{t-1}$. The converse to part (a) (or equivalently (b)) is not true, but in continuous time there is a connection to the right-continuous refinement of the filtration.

Suppose that $T = [0, \infty)$ and that $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is a filtration on (Ω, \mathcal{F}) . A random time τ is a stopping time relative to \mathfrak{F}_+ if and only if $\{\tau < t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$.

Proof

So restated, we need to show that $\{\tau \leq t\} \in \mathcal{F}_{t+}$ for every $t \in [0, \infty)$ if and only if $\{\tau < t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$. (Note by the way, that this is not the same as the statement that for every $t \in T$, $\{\tau < t\} \in \mathcal{F}_{t+}$ if and only if $\{\tau \leq t\} \in \mathcal{F}_t$, which is not true.) Suppose first that τ is a stopping time relative to \mathfrak{F} . Fix $t \in [0, \infty)$ and let (t_1, t_2, \dots) be a strictly decreasing sequence in $[0, \infty)$ with $t_n \downarrow t$ as $n \rightarrow \infty$. Then for each $k \in \mathbb{N}_+$, $\{\tau \leq t\} = \bigcap_{n=k}^{\infty} \{\tau < t_n\}$. If $s > t$ then there exists $k \in \mathbb{N}_+$ such that $t_n < s$ for each $n \geq k$. Hence $\{\tau < t_n\} \in \mathcal{F}_{t_n} \subseteq \mathcal{F}_s$ for $n \geq k$, and so it follows that $\{\tau \leq t\} \in \mathcal{F}_s$. Since this is true for every $s > t$ it follows $\{\tau \leq t\} \in \mathcal{F}_{t+}$. Conversely, suppose that $\{\tau \leq t\} \in \mathcal{F}_{t+}$ for every $t \in [0, \infty)$. Fix $t \in (0, \infty)$ and let (t_1, t_2, \dots) be a strictly increasing sequence in $(0, \infty)$ with $t_n \uparrow t$ as $n \rightarrow \infty$. Then $\bigcup_{i=1}^{\infty} \{\tau \leq t_n\} = \{\tau < t\}$. But for every $n \in \mathbb{N}_+$

$$\{\tau \leq t_n\} \in \mathcal{F}_{t_n+} = \bigcap \{\mathcal{F}_s : s \in (t_n, t)\} \subseteq \mathcal{F}_t \quad (2.11.6)$$

Hence $\{\tau < t\} \in \mathcal{F}_t$.

If $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is a filtration and τ is a random time that satisfies $\{\tau < t\} \in \mathcal{F}_t$ for every $t \in T$, then some authors call τ a *weak stopping time* or say that τ is *weakly optional* for the filtration \mathfrak{F} . But to me, the increase in jargon is not worthwhile, and it's better to simply say that τ is a stopping time for the filtration \mathfrak{F}_+ . The following corollary now follows.

Suppose that $T = [0, \infty)$ and that $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is a right-continuous filtration. A random time τ is a stopping time relative to \mathfrak{F} if and only if $\{\tau < t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$.

The converse to part (c) of the result [above](#) holds in discrete time.

Suppose that $T = \mathbb{N}$ and that $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ is a filtration on (Ω, \mathcal{F}) . A random time τ is a stopping time for \mathfrak{F} if and only if $\{\tau = n\} \in \mathcal{F}_n$ for every $n \in \mathbb{N}$.

Proof

If τ is a stopping time then as shown [above](#), $\{\tau = n\} \in \mathcal{F}_n$ for every $n \in \mathbb{N}$. Conversely, suppose that this condition holds. For $n \in \mathbb{N}$, $\{\tau \leq n\} = \bigcup_{k=0}^n \{\tau = k\}$. But $\{\tau = k\} \in \mathcal{F}_k \subseteq \mathcal{F}_n$ for $k \in \{0, 1, \dots, n\}$ so $\{\tau \leq n\} \in \mathcal{F}_n$.

Basic Constructions

As noted above, a constant element of T_∞ is a stopping time, but not a very interesting one.

Suppose $s \in T_\infty$ and that $\tau(\omega) = s$ for all $\omega \in \Omega$. The τ is a stopping time relative to any filtration on (Ω, \mathcal{F}) .

Proof

For $t \in T$ note that $\{\tau \leq t\} = \Omega$ if $s \leq t$ and $\{\tau \leq t\} = \emptyset$ if $s > t$.

If the filtration $\{\mathcal{F}_t : t \in T\}$ is complete, then a random time that is almost certainly a constant is also a stopping time. The following theorems give some basic ways of constructing new stopping times from ones we already have.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) and that τ_1 and τ_2 are stopping times relative to \mathfrak{F} . Then each of the following is also a stopping time relative to \mathfrak{F} :

1. $\tau_1 \vee \tau_2 = \max\{\tau_1, \tau_2\}$
2. $\tau_1 \wedge \tau_2 = \min\{\tau_1, \tau_2\}$
3. $\tau_1 + \tau_2$

Proof

1. Note that $\{\tau_1 \vee \tau_2 \leq t\} = \{\tau_1 \leq t\} \cap \{\tau_2 \leq t\} \in \mathcal{F}_t$ for $t \in T$, so the result follows from the definition.
2. Note that $\{\tau_1 \wedge \tau_2 > t\} = \{\tau_1 > t\} \cap \{\tau_2 > t\} \in \mathcal{F}_t$ for $t \in T$, so the result follows from the [result above](#).
3. This is simple when $T = \mathbb{N}$. In this case, $\{\tau_1 + \tau_2 \leq t\} = \bigcup_{n=0}^t \{\tau_1 = n\} \cap \{\tau_2 \leq t - n\}$. But for $n \leq t$, $\{\tau_1 = n\} \in \mathcal{F}_n \subseteq \mathcal{F}_t$ and $\{\tau_2 \leq t - n\} \in \mathcal{F}_{t-n} \subseteq \mathcal{F}_t$. Hence $\{\tau_1 + \tau_2 \leq t\} \in \mathcal{F}_t$. Suppose instead that $T = [0, \infty)$ and $t \in T$. Then $\tau_1 + \tau_2 > t$ if and only if either $\tau_1 \leq t$ and $\tau_2 > t - \tau_1$ or $\tau_1 > t$. Of course $\{\tau_1 > t\} \in \mathcal{F}_t$ so we just need to show that the first event is also in \mathcal{F}_t . Note that $\tau_1 \leq t$ and $\tau_2 > t - \tau_1$ if and only if there exists a rational $q \in [0, t]$ such that $q \leq \tau_1 \leq t$ and $\tau_2 \geq t - q$. Each of these events is in \mathcal{F}_t and hence so is the union of the events over the countable collection of rational $q \in [0, t]$.

It follows that if $(\tau_1, \tau_2, \dots, \tau_n)$ is a finite sequence of stopping times relative to \mathfrak{F} , then each of the following is also a stopping time relative to \mathfrak{F} :

- $\tau_1 \vee \tau_2 \vee \dots \vee \tau_n$
- $\tau_1 \wedge \tau_2 \wedge \dots \wedge \tau_n$
- $\tau_1 + \tau_2 + \dots + \tau_n$

We have to be careful when we try to extend these results to infinite sequences.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) , and that $(\tau_n : n \in \mathbb{N}_+)$ is a sequence of stopping times relative to \mathfrak{F} . Then $\sup\{\tau_n : n \in \mathbb{N}_+\}$ is also a stopping time relative to \mathfrak{F} .

Proof

Let $\tau = \sup\{\tau_n : n \in \mathbb{N}_+\}$. Note that τ exists in T_∞ and is a random time. For $t \in T$, $\{\tau \leq t\} = \bigcap_{n=1}^\infty \{\tau_n \leq t\}$. But each event in the intersection is in \mathcal{F}_t and hence so is the intersection.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) , and that $(\tau_n : n \in \mathbb{N}_+)$ is an increasing sequence of stopping times relative to \mathfrak{F} . Then $\lim_{n \rightarrow \infty} \tau_n$ is a stopping time relative to \mathfrak{F} .

Proof

This is a corollary of the previous theorem. Since the sequence is increasing, $\lim_{n \rightarrow \infty} \tau_n = \sup\{\tau_n : n \in \mathbb{N}_+\}$.

Suppose that $T = [0, \infty)$ and that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) . If $(\tau_n : n \in \mathbb{N}_+)$ is a sequence of stopping times relative to \mathfrak{F} , then each of the following is a stopping time relative to \mathfrak{F}_+ :

1. $\inf\{\tau_n : n \in \mathbb{N}_+\}$
2. $\liminf_{n \rightarrow \infty} \tau_n$
3. $\limsup_{n \rightarrow \infty} \tau_n$

Proof

1. Let $\tau = \inf\{\tau_n : n \in \mathbb{N}_+\}$. Then $\{\tau \geq t\} = \bigcap_{n=1}^\infty \{\tau_n \geq t\} \in \mathcal{F}_t$ for $t \in T$. Hence τ is a stopping time relative to \mathfrak{F}_+ by the result above.
2. Recall that $\liminf_{n \rightarrow \infty} \tau_n = \sup\{\inf\{\tau_k : k \geq n\} : n \in \mathbb{N}_+\}$ and so this is a stopping time relative to \mathfrak{F}_+ by part (a) and the [result above](#) on supremums.
3. Similarly note that $\limsup_{n \rightarrow \infty} \tau_n = \inf\{\sup\{\tau_k : k \geq n\} : n \in \mathbb{N}_+\}$ and so this is a stopping time relative to \mathfrak{F}_+ by part (a) and the [result above](#) on supremums.

As a simple corollary, we have the following results:

Suppose that $T = [0, \infty)$ and that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a right-continuous filtration on (Ω, \mathcal{F}) . If $(\tau_n : n \in \mathbb{N}_+)$ is a sequence of stopping times relative to \mathfrak{F} , then each of the following is also a stopping time relative to \mathfrak{F} :

1. $\inf\{\tau_n : n \in \mathbb{N}_+\}$
2. $\liminf_{n \rightarrow \infty} \tau_n$

3. $\limsup_{n \rightarrow \infty} \tau_n$

The σ -Algebra of a Stopping Time

Consider again the general setting of a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ on the sample space (Ω, \mathcal{F}) , and suppose that τ is a stopping time relative to \mathfrak{F} . We want to define the σ -algebra \mathcal{F}_τ of events up to the random time τ , analogous to \mathcal{F}_t the σ -algebra of events up to a fixed time $t \in T$. Here is the appropriate definition:

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) and that τ is a stopping time relative to \mathfrak{F} . Define $\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \in T\}$. Then \mathcal{F}_τ is a σ -algebra.

Proof

First $\Omega \in \mathcal{F}_\tau$ since $\Omega \cap \{\tau \leq t\} = \{\tau \leq t\} \in \mathcal{F}_t$ for $t \in T$. If $A \in \mathcal{F}_\tau$ then $A^c \cap \{\tau \leq t\} = \{\tau \leq t\} \setminus (A \cap \{\tau \leq t\}) \in \mathcal{F}_t$ for $t \in T$. Finally, suppose that $A_i \in \mathcal{F}_\tau$ for i in a countable index set I . Then $(\bigcup_{i \in I} A_i) \cap \{\tau \leq t\} = \bigcup_{i \in I} (A_i \cap \{\tau \leq t\}) \in \mathcal{F}_t$ for $t \in T$.

Thus, an event A is in \mathcal{F}_τ if we can determine if A and $\tau \leq t$ both occurred given our information at time t . If τ is constant, then \mathcal{F}_τ reduces to the corresponding member of the original filtration, which clearly should be the case, and is additional motivation for the definition.

Suppose again that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) . Fix $s \in T$ and define $\tau(\omega) = s$ for all $\omega \in \Omega$. Then $\mathcal{F}_\tau = \mathcal{F}_s$.

Proof

Suppose that $A \in \mathcal{F}_s$. Then $A \in \mathcal{F}$ and for $t \in T$, $A \cap \{\tau \leq t\} = A$ if $s \leq t$ and $A \cap \{\tau \leq t\} = \emptyset$ if $s > t$. In either case, $A \cap \{\tau \leq t\} \in \mathcal{F}_t$ and hence $A \in \mathcal{F}_\tau$. Conversely, suppose that $A \in \mathcal{F}_\tau$. Then $A = A \cap \{\tau \leq s\} \in \mathcal{F}_s$.

Clearly, if we have the information available in \mathcal{F}_τ , then we should know the value of τ itself. This is also true:

Suppose again that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) and that τ is a stopping time relative to \mathfrak{F} . Then τ is measurable with respect to \mathcal{F}_τ .

Proof

It suffices to show that $\{\tau \leq s\} \in \mathcal{F}_\tau$ for each $s \in T$. For $s, t \in T$,

$$\{\tau \leq t\} \cap \{\tau \leq s\} = \{\tau \leq s \wedge t\} \in \mathcal{F}_{s \wedge t} \subseteq \mathcal{F}_t \quad (2.11.7)$$

Hence $\{\tau \leq s\} \in \mathcal{F}_\tau$.

Here are other results that relate the σ -algebra of a stopping time to the original filtration.

Suppose again that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) and that τ is a stopping time relative to \mathfrak{F} . If $A \in \mathcal{F}_\tau$ then for $t \in T$,

1. $A \cap \{\tau < t\} \in \mathcal{F}_t$
2. $A \cap \{\tau = t\} \in \mathcal{F}_t$

Proof

1. By definition, $A \cap \{\tau \leq t\} \in \mathcal{F}_t$. But $\{\tau < t\} \subseteq \{\tau \leq t\}$ and $\{\tau < t\} \in \mathcal{F}_t$. Hence

$$A \cap \{\tau < t\} = A \cap \{\tau \leq t\} \cap \{\tau < t\} \in \mathcal{F}_t.$$

2. similarly $\{\tau = t\} \subseteq \{\tau \leq t\}$ and $\{\tau = t\} \in \mathcal{F}_t$. Hence $A \cap \{\tau = t\} = A \cap \{\tau \leq t\} \cap \{\tau = t\} \in \mathcal{F}_t$

The σ -algebra of a stopping time relative to a filtration is related to the σ -algebra of the stopping time relative to a finer filtration in the natural way.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ and $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ are filtrations on (Ω, \mathcal{F}) and that \mathfrak{G} is finer than \mathfrak{F} . If τ is a stopping time relative to \mathfrak{F} then $\mathcal{F}_\tau \subseteq \mathcal{G}_\tau$.

Proof

From the result above, τ is also a stopping time relative to \mathfrak{G} , so the statement makes sense. If $A \in \mathcal{F}_\tau$ then for $t \in T$, $A \cap \{\tau \leq t\} \in \mathcal{F}_t \subseteq \mathcal{G}_t$, so $A \in \mathcal{G}_\tau$.

When two stopping times are ordered, their σ -algebras are also ordered.

Suppose that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) and that ρ and τ are stopping times for \mathcal{F} with $\rho \leq \tau$. Then $\mathcal{F}_\rho \subseteq \mathcal{F}_\tau$.

Proof

Suppose that $A \in \mathcal{F}_\rho$ and $t \in T$. Note that $\{\tau \leq t\} \subseteq \{\rho \leq t\}$. By definition, $A \cap \{\rho \leq t\} \in \mathcal{F}_t$ and $\{\tau \leq t\} \in \mathcal{F}_t$. Hence $A \cap \{\tau \leq t\} = A \cap \{\rho \leq t\} \cap \{\tau \leq t\} \in \mathcal{F}_t$, so $A \in \mathcal{F}_\tau$.

Suppose again that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) , and that ρ and τ are stopping times for \mathcal{F} . Then each of the following events is in \mathcal{F}_τ and in \mathcal{F}_ρ .

1. $\{\rho < \tau\}$
2. $\{\rho = \tau\}$
3. $\{\rho > \tau\}$
4. $\{\rho \leq \tau\}$
5. $\{\rho \geq \tau\}$

Proof

The proofs are easy when $T = \mathbb{N}$.

1. Let $t \in T$. Then

$$\{\rho < \tau\} \cap \{\tau \leq t\} = \bigcup_{n=0}^t \bigcup_{k=0}^{n-1} \{\tau = n, \rho = k\} \quad (2.11.8)$$

But each event in the union is in \mathcal{F}_t .

2. Similarly, let $t \in T$. Then

$$\{\rho = \tau\} \cap \{\tau \leq t\} = \bigcup_{n=0}^t \{\rho = n, \tau = n\} \quad (2.11.9)$$

and again each event in the union is in \mathcal{F}_t .

3. This follows from symmetry, reversing the roles of ρ and τ in part (a).
4. Note that $\{\rho \leq \tau\} = \{\rho < \tau\} \cup \{\rho = \tau\} \in \mathcal{F}_\tau$.
5. Similarly, note that $\{\rho \geq \tau\} = \{\rho > \tau\} \cup \{\rho = \tau\} \in \mathcal{F}_\tau$.

We can “stop” a filtration at a stopping time. In the next subsection, we will stop a stochastic process in the same way.

Suppose again that $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on (Ω, \mathcal{F}) , and that τ is a stopping times for \mathcal{F} . For $t \in T$ define $\mathcal{F}_t^\tau = \mathcal{F}_{t \wedge \tau}$. Then $\mathcal{F}^\tau = \{\mathcal{F}_t^\tau : t \in T\}$ is a filtration and is coarser than \mathcal{F} .

Proof

The random time $t \wedge \tau$ is a stopping time for each $t \in T$ by the result above, so \mathcal{F}_t^τ is a sub σ -algebra of \mathcal{F} . If $t \in T$, then by definition, $A \in \mathcal{F}_t^\tau$ if and only if $A \cap \{t \wedge \tau \leq r\} \in \mathcal{F}_r$ for every $r \in T$. But for $r \in T$, $\{t \wedge \tau \leq r\} = \Omega$ if $r \geq t$ and $\{t \wedge \tau \leq r\} = \{\tau \leq r\}$ if $r < t$. Hence $A \in \mathcal{F}_t^\tau$ if and only if $A \cap \{\tau \leq r\} \in \mathcal{F}_r$ for $r < t$ and $A \in \mathcal{F}_t$. So in particular, \mathcal{F}^τ is coarser than \mathcal{F} . Further, suppose $s, t \in T$ with $s \leq t$, and that $A \in \mathcal{F}_s^\tau$. Let $r \in T$. If $r < s$ then $A \cap \{\tau \leq r\} \in \mathcal{F}_r$. If $s \leq r < t$ then $A \in \mathcal{F}_s \subseteq \mathcal{F}_r$ and $\{\tau \leq r\} \in \mathcal{F}_r$ so again $A \cap \{\tau \leq r\} \in \mathcal{F}_r$. Finally if $r \geq t$ then $A \in \mathcal{F}_s \subseteq \mathcal{F}_t$. Hence $A \in \mathcal{F}_t^\tau$.

Stochastic Processes

As usual, the most common setting is when we have a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ defined on our sample space (Ω, \mathcal{F}) and with state space (S, \mathcal{S}) . If τ is a random time, we are often interested in the state X_τ at the random time. But there are two issues. First, τ may take the value infinity, in which case X_τ is not defined. The usual solution is to introduce a new “death state” δ , and define $X_\infty = \delta$. The σ -algebra \mathcal{S} on S is extended to $S_\delta = S \cup \{\delta\}$ in the natural way, namely $\mathcal{S}_\delta = \sigma(S \cup \{\delta\})$.

Our other problem is that we naturally expect X_τ to be a random variable (that is, measurable), just as X_t is a random variable for a deterministic $t \in T$. Moreover, if \mathbf{X} is adapted to a filtration $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$, then we would naturally also expect X_τ to be

measurable with respect to \mathcal{F}_τ , just as X_t is measurable with respect to \mathcal{F}_t for deterministic $t \in T$. But this is not obvious, and in fact is not true without additional assumptions. Note that X_τ is a random state at a random time, and so depends on an outcome $\omega \in \Omega$ in two ways: $X_{\tau(\omega)}(\omega)$.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on the sample space (Ω, \mathcal{F}) with state space (S, \mathcal{S}) , and that \mathbf{X} is measurable. If τ is a finite random time, then X_τ is measurable. That is, X_τ is a random variable with values in S .

Proof

Note that $X_\tau : \Omega \rightarrow S$ is the composition of the function $\omega \mapsto (\omega, \tau(\omega))$ from Ω to $\Omega \times T$ with the function $(\omega, t) \mapsto X_t(\omega)$ from $\Omega \times T$ to S . The first function is measurable because the two coordinate functions are measurable. The second function is measurable by assumption.

This result is one of the main reasons for the definition of a measurable process in the first place. Sometimes we literally want to *stop* the random process at a random time τ . As you might guess, this is the origin of the term *stopping time*.

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on the sample space (Ω, \mathcal{F}) with state space (S, \mathcal{S}) , and that \mathbf{X} is measurable. If τ is a random time, then the process $\mathbf{X}^\tau = \{X_t^\tau : t \in T\}$ defined by $X_t^\tau = X_{t \wedge \tau}$ for $t \in T$ is the process \mathbf{X} *stopped at τ* .

Proof

For each $t \in T$, note that $t \wedge \tau$ is a finite random time, and hence $X_{t \wedge \tau}$ is measurable by the previous result. Thus \mathbf{X}^τ is a well-defined stochastic process on (Ω, \mathcal{F}) with state space (S, \mathcal{S}) .

When the original process is progressively measurable, so is the stopped process.

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on the sample space (Ω, \mathcal{F}) with state space (S, \mathcal{S}) , and that \mathbf{X} is progressively measurable with respect to a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$. If τ is a stopping time relative to \mathfrak{F} , then the stopped process $\mathbf{X}^\tau = \{X_t^\tau : t \in T\}$ is progressively measurable with respect to the stopped filtration \mathfrak{F}^τ .

Since \mathfrak{F} is finer than \mathfrak{F}^τ , it follows that \mathbf{X}^τ is also progressively measurable with respect to \mathfrak{F} .

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on the sample space (Ω, \mathcal{F}) with state space (S, \mathcal{S}) , and that \mathbf{X} is progressively measurable with respect to a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ on (Ω, \mathcal{F}) . If τ is a finite stopping time relative to \mathfrak{F} then X_τ is measurable with respect to \mathcal{F}_τ .

For many random processes, the first time that the process enters or hits a set of states is particularly important. In the discussion that follows, let $T_+ = \{t \in T : t > 0\}$, the set of positive times.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on (Ω, \mathcal{F}) with state space (S, \mathcal{S}) . For $A \in \mathcal{S}$, define

1. $\rho_A = \inf\{t \in T : X_t \in A\}$, the first *entry time* to A .
2. $\tau_A = \inf\{t \in T_+ : X_t \in A\}$, the first *hitting time* to A .

As usual, $\inf(\emptyset) = \infty$ so $\rho_A = \infty$ if $X_t \notin A$ for all $t \in T$, so that the process never enters A , and $\tau_A = \infty$ if $X_t \notin A$ for all $t \in T_+$, so that the process never hits A . In discrete time, it's easy to see that these are stopping times.

Suppose that $\{X_n : n \in \mathbb{N}\}$ is a stochastic process on (Ω, \mathcal{F}) with state space (S, \mathcal{S}) . If $A \in \mathcal{S}$ then τ_A and ρ_A are stopping times relative to the natural filtration \mathfrak{F}^0 .

Proof

Let $n \in \mathbb{N}$. Note that $\{\rho_A > n\} = \{X_0 \notin A, X_1 \notin A, \dots, X_n \notin A\} \in \sigma\{X_0, X_1, \dots, X_n\}$. Similarly, $\{\tau_A > n\} = \{X_1 \notin A, X_2 \notin A, \dots, X_n \notin A\} \subseteq \sigma\{X_0, X_1, \dots, X_n\}$.

So of course in discrete time, τ_A and ρ_A are stopping times relative to any filtration \mathfrak{F} to which \mathbf{X} is adapted. You might think that τ_A and ρ_A should always be stopping times, since $\tau_A \leq t$ if and only if $X_s \in A$ for some $s \in T_+$ with $s \leq t$, and $\rho_A \leq t$ if and only if $X_s \in A$ for some $s \in T$ with $s \leq t$. It would seem that these events are known if one is allowed to observe the process up to time t .

The problem is that when $T = [0, \infty)$, these are uncountable unions, so we need to make additional assumptions on the stochastic process \mathbf{X} or the filtration \mathfrak{F} , or both.

Suppose that S has an LCCB topology, and that \mathcal{S} is the σ -algebra of Borel sets. Suppose also that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is right continuous and has left limits. Then τ_A and ρ_A are stopping times relative to \mathfrak{F}_+^0 for every open $A \in \mathcal{S}$.

Here is another result that requires less of the stochastic process \mathbf{X} , but more of the filtration \mathfrak{F} .

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a stochastic process on (Ω, \mathcal{F}) that is progressively measurable relative to a complete, right-continuous filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$. If $A \in \mathcal{S}$ then ρ_A and τ_A are stopping times relative to \mathfrak{F} .

This page titled [2.11: Filtrations and Stopping Times](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

3: Distributions

Recall that a *probability distribution* is just another name for a [probability measure](#). Most distributions are associated with [random variables](#), and in fact every distribution *can* be associated with a random variable. In this chapter we explore the basic types of probability distributions (discrete, continuous, mixed), and the ways that distributions can be defined using density functions, distribution functions, and quantile functions. We also study the relationship between the distribution of a random vector and the distributions of its components, conditional distributions, and how the distribution of a random variable changes when the variable is transformed.

In the advanced sections, we study convergence in distribution, one of the most important types of convergence. We also construct the abstract integral with respect to a positive measure and study the basic properties of the integral. This leads in turn to general (signed measures), absolute continuity and singularity, and the existence of density functions. Finally, we study various vector spaces of functions that are defined by integral pro

[3.1: Discrete Distributions](#)

[3.2: Continuous Distributions](#)

[3.3: Mixed Distributions](#)

[3.4: Joint Distributions](#)

[3.5: Conditional Distributions](#)

[3.6: Distribution and Quantile Functions](#)

[3.7: Transformations of Random Variables](#)

[3.8: Convergence in Distribution](#)

[3.9: General Distribution Functions](#)

[3.10: The Integral With Respect to a Measure](#)

[3.11: Properties of the Integral](#)

[3.12: General Measures](#)

[3.13: Absolute Continuity and Density Functions](#)

[3.14: Function Spaces](#)

This page titled [3: Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.1: Discrete Distributions

Basic Theory

Definitions and Basic Properties

As usual, our starting point is a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$. So to review, S is the set of outcomes, \mathcal{S} the collection of events, and \mathbb{P} the probability measure on the sample space (S, \mathcal{S}) . We use the terms *probability measure* and *probability distribution* synonymously in this text. Also, since we use a general definition of random variable, every probability measure can be thought of as the probability distribution of a random variable, so we can always take this point of view if we like. Indeed, most probability measures naturally have random variables associated with them.

Recall that the sample space (S, \mathcal{S}) is *discrete* if S is countable and $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S . In this case, \mathbb{P} is a *discrete distribution* and $(S, \mathcal{S}, \mathbb{P})$ is a *discrete probability space*.

For the remainder of our discussion we assume that $(S, \mathcal{S}, \mathbb{P})$ is a discrete probability space. In the picture below, the blue dots are intended to represent points of positive probability.




Figure 3.1.1: A discrete distribution

It's very simple to describe a discrete probability distribution with the function that assigns probabilities to the individual points in S .

The function f on S defined by $f(x) = \mathbb{P}(\{x\})$ for $x \in S$ is the *probability density function* of \mathbb{P} , and satisfies the following properties:

1. $f(x) \geq 0$, $x \in S$
2. $\sum_{x \in S} f(x) = 1$
3. $\sum_{x \in A} f(x) = \mathbb{P}(A)$ for $A \subseteq S$

Proof

These properties follow from the axioms of a probability measure.

1. $f(x) = \mathbb{P}(\{x\}) \geq 0$ since probabilities are nonnegative.
2. $\sum_{x \in S} f(x) = \sum_{x \in S} \mathbb{P}(\{x\}) = \mathbb{P}(S) = 1$ by the countable additivity axiom.
3. $\sum_{x \in A} f(x) = \sum_{x \in A} \mathbb{P}(\{x\}) = \mathbb{P}(A)$ for $A \subseteq S$ again, by the countable additivity axiom.

Property (c) is particularly important since it shows that a discrete probability distribution is completely determined by its probability density function. Conversely, any function that satisfies properties (a) and (b) can be used to construct a discrete probability distribution on S via property (c).

A nonnegative function f on S that satisfies $\sum_{x \in S} f(x) = 1$ is a (discrete) *probability density function* on S , and then \mathbb{P} defined as follows is a probability measure on S .

$$\mathbb{P}(A) = \sum_{x \in A} f(x), \quad A \subseteq S \quad (3.1.1)$$

Proof

1. $\mathbb{P}(A) = \sum_{x \in A} f(x) \geq 0$ since f is nonnegative.
2. $\mathbb{P}(S) = \sum_{x \in S} f(x) = 1$ by property (b)
3. Suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of subsets of S , and let $A = \bigcup_{i \in I} A_i$. Then

$$\mathbb{P}(A) = \sum_{x \in A} f(x) = \sum_{i \in I} \sum_{x \in A_i} f(x) = \sum_{i \in I} \mathbb{P}(A_i) \quad (3.1.2)$$

Note that since f is nonnegative, the order of the terms in the sum do not matter.

Technically, f is the density of \mathbb{P} relative to counting measure $\#$ on S . The technicalities are discussed in detail in the advanced section on absolute continuity and density functions.




Figure 3.1.2: A discrete distribution is completely determined by its probability density function.

The set of outcomes S is often a countable subset of some larger set, such as \mathbb{R}^n for some $n \in \mathbb{N}_+$. But not always. We might want to consider a random variable with values in a deck of cards, or a set of words, or some other discrete population of objects. Of course, we can always map a countable set S one-to-one into a Euclidean set, but it might be contrived or unnatural to do so. In any event, if S is a subset of a larger set, we can always extend a probability density function f , if we want, to the larger set by defining $f(x) = 0$ for $x \notin S$. Sometimes this extension simplifies formulas and notation. Put another way, the “set of values” is often a convenience set that includes the points with positive probability, but perhaps other points as well.

Suppose that f is a probability density function on S . Then $\{x \in S : f(x) > 0\}$ is the *support set* of the distribution.

Values of x that maximize the probability density function are important enough to deserve a name.

Suppose again that f is a probability density function on S . An element $x \in S$ that maximizes f is a *mode* of the distribution.

When there is only one mode, it is sometimes used as a measure of the *center* of the distribution.

A discrete probability distribution defined by a probability density function f is equivalent to a *discrete mass distribution*, with total mass 1. In this analogy, S is the (countable) set of point masses, and $f(x)$ is the mass of the point at $x \in S$. Property (c) in (2) above simply means that the mass of a set A can be found by adding the masses of the points in A .

But let's consider a probabilistic interpretation, rather than one from physics. We start with a basic random variable X for an experiment, defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Suppose that X has a discrete distribution on S with probability density function f . So in this setting, $f(x) = \mathbb{P}(X = x)$ for $x \in S$. We create a new, compound experiment by conducting independent repetitions of the original experiment. So in the compound experiment, we have a sequence of independent random variables (X_1, X_2, \dots) each with the same distribution as X ; in statistical terms, we are *sampling* from the distribution of X . Define

$$f_n(x) = \frac{1}{n} \# \{i \in \{1, 2, \dots, n\} : X_i = x\} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i = x), \quad x \in S \quad (3.1.3)$$

Note that $f_n(x)$ is the *relative frequency* of outcome $x \in S$ in the first n runs. Note also that $f_n(x)$ is a random variable for the compound experiment for each $x \in S$. By the law of large numbers, $f_n(x)$ should converge to $f(x)$, in some sense, as $n \rightarrow \infty$. The function f_n is called the *empirical probability density function*, and it is in fact a (random) probability density function, since it satisfies properties (a) and (b) of (2). Empirical probability density functions are displayed in most of the simulation apps that deal with discrete variables.

It's easy to construct discrete probability density functions from other nonnegative functions defined on a countable set.

Suppose that g is a nonnegative function defined on S , and let

$$c = \sum_{x \in S} g(x) \quad (3.1.4)$$

If $0 < c < \infty$, then the function f defined by $f(x) = \frac{1}{c} g(x)$ for $x \in S$ is a discrete probability density function on S .

Proof

Clearly $f(x) \geq 0$ for $x \in S$. also

$$\sum_{x \in S} f(x) = \frac{1}{c} \sum_{x \in S} g(x) = \frac{c}{c} = 1 \quad (3.1.5)$$

Note that since we are assuming that g is nonnegative, $c = 0$ if and only if $g(x) = 0$ for every $x \in S$. At the other extreme, $c = \infty$ could only occur if S is infinite (and the infinite series diverges). When $0 < c < \infty$ (so that we can construct the probability density function f), c is sometimes called the *normalizing constant*. This result is useful for constructing probability density functions with desired functional properties (domain, shape, symmetry, and so on).

Conditional Densities

Suppose again that X is a random variable on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and that X takes values in our discrete set S . The distribution of X (and hence the probability density function of X) is based on the underlying probability measure on the sample space (Ω, \mathcal{F}) . This measure could be a conditional probability measure, conditioned on a given event $E \in \mathcal{F}$ (with $\mathbb{P}(E) > 0$). The probability density function in this case is

$$f(x | E) = \mathbb{P}(X = x | E), \quad x \in S \quad (3.1.6)$$

Except for notation, no new concepts are involved. Therefore, all results that hold for discrete probability density functions in general have analogies for conditional discrete probability density functions.

For fixed $E \in \mathcal{F}$ with $\mathbb{P}(E) > 0$ the function $x \mapsto f(x | E)$ is a discrete probability density function on S . That is,

1. $f(x | E) \geq 0$ for $x \in S$.
2. $\sum_{x \in S} f(x | E) = 1$
3. $\sum_{x \in A} f(x | E) = \mathbb{P}(X \in A | E)$ for $A \subseteq S$

Proof

This is a consequence of the fact that $A \mapsto \mathbb{P}(A | E)$ is a probability measure on (Ω, \mathcal{F}) . The function $x \mapsto f(x | E)$ plays the same role for the conditional probability measure that f does for the original probability measure \mathbb{P} .

In particular, the event E could be an event defined in terms of the random variable X itself.

Suppose that $B \subseteq S$ and $\mathbb{P}(X \in B) > 0$. The conditional probability density function of X given $X \in B$ is the function on B defined by

$$f(x | X \in B) = \frac{f(x)}{\mathbb{P}(X \in B)} = \frac{f(x)}{\sum_{y \in B} f(y)}, \quad x \in B \quad (3.1.7)$$

Proof

This follows from the previous theorem. $f(x | X \in B) = \mathbb{P}(X = x, X \in B) / \mathbb{P}(X \in B)$. The numerator is $f(x)$ if $x \in B$ and is 0 if $x \notin B$.

Note that the denominator is simply the [normalizing constant](#) for f restricted to B . Of course, $f(x | B) = 0$ for $x \in B^c$.

Conditioning and Bayes' Theorem

Suppose again that X is a random variable defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and that X has a discrete distribution on S , with probability density function f . We assume that $f(x) > 0$ for $x \in S$ so that the distribution has support S . The versions of the law of total probability and Bayes' theorem given in the following theorems follow immediately from the corresponding results in the section on Conditional Probability. Only the notation is different.

Law of Total Probability. If $E \in \mathcal{F}$ is an event then

$$\mathbb{P}(E) = \sum_{x \in S} f(x) \mathbb{P}(E | X = x) \quad (3.1.8)$$

Proof

Note that $\{\{X = x\} : x \in S\}$ is a countable partition of the sample space Ω . That is, these events are disjoint and their union is the entire sample space Ω . Hence

$$\mathbb{P}(E) = \sum_{x \in S} \mathbb{P}(E \cap \{X = x\}) = \sum_{x \in S} \mathbb{P}(X = x) \mathbb{P}(E | X = x) = \sum_{x \in S} f(x) \mathbb{P}(E | X = x) \quad (3.1.9)$$

This result is useful, naturally, when the distribution of X and the conditional probability of E given the values of X are known. When we compute $\mathbb{P}(E)$ in this way, we say that we are *conditioning* on X . Note that $\mathbb{P}(E)$, as expressed by the formula, is a weighted average of $\mathbb{P}(E | X = x)$, with weight factors $f(x)$, over $x \in S$.

Bayes' Theorem. If $E \in \mathcal{F}$ is an event with $\mathbb{P}(E) > 0$ then

$$f(x | E) = \frac{f(x) \mathbb{P}(E | X = x)}{\sum_{y \in S} f(y) \mathbb{P}(E | X = y)}, \quad x \in S \quad (3.1.10)$$

Proof

Note that the numerator of the fraction on the right is $\mathbb{P}(X = x) \mathbb{P}(E | X = x) = \mathbb{P}(\{X = x\} \cap E)$. The denominator is $\mathbb{P}(E)$ by the previous theorem. Hence the ratio is $\mathbb{P}(X = x | E) = f(x | E)$.

Bayes' theorem, named for Thomas Bayes, is a formula for the conditional probability density function of X given E . Again, it is useful when the quantities on the right are known. In the context of Bayes' theorem, the (unconditional) distribution of X is referred to as the *prior* distribution and the conditional distribution as the *posterior* distribution. Note that the denominator in Bayes' formula is $\mathbb{P}(E)$ and is simply the [normalizing constant](#) for the function $x \mapsto f(x) \mathbb{P}(E | X = x)$.

Examples and Special Cases

We start with some simple (albeit somewhat artificial) discrete distributions. After that, we study three special parametric models—the [discrete uniform distribution](#), [hypergeometric distributions](#), and [Bernoulli trials](#). These models are very important, so when working the computational problems that follow, try to see if the problem fits one of these models. As always, be sure to try the problems yourself before looking at the answers and proofs in the text.

Simple Discrete Distributions

Let g be the function defined by $g(n) = n(10 - n)$ for $n \in \{1, 2, \dots, 9\}$.

1. Find the probability density function f that is [proportional](#) to g as in .
2. Sketch the graph of f and find the mode of the distribution.
3. Find $\mathbb{P}(3 \leq N \leq 6)$ where N has probability density function f .

Answer

1. $f(n) = \frac{1}{165} n(10 - n)$ for $n \in \{1, 2, \dots, 9\}$
2. mode $n = 5$
3. $\frac{94}{165}$

Let g be the function defined by $g(n) = n^2(10 - n)$ for $n \in \{1, 2, \dots, 10\}$.

1. Find the probability density function f that is [proportional](#) to g .
2. Sketch the graph of f and find the mode of the distribution.
3. Find $\mathbb{P}(3 \leq N \leq 6)$ where N has probability density function f .

Answer

1. $f(n) = \frac{1}{825} n^2(10 - n)$ for $n \in \{1, 2, \dots, 9\}$
2. mode $n = 7$
3. $\frac{428}{825}$

Let g be the function defined by $g(x, y) = x + y$ for $(x, y) \in \{1, 2, 3\}^2$.

1. Sketch the domain of g .

2. Find the probability density function f that is **proportional** to g .
3. Find the mode of the distribution.
4. Find $\mathbb{P}(X > Y)$ where (X, Y) has probability density function f .

Answer

1. $f(x, y) = \frac{1}{36}(x + y)$ for $(x, y) \in \{1, 2, 3\}^2$
2. mode $(3, 3)$
3. $\frac{2}{9}$

Let g be the function defined by $g(x, y) = xy$ for $(x, y) \in \{(1, 1), (1, 2), (1, 3), (2, 2), (2, 3), (3, 3)\}$

1. Sketch the domain of g .
2. Find the probability density function f that is **proportional** to g .
3. Find the mode of the distribution.
4. Find $\mathbb{P}[(X, Y) \in \{(1, 2), (1, 3), (2, 2), (2, 3)\}]$ where (X, Y) has probability density function f .

Answer

1. $f(x, y) = \frac{1}{25}xy$ for $(x, y) \in \{(1, 1), (1, 2), (1, 3), (2, 2), (2, 3), (3, 3)\}$
2. mode $(3, 3)$
3. $\frac{3}{5}$

Consider the following game: An urn initially contains one red and one green ball. A ball is selected at random, and if the ball is green, the game is over. If the ball is red, the ball is returned to the urn, another red ball is added, and the game continues. At each stage, a ball is selected at random, and if the ball is green, the game is over. If the ball is red, the ball is returned to the urn, another red ball is added, and the game continues. Let X denote the length of the game (that is, the number of selections required to obtain a green ball). Find the probability density function of X .

Solution

Note that X takes values in \mathbb{N}_+ . Using the multiplication rule for conditional probabilities, the PDF f of X is given by

$$f(1) = \frac{1}{2} = \frac{1}{1 \cdot 2}, \quad f(2) = \frac{1}{2} \cdot \frac{1}{3} = \frac{1}{2 \cdot 3}, \quad f(3) = \frac{1}{2} \cdot \frac{2}{3} \cdot \frac{1}{4} = \frac{1}{3 \cdot 4} \quad (3.1.11)$$

and in general, $f(x) = \frac{1}{x(x+1)}$ for $x \in \mathbb{N}_+$. By partial fractions, $f(x) = \frac{1}{x} - \frac{1}{x+1}$ for $x \in \mathbb{N}_+$ so we can check that f is a valid PDF:

$$\sum_{x=1}^{\infty} \left(\frac{1}{x} - \frac{1}{x+1} \right) = \lim_{n \rightarrow \infty} \sum_{x=1}^n \left(\frac{1}{x} - \frac{1}{x+1} \right) = \lim_{n \rightarrow \infty} \left(1 - \frac{1}{n+1} \right) = 1 \quad (3.1.12)$$

Discrete Uniform Distributions

An element X is chosen at random from a finite set S . The distribution of X is the *discrete uniform distribution* on S .

1. X has probability density function f given by $f(x) = 1/\#(S)$ for $x \in S$.
2. $\mathbb{P}(X \in A) = \#(A)/\#(S)$ for $A \subseteq S$.

Proof

The phrase *at random* means that all outcomes are equally likely.

Many random variables that arise in sampling or combinatorial experiments are transformations of uniformly distributed variables. The next few exercises review the standard methods of sampling from a finite population. The parameters m and n are positive integers.

Suppose that n elements are chosen at random, with replacement from a set D with m elements. Let \mathbf{X} denote the ordered sequence of elements chosen. Then \mathbf{X} is uniformly distributed on the Cartesian power set $S = D^n$, and has probability density function f given by

$$f(\mathbf{x}) = \frac{1}{m^n}, \quad \mathbf{x} \in S \quad (3.1.13)$$

Proof

Recall that $\#(D^n) = m^n$.

Suppose that n elements are chosen at random, without replacement from a set D with m elements (so $n \leq m$). Let \mathbf{X} denote the ordered sequence of elements chosen. Then \mathbf{X} is uniformly distributed on the set S of permutations of size n chosen from D , and has probability density function f given by

$$f(\mathbf{x}) = \frac{1}{m^{(n)}}, \quad \mathbf{x} \in S \quad (3.1.14)$$

Proof

Recall that the number of permutations of size n from D is $m^{(n)}$.

Suppose that n elements are chosen at random, without replacement, from a set D with m elements (so $n \leq m$). Let \mathbf{W} denote the unordered set of elements chosen. Then \mathbf{W} is uniformly distributed on the set T of combinations of size n chosen from D , and has probability density function f given by

$$f(\mathbf{w}) = \frac{1}{\binom{m}{n}}, \quad \mathbf{w} \in T \quad (3.1.15)$$

Proof

Recall that the number of combinations of size n from D is $\binom{m}{n}$.

Suppose that X is uniformly distributed on a finite set S and that B is a nonempty subset of S . Then the conditional distribution of X given $X \in B$ is uniform on B .

Proof

From (7), the conditional probability density function of X given $X \in B$ is

$$f(x | B) = \frac{f(x)}{\mathbb{P}(X \in B)} = \frac{1/\#(S)}{\#(B)/\#(S)} = \frac{1}{\#(B)}, \quad x \in B \quad (3.1.16)$$

Hypergeometric Models

Suppose that a *dichotomous* population consists of m objects of two different types: r of the objects are type 1 and $m - r$ are type 0. Here are some typical examples:

- The objects are persons, each either *male* or *female*.
- The objects are voters, each either a *democrat* or a *republican*.
- The objects are devices of some sort, each either *good* or *defective*.
- The objects are fish in a lake, each either *tagged* or *untagged*.
- The objects are balls in an urn, each either *red* or *green*.

A sample of n objects is chosen at random (without replacement) from the population. Recall that this means that the samples, either ordered or unordered are equally likely. Note that this probability model has three parameters: the population size m , the number of type 1 objects r , and the sample size n . Each is a nonnegative integer with $r \leq m$ and $n \leq m$. Now, suppose that we keep track of order, and let X_i denote the type of the i th object chosen, for $i \in \{1, 2, \dots, n\}$. Thus, X_i is an indicator variable (that is, a variable that just takes values 0 and 1).

$\mathbf{X} = (X_1, X_2, \dots, X_n)$ has probability density function f given by

$$f(x_1, x_2, \dots, x_n) = \frac{r^{(y)}(m-r)^{(n-y)}}{m^{(n)}}, \quad (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \text{ where } y = x_1 + x_2 + \dots + x_n \quad (3.1.17)$$

Proof

Recall again that the ordered samples are equally likely, and there are $m^{(n)}$ such samples. The number of ways to select the y type 1 objects and place them in the positions where $x_i = 1$ is $r^{(y)}$. The number of ways to select the $n - y$ type 0 objects and place them in the positions where $x_i = 0$ is $(m - r)^{(n-y)}$. Thus the result follows from the multiplication principle.

Note that the value of $f(x_1, x_2, \dots, x_n)$ depends only on $y = x_1 + x_2 + \dots + x_n$, and hence is unchanged if (x_1, x_2, \dots, x_n) is permuted. This means that (X_1, X_2, \dots, X_n) is exchangeable. In particular, the distribution of X_i is the same as the distribution of X_1 , so $\mathbb{P}(X_i = 1) = \frac{r}{m}$. Thus, the variables are *identically distributed*. Also the distribution of (X_i, X_j) is the same as the distribution of (X_1, X_2) , so $\mathbb{P}(X_i = 1, X_j = 1) = \frac{r(r-1)}{m(m-1)}$. Thus, X_i and X_j are not independent, and in fact are negatively correlated.

Now let Y denote the number of type 1 objects in the sample. Note that $Y = \sum_{i=1}^n X_i$. Any counting variable can be written as a sum of indicator variables.

Y has probability density function g given by.

$$g(y) = \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}}, \quad y \in \{0, 1, \dots, n\} \quad (3.1.18)$$

1. $g(y-1) < g(y)$ if and only if $y < t$ where $t = (r+1)(n+1)/(m+2)$.
2. If t is not a positive integer, there is a single mode at $\lfloor t \rfloor$.
3. If t is a positive integer, then there are two modes, at $t-1$ and t .

Proof

Recall again that the unordered samples of size n chosen from the population are equally likely. By the multiplication principle, the number of samples with exactly y type 1 objects and $n - y$ type 0 objects is $\binom{m}{y} \binom{m-r}{n-y}$. The total number of samples is $\binom{m}{n}$.

1. Note that $g(y-1) < g(y)$ if and only if $\binom{r}{y-1} \binom{m-r}{n+1-y} < \binom{r}{y} \binom{m-r}{n-y}$. Writing the binomial coefficients in terms of factorials and canceling terms gives $g(y-1) < g(y)$ if and only if $y < t$, where t is given above.
2. By the same argument, $f(y-1) = f(y)$ if and only if $y = t$. If t is not an integer then this cannot happen. Letting $z = \lfloor t \rfloor$, it follows from (a) that $g(y) < g(z)$ if $y < z$ or $y > z$.
3. If t is a positive integer, then by (b), $g(t-1) = g(t)$ and by (a) $g(y) < g(t-1)$ if $y < t-1$ and $g(y) < g(t)$ if $y > t$.

The distribution defined by the probability density function in the last result is the *hypergeometric distributions* with parameters m , r , and n . The term *hypergeometric* comes from a certain class of special functions, but is not particularly helpful in terms of remembering the model. Nonetheless, we are stuck with it. The set of values $\{0, 1, \dots, n\}$ is a *convenience set*: it contains all of the values that have positive probability, but depending on the parameters, some of the values may have probability 0. Recall our convention for binomial coefficients: for $j, k \in \mathbb{N}_+$, $\binom{k}{j} = 0$ if $j > k$. Note also that the hypergeometric distribution is unimodal: the probability density function increases and then decreases, with either a single mode or two adjacent modes.

We can extend the hypergeometric model to a population of three types. Thus, suppose that our population consists of m objects; r of the objects are type 1, s are type 2, and $m - r - s$ are type 0. Here are some examples:

- The objects are voters, each a *democrat*, a *republican*, or an *independent*.
- The objects are cicadas, each one of three species: *tredecula*, *tredecassini*, or *tredecim*.
- The objects are peaches, each classified as *small*, *medium*, or *large*.
- The objects are faculty members at a university, each an *assistant professor*, or an *associate professor*, or a *full professor*.

Once again, a sample of n objects is chosen at random (without replacement). The probability model now has four parameters: the population size m , the type sizes r and s , and the sample size n . All are nonnegative integers with $r + s \leq m$ and $n \leq m$. Moreover, we now need two random variables to keep track of the counts for the three types in the sample. Let Y denote the number of type 1 objects in the sample and Z the number of type 2 objects in the sample.

(Y, Z) has probability density function h given by

$$h(y, z) = \frac{\binom{r}{y} \binom{s}{z} \binom{m-r-s}{n-y-z}}{\binom{m}{n}}, \quad (y, z) \in \{0, 1, \dots, n\}^2 \text{ with } y + z \leq n \quad (3.1.19)$$

Proof

Once again, by the multiplication principle, the number of samples of size n from the population with exactly y type 1 objects, z type 2 objects, and $n - y - z$ type 0 objects is $\binom{r}{y} \binom{s}{z} \binom{m-r-s}{n-y-z}$. The total number of samples of size n is $\binom{m}{n}$.

The distribution defined by the density function in the last exercise is the *bivariate hypergeometric distribution* with parameters m , r , s , and n . Once again, the domain given is a convenience set; it includes the set of points with positive probability, but depending on the parameters, may include points with probability 0. Clearly, the same general pattern applies to populations with even more types. However, because of all of the parameters, the formulas are not worth remembering in detail; rather, just note the pattern, and remember the combinatorial meaning of the binomial coefficient. The hypergeometric model will be revisited later in this chapter, in the section on joint distributions and in the section on conditional distributions. The hypergeometric distribution and the multivariate hypergeometric distribution are studied in detail in the chapter on Finite Sampling Models. This chapter contains a variety of distributions that are based on discrete uniform distributions.

Bernoulli Trials

A *Bernoulli trials sequence* is a sequence (X_1, X_2, \dots) of independent, identically distributed indicator variables. Random variable X_i is the outcome of trial i , where in the usual terminology of reliability, 1 denotes *success* while 0 denotes *failure*. The process is named for Jacob Bernoulli. Let $p = \mathbb{P}(X_i = 1) \in [0, 1]$ denote the success parameter of the process. Note that the indicator variables in the [hypergeometric model](#) satisfy one of the assumptions of Bernoulli trials (identical distributions) but not the other (independence).

$\mathbf{X} = (X_1, X_2, \dots, X_n)$ has probability density function f given by

$$f(x_1, x_2, \dots, x_n) = p^y (1-p)^{n-y}, \quad (x_1, x_2, \dots, x_n) \in \{0, 1\}^n, \text{ where } y = x_1 + x_2 + \dots + x_n \quad (3.1.20)$$

Proof

By definition, $\mathbb{P}(X_i = 1) = p$ and $\mathbb{P}(X_i = 0) = 1 - p$. Equivalently, $\mathbb{P}(X_i = x) = p^x (1-p)^{1-x}$ for $x \in \{0, 1\}$. The formula for f then follows by independence.

Now let Y denote the number of successes in the first n trials. Note that $Y = \sum_{i=1}^n X_i$, so we see again that a complicated random variable can be written as a sum of simpler ones. In particular, a counting variable can always be written as a sum of indicator variables.

Y has probability density function g given by

$$g(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (3.1.21)$$

1. $g(y-1) < g(y)$ if and only if $y < t$, where $t = (n+1)p$.
2. If t is not a positive integer, there is a single mode at $\lfloor t \rfloor$.
3. If t is a positive integer, then there are two modes, at $t-1$ and t .

Proof

From the previous result, any particular sequence of n Bernoulli trials with y successes and $n-y$ failures has probability $p^y (1-p)^{n-y}$. The number of such sequences is $\binom{n}{y}$, so the formula for g follows by the additivity of probability.

1. Note that $g(y-1) < g(y)$ if and only if $\binom{n}{y-1} p^{y-1} (1-p)^{n-y+1} < \binom{n}{y} p^y (1-p)^{n-y}$. Writing the binomial coefficients in terms of factorials and canceling gives $g(y-1) < g(y)$ if and only if $y < t$ where $t = (n+1)p$.
2. By the same argument, $g(y-1) = g(y)$ if and only if $y = t$. If t is not an integer, this cannot happen. Letting $z = \lfloor t \rfloor$, it follows from (a) that $g(y) < g(z)$ if $y < z$ or $y > z$.
3. If t is a positive integer, then by (b), $g(t-1) = g(t)$ and by (a) $g(y) < g(t-1)$ if $y < t-1$ and $g(y) < g(t)$ if $y > t$.

The distribution defined by the probability density function in the last theorem is called the *binomial distribution* with parameters n and p . The distribution is unimodal: the probability density function at first increases and then decreases, with either a single mode or two adjacent modes. The binomial distribution is studied in detail in the chapter on Bernoulli Trials.

Suppose that $p > 0$ and let N denote the trial number of the first success. Then N has probability density function h given by

$$h(n) = (1-p)^{n-1}p, \quad n \in \mathbb{N}_+ \quad (3.1.22)$$

The probability density function h is decreasing and the mode is $n = 1$.

Proof

For $n \in \mathbb{N}_+$, the event $\{N = n\}$ means that the first $n - 1$ trials were failures and trial n was a success. Each trial results in failure with probability $1 - p$ and success with probability p , and the trials are independent, so $\mathbb{P}(N = n) = (1 - p)^{n-1}p$. Using geometric series, we can check that

$$\sum_{n=1}^{\infty} h(n) = \sum_{n=1}^{\infty} p(1-p)^{n-1} = \frac{p}{1-(1-p)} = 1 \quad (3.1.23)$$

The distribution defined by the probability density function in the last exercise is the *geometric distribution* on \mathbb{N}_+ with parameter p . The geometric distribution is studied in detail in the chapter on Bernoulli Trials.

Sampling Problems

In the following exercises, be sure to check if the problem fits one of the general models above.

An urn contains 30 red and 20 green balls. A sample of 5 balls is selected at random, without replacement. Let Y denote the number of red balls in the sample.

1. Compute the probability density function of Y explicitly and identify the distribution by name and parameter values.
2. Graph the probability density function and identify the mode(s).
3. Find $\mathbb{P}(Y > 3)$.

Answer

1. $f(0) = 0.0073$, $f(1) = 0.0686$, $f(2) = 0.2341$, $f(3) = 0.3641$, $f(4) = 0.2587$, $f(5) = 0.0673$ Hypergeometric with $m = 50$, $r = 30$, $n = 5$
2. mode: $y = 3$
3. $\mathbb{P}(Y > 3) = 0.3260$

In the ball and urn experiment, select sampling without replacement and set $m = 50$, $r = 30$, and $n = 5$. Run the experiment 1000 times and note the agreement between the empirical density function of Y and the probability density function.

An urn contains 30 red and 20 green balls. A sample of 5 balls is selected at random, *with* replacement. Let Y denote the number of red balls in the sample.

1. Compute the probability density function of Y explicitly and identify the distribution by name and parameter values.
2. Graph the probability density function and identify the mode(s).
3. Find $\mathbb{P}(Y > 3)$.

Answer

1. $f(0) = 0.0102$, $f(1) = 0.0768$, $f(2) = 0.2304$, $f(3) = 0.3456$, $f(4) = 0.2592$, $f(5) = 0.0778$ Binomial with $n = 5$, $p = 3/5$
2. mode: $y = 3$
3. $\mathbb{P}(Y > 3) = 0.3370$

In the ball and urn experiment, select sampling with replacement and set $m = 50$, $r = 30$, and $n = 5$. Run the experiment 1000 times and note the agreement between the empirical density function of Y and the probability density function.

A group of voters consists of 50 democrats, 40 republicans, and 30 independents. A sample of 10 voters is chosen at random, without replacement. Let X denote the number of democrats in the sample and Y the number of republicans in the sample.

1. Give the probability density function of X .
2. Give the probability density function of Y .
3. Give the probability density function of (X, Y) .
4. Find the probability that the sample has at least 4 democrats and at least 4 republicans.

Answer

1. $g(x) = \frac{\binom{50}{x} \binom{70}{10-x}}{\binom{120}{10}}$ for $x \in \{0, 1, \dots, 10\}$ This is the hypergeometric distribution with parameters $m = 120$, $r = 50$ and $n = 10$.
2. $h(y) = \frac{\binom{40}{y} \binom{80}{10-y}}{\binom{120}{10}}$ for $y \in \{0, 1, \dots, 10\}$ This is the hypergeometric distribution with parameters $m = 120$, $r = 40$ and $n = 10$.
3. $f(x, y) = \frac{\binom{50}{x} \binom{40}{y} \binom{30}{10-x-y}}{\binom{120}{10}}$ for $(x, y) \in \{0, 1, \dots, 10\}^2$ with $x + y \leq 10$. This is the bivariate hypergeometric distribution with parameters $m = 120$, $r = 50$, $s = 40$ and $n = 10$.
4. $\mathbb{P}(X \geq 4, Y \geq 4) = \frac{15\,137\,200}{75\,597\,113} \approx 0.200$

The Math Club at Enormous State University (ESU) has 20 freshmen, 40 sophomores, 30 juniors, and 10 seniors. A committee of 8 club members is chosen at random, without replacement to organize π -day activities. Let X denote the number of freshman in the sample, Y the number of sophomores, and Z the number of juniors.

1. Give the probability density function of X .
2. Give the probability density function of Y .
3. Give the probability density function of Z .
4. Give the probability density function of (X, Y) .
5. Give the probability density function of (X, Y, Z) .
6. Find the probability that the committee has no seniors.

Answer

1. $f_X(x) = \frac{\binom{20}{x} \binom{80}{8-x}}{\binom{100}{8}}$ for $x \in \{0, 1, \dots, 8\}$ This is the hypergeometric distribution with parameters $m = 100$, $r = 20$, and $n = 8$.
2. $f_Y(y) = \frac{\binom{40}{y} \binom{60}{8-y}}{\binom{100}{8}}$ for $y \in \{0, 1, \dots, 8\}$ This is the hypergeometric distribution with parameters $m = 100$, $r = 40$, and $n = 8$.
3. $f_Z(z) = \frac{\binom{30}{z} \binom{70}{8-z}}{\binom{100}{8}}$ for $z \in \{0, 1, \dots, 8\}$ This is the hypergeometric distribution with parameters $m = 100$, $r = 30$, and $n = 8$.
4. $f_{X,Y}(x, y) = \frac{\binom{20}{x} \binom{40}{y} \binom{40}{8-x-y}}{\binom{100}{8}}$ for $(x, y) \in \{0, 1, \dots, 8\}^2$ with $x + y \leq 8$. This is the bivariate hypergeometric distribution with parameters $m = 100$, $r = 20$, $s = 40$ and $n = 8$.
5. $f_{X,Y,Z}(x, y, z) = \frac{\binom{20}{x} \binom{40}{y} \binom{30}{z} \binom{10}{8-x-y-z}}{\binom{100}{8}}$ for $(x, y, z) \in \{0, 1, \dots, 8\}^3$ with $x + y + z \leq 8$. This is the tri-variate hypergeometric distribution with parameters $m = 100$, $r = 20$, $s = 40$, $t = 30$, and $n = 8$.
6. $\mathbb{P}(X + Y + Z = 8) = \frac{156\,597\,013}{275\,935\,140} \approx 0.417$

Coins and Dice

Suppose that a coin with probability of heads p is tossed repeatedly, and the sequence of heads and tails is recorded.

1. Identify the underlying probability model by name and parameter.
2. Let Y denote the number of heads in the first n tosses. Give the probability density function of Y and identify the distribution by name and parameters.
3. Let N denote the number of tosses needed to get the first head. Give the probability density function of N and identify the distribution by name and parameter.

Answer

1. Bernoulli trials with success parameter p .
2. $f(k) = \binom{n}{k} p^k (1-p)^{n-k}$ for $k \in \{0, 1, \dots, n\}$. This is the binomial distribution with trial parameter n and success parameter p .
3. $g(n) = p(1-p)^{n-1}$ for $n \in \mathbb{N}_+$. This is the geometric distribution with success parameter p .

Suppose that a coin with probability of heads $p = 0.4$ is tossed 5 times. Let Y denote the number of heads.

1. Compute the probability density function of Y explicitly.
2. Graph the probability density function and identify the mode.
3. Find $\mathbb{P}(Y > 3)$.

Answer

1. $f(0) = 0.0778$ $f(1) = 0.2592$ $f(2) = 0.3456$ $f(3) = 0.2304$ $f(4) = 0.0768$ $f(5) = 0.0102$
2. mode: $k = 2$
3. $\mathbb{P}(Y > 3) = 0.0870$

In the binomial coin experiment, set $n = 5$ and $p = 0.4$. Run the experiment 1000 times and compare the empirical density function of Y with the probability density function.

Suppose that a coin with probability of heads $p = 0.2$ is tossed until heads occurs. Let N denote the number of tosses.

1. Find the probability density function of N .
2. Find $\mathbb{P}(N \leq 5)$.

Answer

1. $f(n) = (0.8)^{n-1} 0.2$ for $n \in \mathbb{N}_+$
2. $\mathbb{P}(N \leq 5) = 0.67232$

In the negative binomial experiment, set $k = 1$ and $p = 0.2$. Run the experiment 1000 times and compare the empirical density function with the probability density function.

Suppose that two fair, standard dice are tossed and the sequence of scores (X_1, X_2) recorded. Let $Y = X_1 + X_2$ denote the sum of the scores, $U = \min\{X_1, X_2\}$ the minimum score, and $V = \max\{X_1, X_2\}$ the maximum score.

1. Find the probability density function of (X_1, X_2) . Identify the distribution by name.
2. Find the probability density function of Y .
3. Find the probability density function of U .
4. Find the probability density function of V .
5. Find the probability density function of (U, V) .

Answer

We denote the PDFs by f , g , h_1 , h_2 , and h respectively.

1. $f(x_1, x_2) = \frac{1}{36}$ for $(x_1, x_2) \in \{1, 2, 3, 4, 5, 6\}^2$. This is the uniform distribution on $\{1, 2, 3, 4, 5, 6\}^2$.
2. $g(2) = g(12) = \frac{1}{36}$, $g(3) = g(11) = \frac{2}{36}$, $g(4) = g(10) = \frac{3}{36}$, $g(5) = g(9) = \frac{4}{36}$, $g(6) = g(8) = \frac{5}{36}$, $g(7) = \frac{6}{36}$
3. $h_1(1) = \frac{11}{36}$, $h_1(2) = \frac{9}{36}$, $h_1(3) = \frac{7}{36}$, $h_1(4) = \frac{5}{36}$, $h_1(5) = \frac{3}{36}$, $h_1(6) = \frac{1}{36}$
4. $h_2(1) = \frac{1}{36}$, $h_2(2) = \frac{3}{36}$, $h_2(3) = \frac{5}{36}$, $h_2(4) = \frac{7}{36}$, $h_2(5) = \frac{9}{36}$, $h_2(6) = \frac{11}{36}$
5. $h(u, v) = \frac{2}{36}$ if $u < v$, $h(u, v) = \frac{1}{36}$ if $u = v$ where $(u, v) \in \{1, 2, 3, 4, 5, 6\}^2$ with $u \leq v$

Note that (U, V) in the last exercise could serve as the outcome of the experiment that consists of throwing two standard dice if we did not bother to record order. Note from the previous exercise that this random vector does not have a uniform distribution when the dice are fair. The mistaken idea that this vector *should* have the uniform distribution was the cause of difficulties in the early development of probability.

In the dice experiment, select $n = 2$ fair dice. Select the following random variables and note the shape and location of the probability density function. Run the experiment 1000 times. For each of the following variables, compare the empirical density function with the probability density function.

1. Y , the sum of the scores.
2. U , the minimum score.
3. V , the maximum score.

In the die-coin experiment, a fair, standard die is rolled and then a fair coin is tossed the number of times showing on the die. Let N denote the die score and Y the number of heads.

1. Find the probability density function of N . Identify the distribution by name.
2. Find the probability density function of Y .

Answer

1. $g(n) = \frac{1}{6}$ for $n \in \{1, 2, 3, 4, 5, 6\}$ This is the uniform distribution on $\{1, 2, 3, 4, 5, 6\}$
2. $h(0) = \frac{63}{384}$, $h(1) = \frac{120}{384}$, $h(2) = \frac{90}{384}$, $h(3) = \frac{64}{384}$, $h(4) = \frac{29}{384}$, $h(5) = \frac{8}{384}$, $h(6) = \frac{1}{384}$

Run the die-coin experiment 1000 times. For the number of heads, compare the empirical density function with the probability density function.

Suppose that a bag contains 12 coins: 5 are fair, 4 are biased with probability of heads $\frac{1}{3}$; and 3 are two-headed. A coin is chosen at random from the bag and tossed 5 times. Let V denote the probability of heads of the selected coin and let Y denote the number of heads.

1. Find the probability density function of V .
2. Find the probability density function of Y .

Answer

1. $g(1/2) = 5/12$, $g(1/3) = 4/12$, $g(1) = 3/12$
2. $h(0) = 5311/93312$, $h(1) = 16315/93312$, $h(2) = 22390/93312$, $h(3) = 17270/93312$, $h(4) = 7355/93312$, $h(5) = 24671/93312$

Compare the [die-coin experiment](#) with the [bag of coins experiment](#). In the first experiment, we toss a coin with a *fixed* probability of heads a *random* number of times. In second experiment, we effectively toss a coin with a *random* probability of heads a *fixed* number of times. In both cases, we can think of starting with a binomial distribution and *randomizing* one of the parameters.

In the coin-die experiment, a fair coin is tossed. If the coin lands tails, a fair die is rolled. If the coin lands heads, an ace-six flat die is tossed (faces 1 and 6 have probability $\frac{1}{4}$ each, while faces 2, 3, 4, 5 have probability $\frac{1}{8}$ each). Find the probability density function of the die score Y .

Answer

$$f(y) = 5/24 \text{ for } y \in \{1, 6\}, f(y) = 7/24 \text{ for } y \in \{2, 3, 4, 5\}$$

Run the coin-die experiment 1000 times, with the settings in the previous exercise. Compare the empirical density function with the probability density function.

Suppose that a standard die is thrown 10 times. Let Y denote the number of times an ace or a six occurred. Give the probability density function of Y and identify the distribution by name and parameter values in each of the following cases:

1. The die is fair.

2. The die is an ace-six flat.

Answer

1. $f(k) = \binom{10}{k} \left(\frac{1}{3}\right)^k \left(\frac{2}{3}\right)^{10-k}$ for $k \in \{0, 1, \dots, 10\}$. This is the binomial distribution with trial parameter $n = 10$ and success parameter $p = \frac{1}{3}$
2. $f(k) = \binom{10}{k} \left(\frac{1}{2}\right)^{10}$ for $k \in \{0, 1, \dots, 10\}$. This is the binomial distribution with trial parameter $n = 10$ and success parameter $p = \frac{1}{2}$

Suppose that a standard die is thrown until an ace or a six occurs. Let N denote the number of throws. Give the probability density function of N and identify the distribution by name and parameter values in each of the following cases:

1. The die is fair.
2. The die is an ace-six flat.

Answer

1. $g(n) = \left(\frac{2}{3}\right)^{n-1} \frac{1}{3}$ for $n \in \mathbb{N}_+$. This is the geometric distribution with success parameter $p = \frac{1}{3}$
2. $g(n) = \left(\frac{1}{2}\right)^n$ for $n \in \mathbb{N}_+$. This is the geometric distribution with success parameter $p = \frac{1}{2}$

Fred and Wilma takes turns tossing a coin with probability of heads $p \in (0, 1)$: Fred first, then Wilma, then Fred again, and so forth. The first person to toss heads wins the game. Let N denote the number of tosses, and W the event that Wilma wins.

1. Give the probability density function of N and identify the distribution by name.
2. Compute $\mathbb{P}(W)$ and sketch the graph of this probability as a function of p .
3. Find the conditional probability density function of N given W .

Answer

1. $f(n) = p(1-p)^{n-1}$ for $n \in \mathbb{N}_+$. This is the geometric distribution with success parameter p .
2. $\mathbb{P}(W) = \frac{1-p}{2-p}$
3. $f(n | W) = p(2-p)(1-p)^{n-2}$ for $n \in \{2, 4, \dots\}$

The *alternating coin tossing game* is studied in more detail in the section on The Geometric Distribution in the chapter on Bernoulli trials.

Suppose that k players each have a coin with probability of heads p , where $k \in \{2, 3, \dots\}$ and where $p \in (0, 1)$.

1. Suppose that the players toss their coins at the same time. Find the probability that there is an *odd man*, that is, one player with a different outcome than all the rest.
2. Suppose now that the players repeat the procedure in part (a) until there is an odd man. Find the probability density function of N , the number of rounds played, and identify the distribution by name.

Answer

1. The probability is $2p(1-p)$ if $k = 2$, and is $kp(1-p)^{k-1} + kp^{k-1}(1-p)$ if $k > 2$.
2. Let r_k denote the probability in part (a). N has PDF $f(n) = (1-r_k)^{n-1} r_k$ for $n \in \mathbb{N}$, and has the geometric distribution with parameter r_k .

The *odd man out* game is treated in more detail in the section on the Geometric Distribution in the chapter on Bernoulli Trials.

Cards

Recall that a *poker hand* consists of 5 cards chosen at random and without replacement from a standard deck of 52 cards. Let X denote the number of spades in the hand and Y the number of hearts in the hand. Give the probability density function of each of the following random variables, and identify the distribution by name:

1. X
2. Y
3. (X, Y)

Answer

1. $g(x) = \frac{\binom{13}{x} \binom{39}{5-x}}{\binom{52}{5}}$ for $x \in \{0, 1, 2, 3, 4, 5\}$ This is the hypergeometric distribution with population size $m = 52$, type parameter $r = 13$, and sample size $n = 5$
2. $h(y) = \frac{\binom{13}{y} \binom{39}{5-y}}{\binom{52}{5}}$ for $y \in \{0, 1, 2, 3, 4, 5\}$ This is the same hypergeometric distribution as in part (a).
3. $f(x, y) = \frac{\binom{13}{x} \binom{13}{y} \binom{26}{5-x-y}}{\binom{52}{5}}$ for $(x, y) \in \{0, 1, 2, 3, 4, 5\}^2$ with $x + y \leq 5$. This is a bivariate hypergeometric distribution.

Recall that a *bridge hand* consists of 13 cards chosen at random and without replacement from a standard deck of 52 cards. An *honor card* is a card of denomination ace, king, queen, jack or 10. Let N denote the number of honor cards in the hand.

1. Find the probability density function of N and identify the distribution by name.
2. Find the probability that the hand has no honor cards. A hand of this kind is known as a Yarborough, in honor of Second Earl of Yarborough.

Answer

1. $f(n) = \frac{\binom{20}{n} \binom{32}{13-n}}{\binom{52}{13}}$ for $n \in \{0, 1, \dots, 13\}$ This is the hypergeometric distribution with population size $m = 52$, type parameter $r = 20$ and sample size $n = 13$.
2. 0.00547

In the most common *high card point system* in bridge, an ace is worth 4 points, a king is worth 3 points, a queen is worth 2 points, and a jack is worth 1 point. Find the probability density function of V , the point value of a random bridge hand.

Reliability

Suppose that in a batch of 500 components, 20 are defective and the rest are good. A sample of 10 components is selected at random and tested. Let X denote the number of defectives in the sample.

1. Find the probability density function of X and identify the distribution by name and parameter values.
2. Find the probability that the sample contains at least one defective component.

Answer

1. $f(x) = \frac{\binom{20}{x} \binom{480}{10-x}}{\binom{500}{10}}$ for $x \in \{0, 1, \dots, 10\}$ This is the hypergeometric distribution with population size $m = 500$, type parameter $r = 20$, and sample size $n = 10$.
2. $\mathbb{P}(X \geq 1) = 1 - \frac{\binom{480}{10}}{\binom{500}{10}} \approx 0.3377$

A plant has 3 assembly lines that produce a certain type of component. Line 1 produces 50% of the components and has a defective rate of 4%; line 2 produces 30% of the components and has a defective rate of 5%; line 3 produces 20% of the components and has a defective rate of 1%. A component is chosen at random from the plant and tested.

1. Find the probability that the component is defective.
2. Given that the component is defective, find the conditional probability density function of the line that produced the component.

Answer

Let D the event that the item is defective, and $f(\cdot | D)$ the PDF of the line number given D .

1. $\mathbb{P}(D) = 0.037$
2. $f(1 | D) = 0.541, f(2 | D) = 0.405, f(3 | D) = 0.054$

Recall that in the standard model of *structural reliability*, a system consists of n components, each of which, independently of the others, is either working or failed. Let X_i denote the state of component i , where 1 means working and 0 means failed. Thus, the state vector is $\mathbf{X} = (X_1, X_2, \dots, X_n)$. The system as a whole is also either working or failed, depending only on the states of the components. Thus, the state of the system is an indicator random variable $U = u(\mathbf{X})$ that depends on the states of the components according to a *structure function* $u : \{0, 1\}^n \rightarrow \{0, 1\}$. In a *series system*, the system works if and only if every component works. In a *parallel system*, the system works if and only if at least one component works. In a *k out of n system*, the system works if and only if at least k of the n components work.

The *reliability* of a device is the probability that it is working. Let $p_i = \mathbb{P}(X_i = 1)$ denote the reliability of component i , so that $\mathbf{p} = (p_1, p_2, \dots, p_n)$ is the vector of component reliabilities. Because of the independence assumption, the system reliability depends only on the component reliabilities, according to a *reliability function* $r(\mathbf{p}) = \mathbb{P}(U = 1)$. Note that when all component reliabilities have the same value p , the states of the components form a sequence of n [Bernoulli trials](#). In this case, the system reliability is, of course, a function of the common component reliability p .

Suppose that the component reliabilities all have the same value p . Let \mathbf{X} denote the state vector and Y denote the number of working components.

1. Give the probability density function of \mathbf{X} .
2. Give the probability density function of Y and identify the distribution by name and parameter.
3. Find the reliability of the k out of n system.

Answer

1. $f(x_1, x_2, \dots, x_n) = p^y(1-p)^{n-y}$ for $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ where $y = x_1 + x_2 + \dots + x_n$
2. $g(k) = \binom{n}{y} p^y (1-p)^{n-y}$ for $y \in \{0, 1, \dots, n\}$ This is the binomial distribution with trial parameter n and success parameter p .
3. $r(p) = \sum_{i=k}^n \binom{n}{i} p^i (1-p)^{n-i}$

Suppose that we have 4 independent components, with common reliability $p = 0.8$. Let Y denote the number of working components.

1. Find the probability density function of Y explicitly.
2. Find the reliability of the parallel system.
3. Find the reliability of the 2 out of 4 system.
4. Find the reliability of the 3 out of 4 system.
5. Find the reliability of the series system.

Answer

1. $g(0) = 0.0016, g(1) = 0.0256, g(2) = 0.1536, g(3) = g(4) = 0.4096$
2. $r_{4,1} = 0.9984$
3. $r_{4,2} = 0.9729$
4. $r_{4,3} = 0.8192$
5. $r_{4,4} = 0.4096$

Suppose that we have 4 independent components, with reliabilities $p_1 = 0.6, p_2 = 0.7, p_3 = 0.8$, and $p_4 = 0.9$. Let Y denote the number of working components.

1. Find the probability density function of Y .
2. Find the reliability of the parallel system.
3. Find the reliability of the 2 out of 4 system.
4. Find the reliability of the 3 out of 4 system.
5. Find the reliability of the series system.

Answer

1. $g(0) = 0.0024, g(1) = 0.0404, g(2) = 0.2144, g(3) = 0.4404, g(4) = 0.3024$
2. $r_{4,1} = 0.9976$
3. $r_{4,2} = 0.9572$

4. $r_{4,3} = 0.7428$
5. $r_{4,4} = 0.3024$

The Poisson Distribution

Suppose that $a > 0$. Define f by

$$f(n) = e^{-a} \frac{a^n}{n!}, \quad n \in \mathbb{N} \quad (3.1.24)$$

1. f is a probability density function.
2. $f(n-1) < f(n)$ if and only if $n < a$.
3. If a is not a positive integer, there is a single mode at $\lfloor a \rfloor$.
4. If a is a positive integer, there are two modes at $a-1$ and a .

Proof

1. Recall from calculus, the exponential series

$$\sum_{n=0}^{\infty} \frac{a^n}{n!} = e^a \quad (3.1.25)$$

Hence f is a probability density function.

2. Note that $f(n-1) < f(n)$ if and only if $\frac{a^{n-1}}{(n-1)!} < \frac{a^n}{n!}$ if and only if $1 < \frac{a}{n}$.
3. By the same argument, $f(n-1) = f(n)$ if and only if $a = n$. If a is not a positive integer this cannot happen. Hence, letting $k = \lfloor a \rfloor$, it follows from (b) that $f(n) < f(k)$ if $n < k$ or $n > k$.
4. If a is a positive integer, then $f(a-1) = f(a)$. From (b), $f(n) < f(a-1)$ if $n < a-1$ and $f(n) < f(a)$ if $n > a$.

The distribution defined by the probability density function in the previous exercise is the *Poisson distribution* with parameter a , named after Simeon Poisson. Note that like the other named distributions we studied above ([hypergeometric](#) and [binomial](#)), the Poisson distribution is unimodal: the probability density function at first increases and then decreases, with either a single mode or two adjacent modes. The Poisson distribution is studied in detail in the Chapter on Poisson Processes, and is used to model the number of “random points” in a region of time or space, under certain ideal conditions. The parameter a is proportional to the size of the region of time or space.

Suppose that the customers arrive at a service station according to the Poisson model, at an average rate of 4 per hour. Thus, the number of customers N who arrive in a 2-hour period has the Poisson distribution with parameter 8.

1. Find the modes.
2. Find $\mathbb{P}(N \geq 6)$.

Answer

1. modes: 7, 8
2. $\mathbb{P}(N > 6) = 0.8088$

In the Poisson experiment, set $r = 4$ and $t = 2$. Run the simulation 1000 times and compare the empirical density function to the probability density function.

Suppose that the number of flaws N in a piece of fabric of a certain size has the Poisson distribution with parameter 2.5.

1. Find the mode.
2. Find $\mathbb{P}(N > 4)$.

Answer

1. mode: 2
2. $\mathbb{P}(N > 4) = 0.1088$

Suppose that the number of raisins N in a piece of cake has the Poisson distribution with parameter 10.

1. Find the modes.
2. Find $\mathbb{P}(8 \leq N \leq 12)$.

Answer

1. modes: 9, 10
2. $\mathbb{P}(8 \leq N \leq 12) = 0.5713$

A Zeta Distribution

Let g be the function defined by $g(n) = \frac{1}{n^2}$ for $n \in \mathbb{N}_+$.

1. Find the probability density function f that is **proportional** to g .
2. Find the mode of the distribution.
3. Find $\mathbb{P}(N \leq 5)$ where N has probability density function f .

Answer

1. $f(n) = \frac{6}{\pi^2 n^2}$ for $n \in \mathbb{N}_+$. Recall that $\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6}$
2. Mode $n = 1$
3. $\mathbb{P}(N \leq 5) = \frac{5269}{600\pi^2}$

The distribution defined in the previous exercise is a member of the *zeta* family of distributions. Zeta distributions are used to model sizes or ranks of certain types of objects, and are studied in more detail in the chapter on Special Distributions.

Benford's Law

Let f be the function defined by $f(d) = \log(d+1) - \log(d) = \log\left(1 + \frac{1}{d}\right)$ for $d \in \{1, 2, \dots, 9\}$. (The logarithm function is the base 10 *common logarithm*, not the base e *natural logarithm*.)

1. Show that f is a probability density function.
2. Compute the values of f explicitly, and sketch the graph.
3. Find $\mathbb{P}(X \leq 3)$ where X has probability density function f .

Answer

1. Note that $\sum_{d=1}^9 f(d) = \log(10) = 1$. The sum collapses.

2.	d	1	2	3	4	5	6	7	8	9
	$f(d)$	0.3010	0.1761	0.1249	0.0969	0.0792	0.0669	0.0580	0.0512	0.0458

3. $\log(4) \approx 0.6020$

The distribution defined in the previous exercise is known as *Benford's law*, and is named for the American physicist and engineer Frank Benford. This distribution governs the leading digit in many real sets of data. Benford's law is studied in more detail in the chapter on Special Distributions.

Data Analysis Exercises

In the M&M data, let R denote the number of red candies and N the total number of candies. Compute and graph the empirical probability density function of each of the following:

1. R
2. N
3. R given $N > 57$

Answer

We denote the PDF of R by f and the PDF of N by g

1.	r	3	4	5	6	8	9	10	11	12	14	15	20
----	-----	---	---	---	---	---	---	----	----	----	----	----	----

$f(r)$	$\frac{1}{30}$	$\frac{3}{30}$	$\frac{3}{30}$	$\frac{2}{30}$	$\frac{4}{30}$	$\frac{5}{30}$	$\frac{2}{30}$	$\frac{1}{30}$	$\frac{3}{30}$	$\frac{3}{30}$	$\frac{3}{30}$	$\frac{1}{30}$
--------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------

2. n	50	53	54	55	56	57	58	59	60	61
$g(n)$	$\frac{1}{30}$	$\frac{1}{30}$	$\frac{1}{30}$	$\frac{4}{30}$	$\frac{4}{30}$	$\frac{3}{30}$	$\frac{9}{30}$	$\frac{3}{30}$	$\frac{2}{30}$	$\frac{2}{30}$

3. r	3	4	6	8	9	11	12	14	15
$f(r \mid N > 57)$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{3}{16}$	$\frac{3}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{3}{16}$	$\frac{2}{16}$

In the Cicada data, let G denotes gender, S species type, and W body weight (in grams). Compute the empirical probability density function of each of the following:

1. G
2. S
3. (G, S)
4. G given $W > 0.20$ grams.

Answer

We denote the PDF of G by g , the PDF of S by h and the PDF of (G, S) by f .

1. $g(0) = \frac{59}{104}, g(1) = \frac{45}{104}$
2. $h(0) = \frac{44}{104}, h(1) = \frac{6}{104}, h(2) = \frac{54}{104}$
3. $f(0, 0) = \frac{16}{104}, f(0, 1) = \frac{3}{104}, f(0, 2) = \frac{40}{104}, f(1, 0) = \frac{28}{104}, f(1, 1) = \frac{3}{104}, f(1, 2) = \frac{14}{104}$
4. $g(0 \mid W > 0.2) = \frac{31}{73}, g(1 \mid W > 0.2) = \frac{42}{73}$

This page titled [3.1: Discrete Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.2: Continuous Distributions

In the previous section, we considered discrete distributions. In this section, we study a complementary type of distribution. As usual, if you are a new student of probability, you may want to skip the technical details.

Basic Theory

Definitions and Basic Properties

As usual, our starting point is a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$. So to review, S is the set of outcomes, \mathcal{S} the collection of events, and \mathbb{P} the probability measure on the sample space (S, \mathcal{S}) . We use the terms *probability measure* and *probability distribution* synonymously in this text. Also, since we use a general definition of random variable, every probability measure can be thought of as the probability distribution of a random variable, so we can always take this point of view if we like. Indeed, most probability measures naturally have random variables associated with them.

In this section, we assume that $S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$.

Details

Technically, S is a measurable subset of \mathbb{R}^n and \mathcal{S} is the σ -algebra measurable subsets of S . Typically in applications, S is defined by a finite number of inequalities involving elementary function.

Here is our first fundamental definition.

The probability measure \mathbb{P} is *continuous* if $\mathbb{P}(\{x\}) = 0$ for all $x \in S$.

The fact that each point is assigned probability 0 might seem impossible or paradoxical at first, but soon we will see very familiar analogies.

If \mathbb{P} is a continuous distribution then $\mathbb{P}(C) = 0$ for every countable $C \subseteq S$.

Proof

Since C is countable, it follows from the additivity axiom of probability that

$$\mathbb{P}(C) = \sum_{x \in C} \mathbb{P}(\{x\}) = 0 \quad (3.2.1)$$

Thus, continuous distributions are in complete contrast with discrete distributions, for which all of the probability mass is concentrated on the points in a discrete set. For a continuous distribution, the probability mass is *continuously* spread over S in some sense. In the picture below, the light blue shading is intended to suggest a continuous distribution of probability.




Figure 3.2.1: A continuous probability distribution on S

Typically, S is a region of \mathbb{R}^n defined by inequalities involving elementary functions, for example an interval in \mathbb{R} , a circular region in \mathbb{R}^2 , and a conical region in \mathbb{R}^3 . Suppose that \mathbb{P} is a continuous probability measure on S . The fact that each point in S has probability 0 is conceptually the same as the fact that an interval of \mathbb{R} can have positive length even though it is composed of points each of which has 0 length. Similarly, a region of \mathbb{R}^2 can have positive area even though it is composed of points (or curves) each of which has area 0. In the one-dimensional case, continuous distributions are used to model random variables that take values in intervals of \mathbb{R} , variables that can, in principle, be measured with any degree of accuracy. Such variables abound in applications and include

- length, area, volume, and distance

- time
- mass and weight
- charge, voltage, and current
- resistance, capacitance, and inductance
- velocity and acceleration
- energy, force, and work

Usually a continuous distribution can usually be described by certain type of function.

Suppose again that \mathbb{P} is a continuous distribution on S . A function $f : S \rightarrow [0, \infty)$ is a *probability density function* for \mathbb{P} if

$$\mathbb{P}(A) = \int_A f(x) dx, \quad A \in \mathcal{S} \quad (3.2.2)$$

Details

Technically, f must be measurable and is a probability density function of \mathbb{P} *with respect* to Lebesgue measure, the standard measure on \mathbb{R}^n . Moreover, the integral is the Lebesgue integral, but the ordinary Riemann integral of calculus suffices for the sets that occur in typical applications.

So the probability distribution \mathbb{P} is completely determined by the probability density function f . As a special case, note that $\int_S f(x) dx = \mathbb{P}(S) = 1$. Conversely, a nonnegative function on S with this property defines a probability measure.

A function $f : S \rightarrow [0, \infty)$ that satisfies $\int_S f(x) dx = 1$ is a *probability density function* on S and then \mathbb{P} defined as follows is a continuous probability measure on S :

$$\mathbb{P}(A) = \int_A f(x) dx, \quad A \in \mathcal{S} \quad (3.2.3)$$

Proof




Figure 3.2.2: A continuous distribution is completely determined by its probability density function

Note that we can always extend f to a probability density function on a subset of \mathbb{R}^n that contains S , or to all of \mathbb{R}^n , by defining $f(x) = 0$ for $x \notin S$. This extension sometimes simplifies notation. Put another way, we can be a bit sloppy about the “set of values” of the random variable. So for example if $a, b \in \mathbb{R}$ with $a < b$ and X has a continuous distribution on the interval $[a, b]$, then we could also say that X has a continuous distribution on (a, b) or $[a, b)$, or $(a, b]$.

The points $x \in S$ that maximize the probability density function f are important, just as in the discrete case.

Suppose that \mathbb{P} is a continuous distribution on S with probability density function f . An element $x \in S$ that maximizes f is a *mode* of the distribution.

If there is only one mode, it is sometimes used as a measure of the *center* of the distribution.

You have probably noticed that probability density functions for continuous distributions are analogous to probability density functions for discrete distributions, with integrals replacing sums. However, there are essential differences. First, every discrete distribution has a unique probability density function f given by $f(x) = \mathbb{P}(\{x\})$ for $x \in S$. For a continuous distribution, the existence of a probability density function is not guaranteed. The advanced section on absolute continuity and density functions has several examples of continuous distribution that do not have density functions, and gives conditions that are necessary and sufficient for the existence of a probability density function. Even if a probability density function f exists, it is *never* unique. Note that the values of f on a finite (or even countably infinite) set of points could be changed to other nonnegative values and the new function would still be a probability density function for the same distribution. The critical fact is that only *integrals* of f are important. Second, the values of the PDF f for a discrete distribution are probabilities, and in particular $f(x) \leq 1$ for $x \in S$. For a

continuous distribution the values are not probabilities and in fact it's possible that $f(x) > 1$ for some or even all $x \in S$. Further, f can be unbounded on S . In the typical calculus interpretation, $f(x)$ really is probability *density* at x . That is, $f(x) dx$ is approximately the probability of a "small" region of size dx about x .

Constructing Probability Density Functions

Just as in the discrete case, a nonnegative function on S can often be scaled to produce a probability density function.

Suppose that $g: S \rightarrow [0, \infty)$ and let

$$c = \int_S g(x) dx \quad (3.2.4)$$

If $0 < c < \infty$ then f defined by $f(x) = \frac{1}{c}g(x)$ for $x \in S$ defines a probability density function for a continuous distribution on S .

Proof

Technically, the function g is measurable. Technicalities aside, the proof is trivial. Clearly $f(x) \geq 0$ for $x \in S$ and

$$\int_S f(x) dx = \frac{1}{c} \int_S g(x) dx = \frac{c}{c} = 1 \quad (3.2.5)$$

Note again that f is just a scaled version of g . So this result can be used to construct probability density functions with desired properties (domain, shape, symmetry, and so on). The constant c is sometimes called the *normalizing constant* of g .

Conditional Densities

Suppose now that X is a random variable defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and that X has a continuous distribution on S . A probability density function for X is based on the underlying probability measure on the sample space (Ω, \mathcal{F}) . This measure could be a conditional probability measure, conditioned on a given event $E \in \mathcal{F}$ with $\mathbb{P}(E) > 0$. Assuming that the conditional probability density function exists, the usual notation is

$$f(x | E), \quad x \in S \quad (3.2.6)$$

Note, however, that except for notation, no new concepts are involved. The defining property is

$$\int_A f(x | E) dx = \mathbb{P}(X \in A | E), \quad A \in \mathcal{S} \quad (3.2.7)$$

and all results that hold for probability density functions in general hold for conditional probability density functions. The event E could be an event described in terms of the random variable X itself:

Suppose that X has a continuous distribution on S with probability density function f and that $B \in \mathcal{S}$ with $\mathbb{P}(X \in B) > 0$. The conditional probability density function of X given $X \in B$ is the function on B given by

$$f(x | X \in B) = \frac{f(x)}{\mathbb{P}(X \in B)}, \quad x \in B \quad (3.2.8)$$

Proof

For $A \in \mathcal{S}$ with $A \subseteq B$,

$$\int_A \frac{f(x)}{\mathbb{P}(X \in B)} dx = \frac{1}{\mathbb{P}(X \in B)} \int_A f(x) dx = \frac{\mathbb{P}(X \in A)}{\mathbb{P}(X \in B)} = \mathbb{P}(X \in A | X \in B) \quad (3.2.9)$$

Of course, $\mathbb{P}(X \in B) = \int_B f(x) dx$ and hence is the normalizing constant for the restriction of f to B , as in (8)

Examples and Applications

As always, try the problems yourself before looking at the answers.

The Exponential Distribution

Let f be the function defined by $f(t) = re^{-rt}$ for $t \in [0, \infty)$, where $r \in (0, \infty)$ is a parameter.

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph of f , and state the important qualitative features.

Proof

1. Note that $f(t) > 0$ for $t \geq 0$. Also $\int_0^\infty e^{-rt} dt = \frac{1}{r}$ so f is a PDF.
2. f is decreasing and concave upward so the mode is 0. $f(x) \rightarrow 0$ as $x \rightarrow \infty$.

The distribution defined by the probability density function in the previous exercise is called the *exponential distribution* with rate parameter r . This distribution is frequently used to model random times, under certain assumptions. Specifically, in the Poisson model of random points in time, the times between successive arrivals have independent exponential distributions, and the parameter r is the average rate of arrivals. The exponential distribution is studied in detail in the chapter on Poisson Processes.

The lifetime T of a certain device (in 1000 hour units) has the exponential distribution with parameter $r = \frac{1}{2}$. Find

1. $\mathbb{P}(T > 2)$
2. $\mathbb{P}(T > 3 \mid T > 1)$

Answer

1. $e^{-1} \approx 0.3679$
2. $e^{-1} \approx 0.3679$

In the gamma experiment, set $n = 1$ to get the exponential distribution. Vary the rate parameter r and note the shape of the probability density function. For various values of r , run the simulation 1000 times and compare the empirical density function with the probability density function.

A Random Angle

In *Bertrand's problem*, a certain random angle Θ has probability density function f given by $f(\theta) = \sin \theta$ for $\theta \in [0, \frac{\pi}{2}]$.

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph f , and state the important qualitative features.
3. Find $\mathbb{P}(\Theta < \frac{\pi}{4})$.

Answer

1. Note that $\sin \theta \geq 0$ for $0 \leq \theta \leq \frac{\pi}{2}$ and $\int_0^{\pi/2} \sin \theta d\theta = 1$.
2. f is increasing and concave downward so the mode is $\frac{\pi}{2}$.
3. $1 - \frac{1}{\sqrt{2}} \approx 0.2929$

Bertrand's problem is named for Joseph Louis Bertrand and is studied in more detail in the chapter on Geometric Models.

In Bertrand's experiment, select the model with uniform distance. Run the simulation 1000 times and compute the empirical probability of the event $\{\Theta < \frac{\pi}{4}\}$. Compare with the true probability in the previous exercise.

Gamma Distributions

Let g_n be the function defined by $g_n(t) = e^{-t} \frac{t^n}{n!}$ for $t \in [0, \infty)$ where $n \in \mathbb{N}$ is a parameter.

1. Show that g_n is a probability density function for each $n \in \mathbb{N}$.
2. Draw a careful sketch of the graph of g_n , and state the important qualitative features.

Proof

1. Note that $g_n(t) \geq 0$ for $t \geq 0$. Also, g_0 is the probability density function of the exponential distribution with parameter 1. For $n \in \mathbb{N}_+$, integration by parts with $u = t^n/n!$ and $dv = e^{-t} dt$ gives $\int_0^\infty g_n(t) dt = \int_0^\infty g_{n-1}(t) dt$. Hence it follows by induction that g_n is a PDF for each $n \in \mathbb{N}_+$.

2. g_0 is decreasing and concave downward, with mode $t = 0$. For $n > 0$, g_n increases and then decreases, with mode $t = n$. g_1 is concave downward and then upward, with inflection point at $t = 2$. For $n > 1$, g_n is concave upward, then downward, then upward again, with inflection points at $n \pm \sqrt{n}$. For all $n \in \mathbb{N}$, $g_n(t) \rightarrow 0$ as $t \rightarrow \infty$.

Interestingly, we showed in the last section on discrete distributions, that $f_t(n) = g_n(t)$ is a probability density function on \mathbb{N} for each $t \geq 0$ (it's the Poisson distribution with parameter t). The distribution defined by the probability density function g_n belongs to the family of *Erlang distributions*, named for Agner Erlang; $n + 1$ is known as the *shape parameter*. The Erlang distribution is studied in more detail in the chapter on the Poisson Process. In turn the Erlang distribution belongs to the more general family of *gamma distributions*. The gamma distribution is studied in more detail in the chapter on Special Distributions.

In the gamma experiment, keep the default rate parameter $r = 1$. Vary the shape parameter and note the shape and location of the probability density function. For various values of the shape parameter, run the simulation 1000 times and compare the empirical density function with the probability density function.

Suppose that the lifetime of a device T (in 1000 hour units) has the gamma distribution [above](#) with $n = 2$. Find each of the following:

1. $\mathbb{P}(T > 3)$.
2. $\mathbb{P}(T \leq 2)$
3. $\mathbb{P}(1 \leq T \leq 4)$

Answer

1. $\frac{17}{2}e^{-3} \approx 0.4232$
2. $1 - 5e^{-2} \approx 0.3233$
3. $\frac{5}{2}e^{-1} - 13e^{-4} \approx 0.6816$

Beta Distributions

Let f be the function defined by $f(x) = 6x(1 - x)$ for $x \in [0, 1]$.

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph of f , and state the important qualitative features.

Answer

1. Note that $f(x) \geq 0$ for $x \in [0, 1]$. Also, $\int_0^1 x(1 - x) dx = \frac{1}{6}$, so f is a PDF
2. f increases and then decreases, with mode at $x = \frac{1}{2}$. f is concave downward. f is symmetric about $x = \frac{1}{2}$ (in fact, the graph is a parabola).

Let f be the function defined by $f(x) = 12x^2(1 - x)$ for $x \in [0, 1]$.

1. Show that f is a probability density function.
2. Draw a careful sketch the graph of f , and state the important qualitative features.

Answer

1. Note that $f(x) \geq 0$ for $0 \leq x \leq 1$. Also $\int_0^1 x^2(1 - x) dx = \frac{1}{12}$, so f is a PDF.
2. f increases and then decreases, with mode at $x = \frac{2}{3}$. f is concave upward and then downward, with inflection point at $x = \frac{1}{3}$.

The distributions defined in the last two exercises are examples of *beta distributions*. These distributions are widely used to model random proportions and probabilities, and physical quantities that take values in bounded intervals (which, after a change of units, can be taken to be $[0, 1]$). Beta distributions are studied in detail in the chapter on Special Distributions.

In the special distribution simulator, select the beta distribution. For the following parameter values, note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function with the probability density function.

1. $a = 2, b = 2$. This gives the first beta distribution [above](#).
2. $a = 3, b = 2$. This gives the second beta distribution [above](#).

Suppose that P is a random proportion. Find $\mathbb{P}\left(\frac{1}{4} \leq P \leq \frac{3}{4}\right)$ in each of the following cases:

1. P has the first beta distribution [above](#).
2. P has the second beta distribution [above](#).

Answer

1. $\frac{11}{16}$
2. $\frac{11}{16}$

Let f be the function defined by

$$f(x) = \frac{1}{\pi \sqrt{x(1-x)}}, \quad x \in (0, 1) \quad (3.2.10)$$

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph of f , and state the important qualitative features.

Answer

1. Note that $f(x) > 0$ for $0 < x < 1$. Using the substitution $u = \sqrt{x}$ gives

$$\int_0^1 \frac{1}{\sqrt{x(1-x)}} dx = \int_0^1 \frac{2}{\sqrt{1-u^2}} du = 2 \arcsin u \Big|_0^1 = \pi \quad (3.2.11)$$

Thus f is a PDF.

2. f is symmetric about $x = \frac{1}{2}$. f decreases and then increases, with minimum at $x = \frac{1}{2}$. $f(x) \rightarrow \infty$ as $x \downarrow 0$ and as $x \uparrow 1$ so the distribution has no mode. f is concave upward.

The distribution defined in the last exercise is also a member of the beta family of distributions. But it is also known as the (standard) *arcsine distribution*, because of the arcsine function that arises in the proof that f is a probability density function. The arcsine distribution has applications to a very important random process known as Brownian motion, named for the Scottish botanist Robert Brown. Arcsine distributions are studied in more generality in the chapter on Special Distributions.

In the special distribution simulator, select the (continuous) arcsine distribution and keep the default parameter values. Run the simulation 1000 times and compare the empirical density function with the probability density function.

Suppose that X_t represents the change in the price of a stock at time t , relative to the value at an initial reference time 0. We treat t as a continuous variable measured in weeks. Let $T = \max\{t \in [0, 1] : X_t = 0\}$, the last time during the first week that the stock price was unchanged over its initial value. Under certain ideal conditions, T will have the arcsine distribution. Find each of the following:

1. $\mathbb{P}\left(T < \frac{1}{4}\right)$
2. $\mathbb{P}\left(T \geq \frac{1}{2}\right)$
3. $\mathbb{P}\left(T \leq \frac{3}{4}\right)$

Answer

1. $\frac{1}{3}$
2. $\frac{1}{2}$
3. $\frac{2}{3}$

Open the Brownian motion experiment and select the *last zero* variable. Run the experiment in single step mode a few times. The random process that you observe models the price of the stock in the previous exercise. Now run the experiment 1000 times and compute the empirical probability of each event in the previous exercise.

The Pareto Distribution

Let g be the function defined by $g(x) = 1/x^b$ for $x \in [1, \infty)$, where $b \in (0, \infty)$ is a parameter.

1. Draw a careful sketch the graph of g , and state the important qualitative features.
2. Find the values of b for which there exists a probability density function f (8)proportional to g . Identify the mode.

Answer

1. g is decreasing and concave upward, with $g(x) \rightarrow 0$ as $x \rightarrow \infty$.
2. Note that if $b \neq 1$

$$\int_1^\infty x^{-b} dx = \frac{x^{1-b}}{1-b} \Big|_1^\infty = \begin{cases} \infty, & 0 < b < 1 \\ \frac{1}{b-1}, & 1 < b < \infty \end{cases} \quad (3.2.12)$$

When $b = 1$ we have $\int_1^\infty x^{-1} dx = \ln x \Big|_1^\infty = \infty$. Thus, when $0 < b \leq 1$, there is no PDF proportional to g . When $b > 1$, the PDF proportional to g is $f(x) = \frac{b-1}{x^b}$ for $x \in [1, \infty)$. The mode is 1.

Note that the qualitative features of g are the same, regardless of the value of the parameter $b > 0$, but only when $b > 1$ can g be normalized into a probability density function. In this case, the distribution is known as the *Pareto distribution*, named for Vilfredo Pareto. The parameter $a = b - 1$, so that $a > 0$, is known as the *shape parameter*. Thus, the Pareto distribution with shape parameter a has probability density function

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (3.2.13)$$

The Pareto distribution is widely used to model certain economic variables and is studied in detail in the chapter on Special Distributions.

In the special distribution simulator, select the Pareto distribution. Leave the scale parameter fixed, but vary the shape parameter, and note the shape of the probability density function. For various values of the shape parameter, run the simulation 1000 times and compare the empirical density function with the probability density function.

Suppose that the income X (in appropriate units) of a person randomly selected from a population has the Pareto distribution with shape parameter $a = 2$. Find each of the following:

1. $\mathbb{P}(X > 2)$
2. $\mathbb{P}(X \leq 4)$
3. $\mathbb{P}(3 \leq X \leq 5)$

Answer

1. $\frac{1}{4}$
2. $\frac{15}{16}$
3. $\frac{16}{225}$

The Cauchy Distribution

Let f be the function defined by

$$f(x) = \frac{1}{\pi(x^2 + 1)}, \quad x \in \mathbb{R} \quad (3.2.14)$$

1. Show that f is a probability density function.
2. Draw a careful sketch the graph of f , and state the important qualitative features.

Answer

1. Note that $f(x) > 0$ for $x \in \mathbb{R}$. Also

$$\int_{-\infty}^\infty \frac{1}{1+x^2} dx = \arctan x \Big|_{-\infty}^\infty = \pi \quad (3.2.15)$$

and hence f is a PDF.

2. f increases and then decreases, with mode $x = 0$. f is concave upward, then downward, then upward again, with inflection points at $x = \pm \frac{1}{\sqrt{3}}$. f is symmetric about $x = 0$.

The distribution constructed in the previous exercise is known as the (standard) *Cauchy distribution*, named after Augustin Cauchy. It might also be called the *arctangent distribution*, because of the appearance of the arctangent function in the proof that f is a probability density function. In this regard, note the similarity to the arcsine distribution [above](#). The Cauchy distribution is studied in more generality in the chapter on Special Distributions. Note also that the Cauchy distribution is obtained by normalizing the function $x \mapsto \frac{1}{1+x^2}$; the graph of this function is known as the *witch of Agnesi*, in honor of Maria Agnesi.

In the special distribution simulator, select the Cauchy distribution with the default parameter values. Run the simulation 1000 times and compare the empirical density function with the probability density function.

A light source is 1 meter away from position 0 on an infinite, straight wall. The angle Θ that the light beam makes with the perpendicular to the wall is randomly chosen from the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$. The position $X = \tan(\Theta)$ of the light beam on the wall has the standard Cauchy distribution. Find each of the following:

1. $\mathbb{P}(-1 < X < 1)$.
2. $\mathbb{P}\left(X \geq \frac{1}{\sqrt{3}}\right)$
3. $\mathbb{P}(X \leq \sqrt{3})$

Answer

1. $\frac{1}{2}$
2. $\frac{1}{3}$
3. $\frac{2}{3}$

The Cauchy experiment (with the default parameter values) is a simulation of the experiment in the last exercise.

1. Run the experiment a few times in single step mode.
2. Run the experiment 1000 times and compare the empirical density function with the probability density function.
3. Using the data from (b), compute the relative frequency of each event in the previous exercise, and compare with the true probability.

The Standard Normal Distribution

Let ϕ be the function defined by $\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ for $z \in \mathbb{R}$.

1. Show that ϕ is a probability density function.
2. Draw a careful sketch the graph of ϕ , and state the important qualitative features.

Proof

2. Note that $\phi(z) > 0$ for $z \in \mathbb{R}$. Let $c = \int_{-\infty}^{\infty} e^{-z^2/2} dz$. Then

$$c^2 = \int_{-\infty}^{\infty} e^{-x^2/2} dx \int_{-\infty}^{\infty} e^{-y^2/2} dy = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-(x^2+y^2)/2} dx dy \quad (3.2.16)$$

Change to polar coordinates: $x = r \cos \theta$, $y = r \sin \theta$ where $r \in [0, \infty)$ and $\theta \in [0, 2\pi)$. Then $x^2 + y^2 = r^2$ and $dx dy = r dr d\theta$. Hence

$$c^2 = \int_0^{2\pi} \int_0^{\infty} r e^{-r^2/2} dr d\theta \quad (3.2.17)$$

Using the simple substitution $u = r^2$, the inner integral is $\int_0^{\infty} e^{-u/2} du = 1$. Then the outer integral is $\int_0^{2\pi} 1 d\theta = 2\pi$. Hence $c = \sqrt{2\pi}$ and so ϕ is a PDF.

3. Note that ϕ is symmetric about 0. ϕ increases and then decreases, with mode $z = 0$. ϕ is concave upward, then downward, then upward again, with inflection points at $z = \pm 1$. $\phi(z) \rightarrow 0$ as $z \rightarrow \infty$ and as $z \rightarrow -\infty$.

The distribution defined in the last exercise is the *standard normal distribution*, perhaps the most important distribution in probability and statistics. Its importance stems largely from the central limit theorem, one of the fundamental theorems in probability. In particular, normal distributions are widely used to model physical measurements that are subject to small, random errors. The family of normal distributions is studied in more generality in the chapter on Special Distributions.

In the special distribution simulator, select the normal distribution and keep the default parameter values. Run the simulation 1000 times and compare the empirical density function and the probability density function.

The function $z \mapsto e^{-z^2/2}$ is a notorious example of an integrable function that does not have an antiderivative that can be expressed in closed form in terms of other elementary functions. (That's why we had to resort to the polar coordinate trick to show that ϕ is a probability density function.) So probabilities involving the normal distribution are usually computed using mathematical or statistical software.

Suppose that the error Z in the length of a certain machined part (in millimeters) has the standard normal distribution. Use mathematical software to approximate each of the following:

1. $\mathbb{P}(-1 \leq Z \leq 1)$
2. $\mathbb{P}(Z > 2)$
3. $\mathbb{P}(Z < -3)$

Answer

1. 0.6827
2. 0.0228
3. 0.0013

The Extreme Value Distribution

Let f be the function defined by $f(x) = e^{-x}e^{-e^{-x}}$ for $x \in \mathbb{R}$.

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph of f , and state the important qualitative features.
3. Find $\mathbb{P}(X > 0)$, where X has probability density function f .

Answer

1. Note that $f(x) > 0$ for $x \in \mathbb{R}$. Using the substitution $u = e^{-x}$,

$$\int_{-\infty}^{\infty} e^{-x} e^{-e^{-x}} dx = \int_0^{\infty} e^{-u} du = 1 \quad (3.2.18)$$

(note that the integrand in the last integral is the exponential PDF with parameter 1.

2. f increases and then decreases, with mode $x = 0$. f is concave upward, then downward, then upward again, with inflection points at $x = \pm \ln[(3 + \sqrt{5})/2]$. Note however that f is not symmetric about 0. $f(x) \rightarrow 0$ as $x \rightarrow \infty$ and as $x \rightarrow -\infty$.
3. $1 - e^{-1} \approx 0.6321$

The distribution in the last exercise is the (standard) *type 1 extreme value distribution*, also known as the *Gumbel distribution* in honor of Emil Gumbel. Extreme value distributions are studied in more generality in the chapter on Special Distributions.

In the special distribution simulator, select the extreme value distribution. Keep the default parameter values and note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function with the probability density function.

The Logistic Distribution

Let f be the function defined by

$$f(x) = \frac{e^x}{(1 + e^x)^2}, \quad x \in \mathbb{R} \quad (3.2.19)$$

1. Show that f is a probability density function.
2. Draw a careful sketch the graph of f , and state the important qualitative features.
3. Find $\mathbb{P}(X > 1)$, where X has probability density function f .

Answer

1. Note that $f(x) > 0$ for $x \in \mathbb{R}$. The substitution $u = e^x$ gives

$$\int_{-\infty}^{\infty} f(x) dx = \int_0^{\infty} \frac{1}{(1+u)^2} du = 1 \quad (3.2.20)$$

2. f is symmetric about 0. f increases and then decreases with mode $x = 0$. f is concave upward, then downward, then upward again, with inflection points at $x = \pm \ln(2 + \sqrt{3})$. $f(x) \rightarrow 0$ as $x \rightarrow \infty$ and as $x \rightarrow -\infty$.
3. $\frac{1}{1+e} \approx 0.2689$

The distribution in the last exercise is the (standard) *logistic distribution*. Logistic distributions are studied in more generality in the chapter on Special Distributions.

In the special distribution simulator, select the logistic distribution. Keep the default parameter values and note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function with the probability density function.

Weibull Distributions

Let f be the function defined by $f(t) = 2te^{-t^2}$ for $t \in [0, \infty)$.

1. Show that f is a probability density function.
2. Draw a careful sketch the graph of f , and state the important qualitative features.

Answer

1. Note that $f(t) \geq 0$ for $t \geq 0$. The substitution $u = t^2$ gives $\int_0^{\infty} f(t) dt = \int_0^{\infty} e^{-u} du = 1$.
2. f increases and then decreases, with mode $t = 1/\sqrt{2}$. f is concave downward and then upward, with inflection point at $t = \sqrt{3/2}$. $f(t) \rightarrow 0$ as $t \rightarrow \infty$.

Let f be the function defined by $f(t) = 3t^2e^{-t^3}$ for $t \geq 0$.

1. Show that f is a probability density function.
2. Draw a careful sketch the graph of f , and state the important qualitative features.

Answer

1. Note that $f(t) \geq 0$ for $t \geq 0$. The substitution $u = t^3$ gives

$$\int_0^{\infty} f(t) dt = \int_0^{\infty} e^{-u} du = 1 \quad (3.2.21)$$

2. f increases and then decreases, with mode $t = (\frac{2}{3})^{1/3}$. f is concave upward, then downward, then upward again, with inflection points at $t = (1 \pm \frac{1}{3}\sqrt{7})^{1/3}$. $f(t) \rightarrow 0$ as $t \rightarrow \infty$.

The distributions in the last two exercises are examples of *Weibull distributions*, name for Waloddi Weibull. Weibull distributions are studied in more generality in the chapter on Special Distributions. They are often used to model random failure times of devices (in appropriately scaled units).

In the special distribution simulator, select the Weibull distribution. For each of the following values of the shape parameter k , note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function with the probability density function.

1. $k = 2$. This gives the first Weibull distribution [above](#).
2. $k = 3$. This gives the second Weibull distribution [above](#).

Suppose that T is the failure time of a device (in 1000 hour units). Find $\mathbb{P}(T > \frac{1}{2})$ in each of the following cases:

1. T has the first Weibull distribution [above](#).
2. T has the second Weibull distribution [above](#).

Answer

1. $e^{-1/4} \approx 0.7788$
2. $e^{-1/8} \approx 0.8825$

Additional Examples

Let f be the function defined by $f(x) = -\ln x$ for $x \in (0, 1]$.

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph of f , and state the important qualitative features.
3. Find $\mathbb{P}(\frac{1}{3} \leq X \leq \frac{1}{2})$ where X has the probability density function in (a).

Answer

1. Note that $-\ln x \geq 0$ for $0 < x \leq 1$. Integration by parts with $u = -\ln x$ and $dv = dx$ gives

$$\int_0^1 -\ln x \, dx = -x \ln x \Big|_0^1 + \int_0^1 1 \, dx = 1 \quad (3.2.22)$$

2. f is decreasing and concave upward, with $f(x) \rightarrow \infty$ as $x \downarrow 0$, so there is no mode.
3. $\frac{1}{2} \ln 2 - \frac{1}{3} \ln 3 + \frac{1}{6} \approx 0.147$

Let f be the function defined by $f(x) = 2e^{-x}(1 - e^{-x})$ for $x \in [0, \infty)$.

1. Show that f is a probability density function.
2. Draw a careful sketch of the graph of f , and give the important qualitative features.
3. Find $\mathbb{P}(X \geq 1)$ where X has the probability density function in (a).

Answer

1. Note that $f(x) > 0$ for $0 < x < \infty$. Also, $\int_0^\infty (e^{-x} - e^{-2x}) \, dx = \frac{1}{2}$, so f is a PDF.
2. f increases and then decreases, with mode $x = \ln(2)$. f is concave downward and then upward, with an inflection point at $x = \ln(4)$. $f(x) \rightarrow 0$ as $x \rightarrow \infty$.
3. $2e^{-1} - e^{-2} \approx 0.6004$

The following problems deal with two and three dimensional random vectors having continuous distributions. The idea of [normalizing](#) a function to form a probability density function is important for some of the problems. The relationship between the distribution of a vector and the distribution of its components will be discussed later, in the section on joint distributions.

Let f be the function defined by $f(x, y) = x + y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

1. Show that f is a probability density function, and identify the mode.
2. Find $\mathbb{P}(Y \geq X)$ where (X, Y) has the probability density function in (a).
3. Find the conditional density of (X, Y) given $\{X < \frac{1}{2}, Y < \frac{1}{2}\}$.

Answer

1. mode $(1, 1)$
2. $\frac{1}{2}$
3. $f(x, y \mid X < \frac{1}{2}, Y < \frac{1}{2}) = 8(x + y)$ for $0 < x < \frac{1}{2}, 0 < y < \frac{1}{2}$

Let g be the function defined by $g(x, y) = x + y$ for $0 \leq x \leq y \leq 1$.

1. Find the probability density function f that is [proportional](#) to g .
2. Find $\mathbb{P}(Y \geq 2X)$ where (X, Y) has the probability density function in (a).

Answer

1. $f(x, y) = 2(x + y), 0 \leq x \leq y \leq 1$
2. $\frac{5}{12}$

Let g be the function defined by $g(x, y) = x^2y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

1. Find the probability density function f that is **proportional** to g .
2. Find $\mathbb{P}(Y \geq X)$ where (X, Y) has the probability density function in (a).

Answer

1. $f(x, y) = 6x^2y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$
2. $\frac{2}{5}$

Let g be the function defined by $g(x, y) = x^2y$ for $0 \leq x \leq y \leq 1$.

1. Find the probability density function f that is **proportional** to g .
2. Find $P(Y \geq 2X)$ where (X, Y) has the probability density function in (a).

Answer

1. $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$
2. $\frac{1}{8}$

Let g be the function defined by $g(x, y, z) = x + 2y + 3z$ for $0 \leq x \leq 1, 0 \leq y \leq 1, 0 \leq z \leq 1$.

1. Find the probability density function f that is **proportional** to g .
2. Find $\mathbb{P}(X \leq Y \leq Z)$ where (X, Y, Z) has the probability density function in (a).

Answer

1. $f(x, y, z) = \frac{1}{3}(x + 2y + 3z)$ for $0 \leq x \leq 1, 0 \leq y \leq 1, 0 \leq z \leq 1$
2. $\frac{7}{36}$

Let g be the function defined by $g(x, y) = e^{-x}e^{-y}$ for $0 \leq x \leq y < \infty$.

1. Find the probability density function f that is **proportional** to g .
2. Find $\mathbb{P}(X + Y < 1)$ where (X, Y) has the probability density function in (a).

Answer

1. $f(x, y) = 2e^{-x}e^{-y}, 0 < x < y < \infty$
2. $1 - 2e^{-1} \approx 0.2642$

Continuous Uniform Distributions

Our next discussion will focus on an important class of continuous distributions that are defined purely in terms of geometry. We need a preliminary definition.

For $n \in \mathbb{N}_+$, the standard measure λ_n on \mathbb{R}^n is given by

$$\lambda_n(A) = \int_A 1 \, dx, \quad A \subseteq \mathbb{R}^n \quad (3.2.23)$$

In particular, $\lambda_1(A)$ is the *length* of $A \subseteq \mathbb{R}$, $\lambda_2(A)$ is the *area* of $A \subseteq \mathbb{R}^2$, and $\lambda_3(A)$ is the *volume* of $A \subseteq \mathbb{R}^3$.

Details

Technically, λ_n is Lebesgue measure on the σ -algebra of measurable subsets of \mathbb{R}^n . The name is in honor of Henri Lebesgue. The representation above in terms of the standard Riemann integral of calculus works for the sets that occur in typical applications. For the remainder of this discussion, we assume that all subsets of \mathbb{R}^n that are mentioned are measurable

Note that if $n > 1$, the integral above is a multiple integral. Generally, $\lambda_n(A)$ is referred to as the *n-dimensional volume* of $A \subseteq \mathbb{R}^n$.

Suppose that $S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$ with $0 < \lambda_n(S) < \infty$.

1. the function f defined by $f(x) = 1/\lambda_n(S)$ for $x \in S$ is a probability density function on S .
2. The probability measure associated with f is given by $\mathbb{P}(A) = \lambda_n(A)/\lambda_n(S)$ for $A \subseteq S$, and is known as the *uniform distribution* on S .

Proof

The proof is simple: Clearly $f(x) > 0$ for $x \in S$ and

$$\int_A f(x) dx = \frac{1}{\lambda_n(S)} \int_A 1 dx = \frac{\lambda_n(A)}{\lambda_n(S)}, \quad A \subseteq S \quad (3.2.24)$$

In particular, when $A = S$ we have $\int_S f(x) dx = 1$.

Note that the probability assigned to a set $A \subseteq \mathbb{R}^n$ is proportional to the size of A , as measured by λ_n . Note also that in both the discrete and continuous cases, the uniform distribution on a set S has constant probability density function on S . The uniform distribution on a set S governs a point X chosen “at random” from S , and in the continuous case, such distributions play a fundamental role in various Geometric Models. Uniform distributions are studied in more generality in the chapter on Special Distributions.

The most important special case is the uniform distribution on an interval $[a, b]$ where $a, b \in \mathbb{R}$ and $a < b$. In this case, the probability density function is

$$f(x) = \frac{1}{b-a}, \quad a \leq x \leq b \quad (3.2.25)$$

This distribution models a point chosen “at random” from the interval. In particular, the uniform distribution on $[0, 1]$ is known as the *standard uniform distribution*, and is very important because of its simplicity and the fact that it can be transformed into a variety of other probability distributions on \mathbb{R} . Almost all computer languages have procedures for simulating independent, standard uniform variables, which are called *random numbers* in this context.

Conditional distributions corresponding to a uniform distribution are also uniform.

Suppose that $R \subseteq S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$, and that $\lambda_n(R) > 0$ and $\lambda_n(S) < \infty$. If \mathbb{P} is the uniform distribution on S , then the conditional distribution given R is uniform on R .

Proof

The proof is very simple: For $A \subseteq R$,

$$\mathbb{P}(A | R) = \frac{\mathbb{P}(A \cap R)}{\mathbb{P}(R)} = \frac{\mathbb{P}(A)}{\mathbb{P}(R)} = \frac{\lambda_n(A)/\lambda_n(S)}{\lambda_n(R)/\lambda_n(S)} = \frac{\lambda_n(A)}{\lambda_n(R)} \quad (3.2.26)$$

The last theorem has important implications for simulations. If we can simulate a random variable that is uniformly distributed on a set, we can simulate a random variable that is uniformly distributed on a subset.

Suppose again that $R \subseteq S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$, and that $\lambda_n(R) > 0$ and $\lambda_n(S) < \infty$. Suppose further that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables, each uniformly distributed on S . Let $N = \min\{k \in \mathbb{N}_+ : X_k \in R\}$. Then

1. N has the geometric distribution on \mathbb{N}_+ with success parameter $p = \lambda_n(R)/\lambda_n(S)$.
2. X_N is uniformly distributed on R .

Proof

1. Since the variables are uniformly distributed on S , $\mathbb{P}(X_k \in R) = \lambda_n(R)/\lambda_n(S)$ for each $k \in \mathbb{N}_+$. Since the variables are independent, each point is in R or not independently. Hence N , the index of the first point to fall in R , has the geometric distribution on \mathbb{N}_+ with success probability $p = \lambda_n(R)/\lambda_n(S)$. That is, $\mathbb{P}(N = k) = (1-p)^{k-1}p$ for $k \in \mathbb{N}_+$.
2. Note that $p \in (0, 1]$, so $\mathbb{P}(N \in \mathbb{N}_+) = 1$ and hence X_N is well defined. We know from our work on independence and conditional probability that the distribution of X_N is the same as the conditional distribution of X given $X \in R$, which by the previous theorem, is uniformly distributed on R .

Suppose in particular that S is a Cartesian product of n bounded intervals. It turns out to be quite easy to simulate a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots)$ each of which is uniformly distributed on S . Thus, the last theorem give an algorithm for simulating a random variable that is uniformly distributed on an irregularly shaped region $R \subseteq S$ (assuming that we have an algorithm for recognizing when a point $x \in \mathbb{R}^n$ falls in R). This method of simulation is known as the *rejection method*, and as we will see in subsequent sections, is more important that might first appear.




Figure 3.2.3: With a sequence of independent, uniformly distributed points in S , the first one to fall in R is uniformly distributed on R .

In the simple probability experiment, random points are uniformly distributed on the rectangular region S . Move and resize the events A and B and note how the probabilities of the 16 events that can be constructed from A and B change. Run the experiment 1000 times and note the agreement between the relative frequencies of the events and the probabilities of the events.

Suppose that (X, Y) is uniformly distributed on the circular region of radius 5, centered at the origin. We can think of (X, Y) as the position of a dart thrown “randomly” at a target. Let $R = \sqrt{X^2 + Y^2}$, the distance from the center to (X, Y) .

1. Give the probability density function of (X, Y) .
2. Find $\mathbb{P}(n \leq R \leq n+1 \text{ for } n \in \{0, 1, 2, 3, 4\})$.

Answer

1. $f(x, y) = \frac{1}{25\pi}$ for $\{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 25\}$
2. $\mathbb{P}(n \leq R \leq n+1) = \frac{2n+1}{25}$ for $n \in \{0, 1, 2, 3, 4\}$

Suppose that (X, Y, Z) is uniformly distributed on the cube $S = [0, 1]^3$. Find $\mathbb{P}(X < Y < Z)$ in two ways:

1. Using the probability density function.
2. Using a combinatorial argument.

Answer

1. $\mathbb{P}(X < Y < Z) = \int_0^1 \int_0^z \int_0^y 1 \, dx \, dy \, dz = \frac{1}{6}$
2. Each of the 6 strict orderings of (X, Y, Z) are equally likely, so $\mathbb{P}(X < Y < Z) = \frac{1}{6}$

The time T (in minutes) required to perform a certain job is uniformly distributed over the interval $[15, 60]$

1. Find the probability that the job requires more than 30 minutes
2. Given that the job is not finished after 30 minutes, find the probability that the job will require more than 15 additional minutes.

Answer

1. $\frac{2}{3}$
2. $\frac{1}{6}$

Data Analysis Exercises

If D is a data set from a variable X with a continuous distribution, then an *empirical density function* can be computed by partitioning the data range into subsets of small size, and then computing the probability density of points in each subset. Empirical probability density functions are studied in more detail in the chapter on Random Samples.

For the cicada data, BW denotes body weight (in grams), BL body length (in millimeters), and G gender (0 for female and 1 for male). Construct an empirical density function for each of the following and display each as a bar graph:

1. BW
2. BL
3. BW given $G = 0$

Answer

1.	BW	(0, 0.1]	(0.1, 0.2]	(0.2, 0.3]	(0.3, 0.4]
	Density	0.8654	5.8654	3.0769	0.1923
2.	BL	(15, 29]	(20, 25]	(25, 30]	(30, 35]
	Density	0.0058	0.1577	0.0346	0.0019
3.	BW	(0, 0.1]	(0.1, 0.2]	(0.2, 0.3]	(0.3, 0.4]
	Density given $G = 0$	0.3390	4.4068	5.0847	0.1695

This page titled [3.2: Continuous Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.3: Mixed Distributions

In the previous two sections, we studied discrete probability measures and continuous probability measures. In this section, we will study probability measures that are combinations of the two types. Once again, if you are a new student of probability you may want to skip the technical details.

Basic Theory

Definitions and Basic Properties

Our starting point is a random experiment modeled by a probability space $(S, \mathcal{S}, \mathbb{P})$. So to review, S is the set of outcomes, \mathcal{S} the collection of events, and \mathbb{P} the probability measure on the sample space (S, \mathcal{S}) . We use the terms *probability measure* and *probability distribution* synonymously in this text. Also, since we use a general definition of random variable, every probability measure can be thought of as the probability distribution of a random variable, so we can always take this point of view if we like. Indeed, most probability measures naturally have random variables associated with them. Here is the main definition:

The probability measure \mathbb{P} is of *mixed type* if S can be partitioned into events D and C with the following properties:

1. D is countable, $0 < \mathbb{P}(D) < 1$ and $\mathbb{P}(\{x\}) > 0$ for every $x \in D$.
2. $C \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$ and $\mathbb{P}(\{x\}) = 0$ for every $x \in C$.

Details

Recall that the term *partition* means that D and C are disjoint and $S = D \cup C$. As always, the collection of events \mathcal{S} is required to be a σ -algebra. The set C is a measurable subset of \mathbb{R}^n and then the elements of \mathcal{S} have the form $A \cup B$ where $A \subseteq D$ and B is a measurable subset of C . Typically in applications, C is defined by a finite number of inequalities involving elementary functions.

Often the discrete set D is a subset of \mathbb{R}^n also, but that's not a requirement. Note that since D and C are complements, $0 < \mathbb{P}(C) < 1$ also. Thus, part of the distribution is concentrated at points in a discrete set D ; the rest of the distribution is continuously spread over C . In the picture below, the light blue shading is intended to represent a continuous distribution of probability while the darker blue dots are intended to represent points of positive probability.




Figure 3.3.1: A mixed distribution on S

The following result is essentially equivalent to the definition.

Suppose that \mathbb{P} is a probability measure on S of mixed type as in (1).

1. The conditional probability measure $A \mapsto \mathbb{P}(A | D) = \mathbb{P}(A)/\mathbb{P}(D)$ for $A \subseteq D$ is a discrete distribution on D
2. The conditional probability measure $A \mapsto \mathbb{P}(A | C) = \mathbb{P}(A)/\mathbb{P}(C)$ for $A \subseteq C$ is a continuous distribution on C .

Proof

In general, conditional probability measures really are probability measures, so the results are obvious since $\mathbb{P}(\{x\} | D) > 0$ for x in the countable set D , and $\mathbb{P}(\{x\} | C) = 0$ for $x \in C$. From another point of view, \mathbb{P} restricted to subsets of D and \mathbb{P} restricted to subsets of C are both finite measures and so can be normalized to produce probability measures.

Note that

$$\mathbb{P}(A) = \mathbb{P}(D)\mathbb{P}(A | D) + \mathbb{P}(C)\mathbb{P}(A | C), \quad A \in \mathcal{S} \quad (3.3.1)$$

Thus, the probability measure P really is a *mixture* of a discrete distribution and a continuous distribution. Mixtures are studied in more generality in the section on conditional distributions. We can define a function on D that is a *partial probability density*

function for the discrete part of the distribution.

Suppose that \mathbb{P} is a probability measure on S of mixed type as in (1). Let g be the function defined by $g(x) = \mathbb{P}(\{x\})$ for $x \in D$. Then

1. $g(x) \geq 0$ for $x \in D$
2. $\sum_{x \in D} g(x) = \mathbb{P}(D)$
3. $\mathbb{P}(A) = \sum_{x \in A} g(x)$ for $A \subseteq D$

Proof

These results follow from the axioms of probability.

1. $g(x) = \mathbb{P}(\{x\}) \geq 0$ since probabilities are nonnegative.
2. $\sum_{x \in D} g(x) = \sum_{x \in D} \mathbb{P}(\{x\}) = \mathbb{P}(D)$ by countable additivity.
3. $\sum_{x \in A} g(x) = \sum_{x \in A} \mathbb{P}(\{x\}) = \mathbb{P}(A)$ for $A \subseteq D$, again by countable additivity.

Technically, g is a density function with respect to counting measure $\#$ on D , the standard measure used for discrete spaces.

Clearly, the normalized function $x \mapsto g(x)/\mathbb{P}(D)$ is the probability density function of the conditional distribution given D , discussed in (2). Often, the continuous part of the distribution is also described by a partial probability density function.

A *partial probability density function* for the continuous part of \mathbb{P} is a nonnegative function $h : C \rightarrow [0, \infty)$ such that

$$\mathbb{P}(A) = \int_A h(x) dx, \quad A \in \mathcal{C} \quad (3.3.2)$$

Details

Technically, h is required to be measurable, and is a density function with respect to Lebesgue measure λ_n on C , the standard measure on \mathbb{R}^n .

Clearly, the normalized function $x \mapsto h(x)/\mathbb{P}(C)$ is the probability density function of the conditional distribution given C discussed in (2). As with purely continuous distributions, the existence of a probability density function for the continuous part of a mixed distribution is not guaranteed. And when it does exist, a density function for the continuous part is not unique. Note that the values of h could be changed to other nonnegative values on a countable subset of C , and the displayed equation above would still hold, because only *integrals* of h are important. The probability measure \mathbb{P} is completely determined by the partial probability density functions.

Suppose that \mathbb{P} has partial probability density functions g and h for the discrete and continuous parts, respectively. Then

$$\mathbb{P}(A) = \sum_{x \in A \cap D} g(x) + \int_{A \cap C} h(x) dx, \quad A \in \mathcal{S} \quad (3.3.3)$$

Proof




Figure 3.3.2: A mixed distribution is completely determined by its partial density functions.

Truncated Variables

Distributions of mixed type occur naturally when a random variable with a continuous distribution is *truncated* in a certain way. For example, suppose that T is the random lifetime of a device, and has a continuous distribution with probability density function f that is positive on $[0, \infty)$. In a test of the device, we can't wait forever, so we might select a positive constant a and record the random variable U , defined by *truncating* T at a , as follows:

$$U = \begin{cases} T, & T < a \\ a, & T \geq a \end{cases} \quad (3.3.4)$$

U has a mixed distribution. In the notation above,

1. $D = \{a\}$ and $g(a) = \int_a^\infty f(t) dt$
2. $C = [0, a)$ and $h(t) = f(t)$ for $t \in [0, a)$

Suppose next that random variable X has a continuous distribution on \mathbb{R} , with probability density function f that is positive on \mathbb{R} . Suppose also that $a, b \in \mathbb{R}$ with $a < b$. The variable is *truncated* on the interval $[a, b]$ to create a new random variable Y as follows:

$$Y = \begin{cases} a, & X \leq a \\ X, & a < X < b \\ b, & X \geq b \end{cases} \quad (3.3.5)$$

Y has a mixed distribution. In the notation above,

1. $D = \{a, b\}$, $g(a) = \int_{-\infty}^a f(x) dx$, $g(b) = \int_b^\infty f(x) dx$
2. $C = (a, b)$ and $h(x) = f(x)$ for $x \in (a, b)$

Another way that a “mixed” probability distribution can occur is when we have a pair of random variables (X, Y) for our experiment, one with a discrete distribution and the other with a continuous distribution. This setting is explored in the next section on Joint Distributions.

Examples and Applications

Suppose that X has probability $\frac{1}{2}$ uniformly distributed on the set $\{1, 2, \dots, 8\}$ and has probability $\frac{1}{2}$ uniformly distributed on the interval $[0, 10]$. Find $\mathbb{P}(X > 6)$.

Answer

$$\frac{13}{40}$$

Suppose that (X, Y) has probability $\frac{1}{3}$ uniformly distributed on $\{0, 1, 2\}^2$ and has probability $\frac{2}{3}$ uniformly distributed on $[0, 2]^2$. Find $\mathbb{P}(Y > X)$.

Answer

$$\frac{4}{9}$$

Suppose that the lifetime T of a device (in 1000 hour units) has the exponential distribution with probability density function $f(t) = e^{-t}$ for $0 \leq t < \infty$. A test of the device is terminated after 2000 hours; the truncated lifetime U is recorded. Find each of the following:

1. $\mathbb{P}(U < 1)$
2. $\mathbb{P}(U = 2)$

Answer

1. $1 - e^{-1} \approx 0.6321$
2. $e^{-2} \approx 0.1353$

This page titled [3.3: Mixed Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.4: Joint Distributions

The purpose of this section is to study how the distribution of a pair of random variables is related to the distributions of the variables individually. If you are a new student of probability you may want to skip the technical details.

Basic Theory

Joint and Marginal Distributions

As usual, we start with a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} is the collection of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . Suppose now that X and Y are random variables for the experiment, and that X takes values in S while Y takes values in T . We can think of (X, Y) as a random variable taking values in the product set $S \times T$. The purpose of this section is to study how the distribution of (X, Y) is related to the distributions of X and Y individually.

Recall that

1. The distribution of (X, Y) is the probability measure on $S \times T$ given by $\mathbb{P}[(X, Y) \in C]$ for $C \subseteq S \times T$.
2. The distribution of X is the probability measure on S given by $\mathbb{P}(X \in A)$ for $A \subseteq S$.
3. The distribution of Y is the probability measure on T given by $\mathbb{P}(Y \in B)$ for $B \subseteq T$.

In this context, the distribution of (X, Y) is called the *joint distribution*, while the distributions of X and of Y are referred to as *marginal distributions*.

Details

The sets S and T come with σ -algebras of admissible subsets \mathcal{S} and \mathcal{T} , respectively, just as the collection of events \mathcal{F} is a σ -algebra. The Cartesian product set $S \times T$ is given the product σ -algebra $\mathcal{S} \otimes \mathcal{T}$ generated by products $A \times B$ where $A \in \mathcal{S}$ and $B \in \mathcal{T}$. The random variables X and Y are measurable, which ensures that (X, Y) is also a random variable (that is, measurable). Moreover, the distribution of (X, Y) is uniquely determined by probabilities of the form $\mathbb{P}[(X, Y) \in A \times B] = \mathbb{P}(X \in A, Y \in B)$ where $A \in \mathcal{S}$ and $B \in \mathcal{T}$. As usual the spaces (S, \mathcal{S}) and (T, \mathcal{T}) each fall into the two classes we have studied in the previous sections:

1. *Discrete*: the set is countable and the σ -algebra consists of all subsets.
2. *Euclidean*: the set is a measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$ and the σ -algebra consists of the measurable subsets.

The first simple but very important point, is that the marginal distributions can be obtained from the joint distribution.

Note that

1. $\mathbb{P}(X \in A) = \mathbb{P}[(X, Y) \in A \times T]$ for $A \subseteq S$
2. $\mathbb{P}(Y \in B) = \mathbb{P}[(X, Y) \in S \times B]$ for $B \subseteq T$

The converse does not hold in general. The joint distribution contains much more information than the marginal distributions separately. However, the converse does hold if X and Y are independent, as we will show [below](#).

Joint and Marginal Densities

Recall that probability distributions are often described in terms of probability density functions. Our goal is to study how the probability density functions of X and Y individually are related to probability density function of (X, Y) . But first we need to make sure that we understand our starting point.

We assume that (X, Y) has density function $f : S \times T \rightarrow (0, \infty)$ in the following sense:

1. If X and Y have discrete distributions on the countable sets S and T respectively, then f is defined by

$$f(x, y) = \mathbb{P}(X = x, Y = y), \quad (x, y) \in S \times T \quad (3.4.1)$$

2. If X and Y have continuous distributions on $S \subseteq \mathbb{R}^j$ and $T \subseteq \mathbb{R}^k$ respectively, then f is defined by the condition

$$\mathbb{P}[(X, Y) \in C] = \int_C f(x, y) d(x, y), \quad C \subseteq S \times T \quad (3.4.2)$$

3. In the mixed case where X has a discrete distribution on the countable set S and Y has a continuous distribution on $T \subseteq \mathbb{R}^k$, then f is defined by the condition

$$\mathbb{P}(X = x, Y \in B) = \int_B f(x, y) dy, \quad x \in S, B \subseteq T \quad (3.4.3)$$

4. In the mixed case where X has a continuous distribution on $S \subseteq \mathbb{R}^j$ and Y has a discrete distribution on the countable set T , then f is defined by the condition

$$\mathbb{P}(X \in A, Y = y) = \int_A f(x, y), dx, \quad A \subseteq S, y \in T \quad (3.4.4)$$

Details

1. In this case, (X, Y) has a discrete distribution on the countable set $S \times T$ and f is the density function with respect to counting measure $\#$ on $S \times T$.
2. In this case, (X, Y) has a continuous distribution on $S \times T \subseteq \mathbb{R}^{j+k}$ and f is the density function with respect to Lebesgue measure λ_{j+k} on $S \times T$. Lebesgue measure, named for Henri Lebesgue is the standard measure on Euclidean spaces.
3. In this case, (X, Y) actually has a continuous distribution:

$$\mathbb{P}[(X, Y) = (x, y)] = \mathbb{P}(X = x, Y = y) \leq \mathbb{P}(Y = y) = 0, \quad (x, y) \in S \times T \quad (3.4.5)$$

The function f is the density function with respect to the product measure formed from counting measure $\#$ on S and Lebesgue measure λ_k on T .

4. This case is just like (c) but with the roles of S and T reversed. Once again, (X, Y) has a continuous distribution and f is the density function with respect to the product measure on $S \times T$ formed by Lebesgue measure λ_j on S and counting measure $\#$ on T .

In cases (b), (c), and (d), the existence of a probability density function is not guaranteed, but is an assumption that we are making. All four cases (and many others) can be unified under the general theories of measure and integration.

First we will see how to obtain the probability density function of one variable when the other has a discrete distribution.

Suppose that (X, Y) has probability density function f as described [above](#).

1. If Y has a discrete distribution on the countable set T , then X has probability density function g given by $g(x) = \sum_{y \in T} f(x, y)$ for $x \in S$
2. If X has a discrete distribution on the countable set S , then Y has probability density function h given by $h(y) = \sum_{x \in S} f(x, y)$, $y \in T$

Proof

The two results are symmetric, so we will prove (a). The main tool is the countable additivity property of probability. Suppose first that X also has a discrete distribution on the countable set S . Then for $x \in S$,

$$g(x) = \mathbb{P}(X = x) = \mathbb{P}(X = x, Y \in T) = \sum_{y \in T} \mathbb{P}(X = x, Y = y) = \sum_{y \in T} f(x, y) \quad (3.4.6)$$

Suppose next that X has a continuous distribution on $S \subseteq \mathbb{R}^j$. Then for $A \subseteq \mathbb{R}^j$,

$$\mathbb{P}(X \in A) = \mathbb{P}(X \in A, Y \in T) = \sum_{y \in T} \mathbb{P}(X \in A, Y = y) = \sum_{y \in T} \int_A f(x, y) dx = \int_A \sum_{y \in T} f(x, y), dx \quad (3.4.7)$$

The interchange of sum and integral is allowed since f is nonnegative. By the meaning of the term, X has probability density function g given by $g(x) = \sum_{y \in T} f(x, y)$ for $x \in S$

Next we will see how to obtain the probability density function of one variable when the other has a continuous distribution.

Suppose again that (X, Y) has probability density function f as described [above](#).

1. If Y has a continuous distribution on $T \subseteq \mathbb{R}^k$ then X has probability density function g given by $g(x) = \int_T f(x, y) dy$, $x \in S$
2. If X has a continuous distribution on $S \subseteq \mathbb{R}^j$ then Y has probability density function h given by $h(y) = \int_S f(x, y) dx$, $y \in T$

Proof

Again, the results are symmetric, so we show (a). Suppose first that X has a discrete distribution on the countable set S . Then for $x \in S$

$$g(x) = \mathbb{P}(X = x) = \mathbb{P}(X = x, Y \in T) = \int_T f(x, y) dy \quad (3.4.8)$$

Next suppose that X has a continuous distribution on $S \subseteq \mathbb{R}^j$. Then for $A \subseteq S$,

$$\mathbb{P}(X \in A) = \mathbb{P}(X \in A, Y \in T) = \mathbb{P}[(X, Y) \in A \times T] = \int_{A \times T} f(x, y) d(x, y) = \int_A \int_T f(x, y) dy \quad (3.4.9)$$

Hence by the very meaning of the term, X has probability density function g given by $g(x) = \int_T f(x, y) dy$ for $x \in S$. Writing the double integral as an iterated integral is a special case of Fubini's theorem, named for Guido Fubini.

In the context of the previous two theorems, f is called the *joint probability density function* of (X, Y) , while g and h are called the *marginal density functions* of X and of Y , respectively. Some of the [computational exercises](#) below will make the term *marginal* clear.

Independence

When the variables are independent, the marginal distributions determine the joint distribution.

If X and Y are independent, then the distribution of X and the distribution of Y determine the distribution of (X, Y) .

Proof

If X and Y are independent then,

$$\mathbb{P}[(X, Y) \in A \times B] = \mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B) \quad A \in \mathcal{S}, B \in \mathcal{T} \quad (3.4.10)$$

and as noted in the details for (1), this completely determines the distribution (X, Y) on $S \times T$.

When the variables are independent, the joint density is the product of the marginal densities.

Suppose that X and Y are independent and have probability density function g and h respectively. Then (X, Y) has probability density function f given by

$$f(x, y) = g(x)h(y), \quad (x, y) \in S \times T \quad (3.4.11)$$

Proof

The main tool is the fact that an event defined in terms of X is independent of an event defined in terms of Y .

1. Suppose that X and Y have discrete distributions on the countable sets S and T respectively. Then for $(x, y) \in S \times T$,

$$\mathbb{P}[(X, Y) = (x, y)] = \mathbb{P}(X = x, Y = y) = \mathbb{P}(X = x)\mathbb{P}(Y = y) = g(x)h(y) \quad (3.4.12)$$

2. Suppose next that X and Y have continuous distributions on $S \subseteq \mathbb{R}^j$ and $T \subseteq \mathbb{R}^k$ respectively. Then for $A \subseteq S$ and $B \subseteq T$.

$$\mathbb{P}[(X, Y) \in A \times B] = \mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B) = \int_A g(x) dx \int_B h(y) dy = \int_{A \times B} g(x)h(y) d(x, y) \quad (3.4.13)$$

As noted in the details for (1), a probability measure on $S \times T$ is completely determined by its values on product sets, so it follows that $\mathbb{P}[(X, Y) \in C] = \int_C f(x, y) d(x, y)$ for general $C \subseteq S \times T$. Hence (X, Y) has PDF f .

3. Suppose next that X has a discrete distribution on the countable set S and that Y has a continuous distribution on $T \subseteq \mathbb{R}^k$. If $x \in S$ and $B \subseteq T$,

$$\mathbb{P}(X = x, Y \in B) = \mathbb{P}(X = x)\mathbb{P}(Y \in B) = g(x) \int_B h(y) dy = \int_B g(x)h(y) dy \quad (3.4.14)$$

so again it follows that (X, Y) has PDF f . The case where X has a continuous distribution on $S \subseteq \mathbb{R}^j$ and Y has a discrete distribution on the countable set T is analogous.

The following result gives a converse to the [last result](#). If the joint probability density factors into a function of x only and a function of y only, then X and Y are independent, and we can *almost* identify the individual probability density functions just from the factoring.

Factoring Theorem. Suppose that (X, Y) has probability density function f of the form

$$f(x, y) = u(x)v(y), \quad (x, y) \in S \times T \quad (3.4.15)$$

where $u: S \rightarrow [0, \infty)$ and $v: T \rightarrow [0, \infty)$. Then X and Y are independent, and there exists a positive constant c such that X and Y have probability density functions g and h , respectively, given by

$$g(x) = c u(x), \quad x \in S \quad (3.4.16)$$

$$h(y) = \frac{1}{c} v(y), \quad y \in T \quad (3.4.17)$$

Proof

Note that the proofs in the various cases are essentially the same, except for sums in the discrete case and integrals in the continuous case.

1. Suppose that X and Y have discrete distributions on the countable sets S and T , respectively, so that (X, Y) has a discrete distribution on $S \times T$. In this case, the assumption is

$$\mathbb{P}(X = x, Y = y) = u(x)v(y), \quad (x, y) \in S \times T \quad (3.4.18)$$

Summing over $y \in T$ in the displayed equation gives $g(x) = \mathbb{P}(X = x) = cu(x)$ for $x \in S$ where $c = \sum_{y \in T} v(y)$. Similarly, summing over $x \in S$ in the displayed equation gives $h(y) = \mathbb{P}(Y = y) = kv(y)$ for $y \in T$ where $k = \sum_{x \in S} u(x)$. Summing over $(x, y) \in S \times T$ in the displayed equation gives $1 = ck$ so $k = 1/c$. Finally, substituting gives $\mathbb{P}(X = x, Y = y) = \mathbb{P}(X = x)\mathbb{P}(Y = y)$ for $(x, y) \in S \times T$ so X and Y are independent.

2. Suppose next that X and Y have continuous distributions on $S \subseteq \mathbb{R}^j$ and $T \subseteq \mathbb{R}^k$ respectively. For $A \subseteq S$ and $B \subseteq T$,

$$\mathbb{P}(X \in A, Y \in B) = \mathbb{P}[(X, Y) \in A \times B] = \int_{A \times B} f(x, y) d(x, y) = \int_A u(x) dx \int_B v(y) dy \quad (3.4.19)$$

Letting $B = T$ in the displayed equation gives $\mathbb{P}(X \in A) = \int_A c u(x) dx$ for $A \subseteq S$, where $c = \int_T v(y) dy$. By definition, X has PDF $g = c u$. Next, letting $A = S$ in the displayed equation gives $\mathbb{P}(Y \in B) = \int_B k v(y) dy$ for $B \subseteq T$, where $k = \int_S u(x) dx$. Thus, Y has PDF $g = k v$. Next, letting $A = S$ and $B = T$ in the displayed equation gives $1 = c k$, so $k = 1/c$. Now note that the displayed equation holds with u replaced by g and v replaced by h , and this in turn gives $\mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B)$, so X and Y are independent.

3. Suppose next that X has a discrete distribution on the countable set S and that Y has a continuous distributions on $T \subseteq \mathbb{R}^k$. For $x \in S$ and $B \subseteq T$,

$$\mathbb{P}(X = x, Y \in B) = \int_B f(x, y) dy = u(x) \int_B v(y) dy \quad (3.4.20)$$

Letting $B = T$ in the displayed equation gives $\mathbb{P}(X = x) = c u(x)$ for $x \in S$, where $c = \int_T v(y) dy$. So X has PDF $g = c u$. Next, summing over $x \in S$ in the displayed equation gives $\mathbb{P}(Y \in B) = k \int_B v(y) dy$ for $B \subseteq T$, where $k = \sum_{x \in S} u(x)$. Thus, Y has PDF $g = k v$. Next, summing over $x \in S$ and letting $B = T$ in the displayed equation gives $1 = c k$, so $k = 1/c$. Now note that the displayed equation holds with u replaced by g and v replaced by h , and this in turn gives $\mathbb{P}(X = x, Y \in B) = \mathbb{P}(X = x)\mathbb{P}(Y \in B)$, so X and Y are independent. The case where X has a continuous distribution on $S \subseteq \mathbb{R}^j$ and Y has a discrete distribution on the countable set T is analogous.

The last two results extend to more than two random variables, because X and Y themselves may be random vectors. Here is the explicit statement:

Suppose that X_i is a random variable taking values in a set R_i with probability density function g_i for $i \in \{1, 2, \dots, n\}$, and that the random variables are independent. Then the random vector $\mathbf{X} = (X_1, X_2, \dots, X_n)$ taking values in $S = R_1 \times R_2 \times \dots \times R_n$ has probability density function f given by

$$f(x_1, x_2, \dots, x_n) = g_1(x_1)g_2(x_2) \cdots g_n(x_n), \quad (x_1, x_2, \dots, x_n) \in S \quad (3.4.21)$$

The special case where the distributions are all the same is particularly important.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent random variables, each taking values in a set R and with common probability density function g . Then the probability density function f of \mathbf{X} on $S = R^n$ is given by

$$f(x_1, x_2, \dots, x_n) = g(x_1)g(x_2) \cdots g(x_n), \quad (x_1, x_2, \dots, x_n) \in S \quad (3.4.22)$$

In probability jargon, \mathbf{X} is a sequence of *independent, identically distributed* variables, a phrase that comes up so often that it is often abbreviated as *IID*. In statistical jargon, \mathbf{X} is a *random sample* of size n from the common distribution. As is evident from the special terminology, this situation is very important in both branches of mathematics. In statistics, the joint probability density function f plays an important role in procedures such as maximum likelihood and the identification of uniformly best estimators.

Recall that (mutual) independence of random variables is a very strong property. If a collection of random variables is independent, then any subcollection is also independent. New random variables formed from disjoint subcollections are independent. For a simple example, suppose that X , Y , and Z are independent real-valued random variables. Then

1. $\sin(X)$, $\cos(Y)$, and e^Z are independent.
2. (X, Y) and Z are independent.
3. $X^2 + Y^2$ and $\arctan(Z)$ are independent.
4. X and Z are independent.
5. Y and Z are independent.

In particular, note that statement 2 in the list above is much stronger than the conjunction of statements 4 and 5. Restated, if X and Z are dependent, then (X, Y) and Z are also dependent.

Examples and Applications

Dice

Recall that a *standard die* is an ordinary six-sided die, with faces numbered from 1 to 6. The answers in the next couple of exercises give the joint distribution in the body of a table, with the marginal distributions literally in the margins. Such tables are the reason for the term *marginal distribution*.

Suppose that two standard, fair dice are rolled and the sequence of scores (X_1, X_2) recorded. Our usual assumption is that the variables X_1 and X_2 are independent. Let $Y = X_1 + X_2$ and $Z = X_1 - X_2$ denote the sum and difference of the scores, respectively.

1. Find the probability density function of (Y, Z) .

- Find the probability density function of Y .
- Find the probability density function of Z .
- Are Y and Z independent?

Answer

Let f denote the PDF of (Y, Z) , g the PDF of Y and h the PDF of Z . The PDFs are given in the following table. Random variables Y and Z are dependent

$f(y, z)$	$y = 2$	3	4	5	6	7	8	9	0	11	12	$h(z)$
$z = -5$	0	0	0	0	0	$\frac{1}{36}$	0	0	0	0	0	$\frac{1}{36}$
-4	0	0	0	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	0	0	0	$\frac{2}{36}$
-3	0	0	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	0	0	$\frac{3}{36}$
-2	0	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	0	$\frac{4}{36}$
-1	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{5}{36}$
0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	$\frac{6}{36}$
1	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{5}{36}$
2	0	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	0	$\frac{4}{36}$
3	0	0	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	0	0	$\frac{3}{36}$
4	0	0	0	0	$\frac{1}{36}$	0	$\frac{1}{36}$	0	0	0	0	$\frac{2}{36}$
5	0	0	0	0	0	$\frac{1}{36}$	0	0	0	0	0	$\frac{1}{36}$
$g(y)$	$\frac{1}{36}$	$\frac{2}{36}$	$\frac{3}{36}$	$\frac{4}{36}$	$\frac{5}{36}$	$\frac{6}{36}$	$\frac{5}{36}$	$\frac{4}{36}$	$\frac{3}{36}$	$\frac{2}{36}$	$\frac{1}{36}$	1

Suppose that two standard, fair dice are rolled and the sequence of scores (X_1, X_2) recorded. Let $U = \min\{X_1, X_2\}$ and $V = \max\{X_1, X_2\}$ denote the minimum and maximum scores, respectively.

- Find the probability density function of (U, V) .
- Find the probability density function of U .
- Find the probability density function of V .
- Are U and V independent?

Answer

Let f denote the PDF of (U, V) , g the PDF of U , and h the PDF of V . The PDFs are given in the following table. Random variables U and V are dependent.

$f(u, v)$	$u = 1$	2	3	4	5	6	$h(v)$
$v = 1$	$\frac{1}{36}$	0	0	0	0	0	$\frac{1}{36}$
2	$\frac{2}{36}$	$\frac{1}{36}$	0	0	0	0	$\frac{3}{36}$
3	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{1}{36}$	0	0	0	$\frac{5}{36}$
4	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{1}{36}$	0	0	$\frac{7}{36}$
5	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{1}{36}$	0	$\frac{9}{36}$
6	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{2}{36}$	$\frac{1}{36}$	$\frac{11}{36}$
$g(u)$	$\frac{11}{36}$	$\frac{9}{36}$	$\frac{7}{36}$	$\frac{5}{36}$	$\frac{3}{36}$	$\frac{1}{36}$	1

The previous two exercises show clearly how little information is given with the marginal distributions compared to the joint distribution. With the marginal PDFs alone, you could not even determine the support set of the joint distribution, let alone the values of the joint PDF.

Simple Continuous Distributions

Suppose that (X, Y) has probability density function f given by $f(x, y) = x + y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = x + \frac{1}{2}$ for $0 \leq x \leq 1$
2. Y has PDF h given by $h(y) = y + \frac{1}{2}$ for $0 \leq y \leq 1$
3. X and Y are dependent.

Suppose that (X, Y) has probability density function f given by $f(x, y) = 2(x + y)$ for $0 \leq x \leq y \leq 1$.

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = (1 + 3x)(1 - x)$ for $0 \leq x \leq 1$.
2. Y has PDF h given by $h(y) = 3y^2$ for $0 \leq y \leq 1$.
3. X and Y are dependent.

Suppose that (X, Y) has probability density function f given by $f(x, y) = 6x^2y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = 3x^2$ for $0 \leq x \leq 1$.
2. Y has PDF h given by $h(y) = 2y$ for $0 \leq y \leq 1$.
3. X and Y are independent.

The last exercise is a good illustration of the [factoring theorem](#). Without any work at all, we can tell that the PDF of X is proportional to $x \mapsto x^2$ on the interval $[0, 1]$, the PDF of Y is proportional to $y \mapsto y$ on the interval $[0, 1]$, and that X and Y are independent. The only thing unclear is how the constant 6 factors.

Suppose that (X, Y) has probability density function f given by $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$.

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = \frac{15}{2}x^2(1 - x^2)$ for $0 \leq x \leq 1$.
2. Y has PDF h given by $h(y) = 5y^4$ for $0 \leq y \leq 1$.
3. X and Y are dependent.

Note that in the last exercise, the [factoring theorem](#) does not apply. Random variables X and Y each take values in $[0, 1]$, but the joint PDF factors only on part of $[0, 1]^2$.

Suppose that (X, Y, Z) has probability density function f given $f(x, y, z) = 2(x + y)z$ for $0 \leq x \leq 1, 0 \leq y \leq 1, 0 \leq z \leq 1$.

1. Find the probability density function of each pair of variables.
2. Find the probability density function of each variable.
3. Determine the dependency relationships between the variables.

Proof

1. (X, Y) has PDF $f_{1,2}$ given by $f_{1,2}(x, y) = x + y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.
2. (X, Z) has PDF $f_{1,3}$ given by $f_{1,3}(x, z) = 2z(x + \frac{1}{2})$ for $0 \leq x \leq 1, 0 \leq z \leq 1$.
3. (Y, Z) has PDF $f_{2,3}$ given by $f_{2,3}(y, z) = 2z(y + \frac{1}{2})$ for $0 \leq y \leq 1, 0 \leq z \leq 1$.
4. X has PDF f_1 given by $f_1(x) = x + \frac{1}{2}$ for $0 \leq x \leq 1$.
5. Y has PDF f_2 given by $f_2(y) = y + \frac{1}{2}$ for $0 \leq y \leq 1$.
6. Z has PDF f_3 given by $f_3(z) = 2z$ for $0 \leq z \leq 1$.
7. Z and (X, Y) are independent; X and Y are dependent.

Suppose that (X, Y) has probability density function f given by $f(x, y) = 2e^{-x}e^{-y}$ for $0 \leq x \leq y < \infty$.

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = 2e^{-2x}$ for $0 \leq x < \infty$.
2. Y has PDF h given by $h(y) = 2(e^{-y} - e^{-2y})$ for $0 \leq y < \infty$.
3. X and Y are dependent.

In the previous exercise, X has an *exponential distribution* with rate parameter 2. Recall that exponential distributions are widely used to model random times, particularly in the context of the Poisson model.

Suppose that X and Y are independent, and that X has probability density function g given by $g(x) = 6x(1-x)$ for $0 \leq x \leq 1$, and that Y has probability density function h given by $h(y) = 12y^2(1-y)$ for $0 \leq y \leq 1$.

1. Find the probability density function of (X, Y) .
2. Find $\mathbb{P}(X + Y \leq 1)$.

Answer

1. (X, Y) has PDF f given by $f(x, y) = 72x(1-x)y^2(1-y)$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.
2. $\mathbb{P}(X + Y \leq 1) = \frac{13}{35}$

In the previous exercise, X and Y have *beta distributions*, which are widely used to model random probabilities and proportions. Beta distributions are studied in more detail in the chapter on Special Distributions.

Suppose that Θ and Φ are independent random angles, with common probability density function g given by $g(t) = \sin(t)$ for $0 \leq t \leq \frac{\pi}{2}$.

1. Find the probability density function of (Θ, Φ) .
2. Find $\mathbb{P}(\Theta \leq \Phi)$.

Answer

1. (Θ, Φ) has PDF f given by $f(\theta, \phi) = \sin(\theta)\sin(\phi)$ for $0 \leq \theta \leq \frac{\pi}{2}, 0 \leq \phi \leq \frac{\pi}{2}$.
2. $\mathbb{P}(\Theta \leq \Phi) = \frac{1}{2}$

The common distribution of X and Y in the previous exercise governs a random angle in Bertrand's problem.

Suppose that X and Y are independent, and that X has probability density function g given by $g(x) = \frac{2}{x^3}$ for $1 \leq x < \infty$, and that Y has probability density function h given by $h(y) = \frac{3}{y^4}$ for $1 \leq y < \infty$.

1. Find the probability density function of (X, Y) .
2. Find $\mathbb{P}(X \leq Y)$.

Answer

1. (X, Y) has PDF f given by $f(x, y) = \frac{6}{x^3y^4}$ for $1 \leq x < \infty, 1 \leq y < \infty$.
2. $\mathbb{P}(X \leq Y) = \frac{2}{5}$

Both X and Y in the previous exercise have *Pareto distributions*, named for Vilfredo Pareto. Recall that Pareto distributions are used to model certain economic variables and are studied in more detail in the chapter on Special Distributions.

Suppose that (X, Y) has probability density function g given by $g(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$, and that Z has probability density function h given by $h(z) = 4z^3$ for $0 \leq z \leq 1$, and that (X, Y) and Z are independent.

1. Find the probability density function of (X, Y, Z) .
2. Find the probability density function of (X, Z) .
3. Find the probability density function of (Y, Z) .
4. Find $\mathbb{P}(Z \leq XY)$.

Answer

1. (X, Y, Z) has PDF f given by $f(x, y, z) = 60x^2yz^3$ for $0 \leq x \leq y \leq 1, 0 \leq z \leq 1$.
2. (X, Z) has PDF $f_{1,3}$ given by $f_{1,3}(x, z) = 30x^2(1-x^2)z^3$ for $0 \leq x \leq 1, 0 \leq z \leq 1$.
3. (Y, Z) has PDF $f_{2,3}$ given by $f_{2,3}(y, z) = 20y^4z^3$ for $0 \leq y \leq 1, 0 \leq z \leq 1$.
4. $\mathbb{P}(Z \leq XY) = \frac{15}{92}$

Multivariate Uniform Distributions

Multivariate uniform distributions give a geometric interpretation of some of the concepts in this section.

Recall first that for $n \in \mathbb{N}_+$, the standard measure on \mathbb{R}^n is

$$\lambda_n(A) = \int_A 1 dx, \quad A \subseteq \mathbb{R}^n \quad (3.4.23)$$

In particular, $\lambda_1(A)$ is the length of $A \subseteq \mathbb{R}$, $\lambda_2(A)$ is the area of $A \subseteq \mathbb{R}^2$, and $\lambda_3(A)$ is the volume of $A \subseteq \mathbb{R}^3$.

Details

Technically λ_n is Lebesgue measure on the measurable subsets of \mathbb{R}^n . The integral representation is valid for the types of sets that occur in applications. In the discussion below, all subsets are assumed to be measurable.

Suppose now that X takes values in \mathbb{R}^j , Y takes values in \mathbb{R}^k , and that (X, Y) is *uniformly distributed* on a set $R \subseteq \mathbb{R}^{j+k}$. So $0 < \lambda_{j+k}(R) < \infty$ and the joint probability density function f of (X, Y) is given by $f(x, y) = 1/\lambda_{j+k}(R)$ for $(x, y) \in R$. Recall that uniform distributions *always* have constant density functions. Now let S and T be the projections of R onto \mathbb{R}^j and \mathbb{R}^k respectively, defined as follows:

$$S = \{x \in \mathbb{R}^j : (x, y) \in R \text{ for some } y \in \mathbb{R}^k\}, \quad T = \{y \in \mathbb{R}^k : (x, y) \in R \text{ for some } x \in \mathbb{R}^j\} \quad (3.4.24)$$

Note that $R \subseteq S \times T$. Next we denote the cross sections at $x \in S$ and at $y \in T$, respectively by

$$T_x = \{t \in T : (x, t) \in R\}, \quad S_y = \{s \in S : (s, y) \in R\} \quad (3.4.25)$$




Figure 3.4.1: The projections S and T , and the cross sections at x and y

X takes values in S and Y takes values in T . The probability density functions g and h of X and Y are proportional to the cross-sectional measures:

1. $g(x) = \lambda_k(T_x) / \lambda_{j+k}(R)$ for $x \in S$
2. $h(y) = \lambda_j(S_y) / \lambda_{j+k}(R)$ for $y \in T$

Proof

From our general theory, X has PDF g given by

$$g(x) = \int_{T_x} f(x, y) dy = \int_{T_x} \frac{1}{\lambda_{j+k}(R)} dy = \frac{\lambda_k(T_x)}{\lambda_{j+k}(R)}, \quad x \in S \quad (3.4.26)$$

Technically, it's possible that $\lambda_k(T_x) = \infty$ for some $x \in S$, but the set of such x has measure 0. That is, $\lambda_j\{x \in S : \lambda_k(T_x) = \infty\} = 0$. The result for Y is analogous.

In particular, note from previous theorem that X and Y are neither independent nor uniformly distributed in general. However, these properties *do* hold if R is a Cartesian product set.

Suppose that $R = S \times T$.

1. X is uniformly distributed on S .
2. Y is uniformly distributed on T .
3. X and Y are independent.

Proof

In this case, $T_x = T$ and $S_y = S$ for every $x \in S$ and $y \in T$. Also, $\lambda_{j+k}(R) = \lambda_j(S)\lambda_k(T)$, so for $x \in S$ and $y \in T$, $f(x, y) = 1/[\lambda_j(S)\lambda_k(T)]$, $g(x) = 1/\lambda_j(S)$, $h(y) = 1/\lambda_k(T)$.

In each of the following cases, find the joint and marginal probability density functions, and determine if X and Y are independent.

1. (X, Y) is uniformly distributed on the square $R = [-6, 6]^2$.
2. (X, Y) is uniformly distributed on the triangle $R = \{(x, y) : -6 \leq y \leq x \leq 6\}$.

3. (X, Y) is uniformly distributed on the circle $R = \{(x, y) : x^2 + y^2 \leq 36\}$.

Answer

In the following, f is the PDF of (X, Y) , g the PDF of X , and h the PDF of Y .

1.
 - $f(x, y) = \frac{1}{144}$ for $-6 \leq x \leq 6, -6 \leq y \leq 6$
 - $g(x) = \frac{1}{12}$ for $-6 \leq x \leq 6$
 - $h(y) = \frac{1}{12}$ for $-6 \leq y \leq 6$
 - X and Y are independent.
2.
 - $f(x, y) = \frac{1}{72}$ for $-6 \leq y \leq x \leq 6$
 - $g(x) = \frac{1}{72}(x+6)$ for $-6 \leq x \leq 6$
 - $h(y) = \frac{1}{72}(y+6)$ for $-6 \leq y \leq 6$
 - X and Y are dependent.
3.
 - $f(x, y) = \frac{1}{36\pi}$ for $x^2 + y^2 \leq 36$
 - $g(x) = \frac{1}{18\pi}\sqrt{36-x^2}$ for $-6 \leq x \leq 6$
 - $h(y) = \frac{1}{18\pi}\sqrt{36-y^2}$ for $-6 \leq y \leq 6$
 - X and Y are dependent.

In the bivariate uniform experiment, run the simulation 1000 times for each of the following cases. Watch the points in the scatter plot and the graphs of the marginal distributions. Interpret what you see in the context of the discussion above.

1. square
2. triangle
3. circle

Suppose that (X, Y, Z) is uniformly distributed on the cube $[0, 1]^3$.

1. Give the joint probability density function of (X, Y, Z) .
2. Find the probability density function of each pair of variables.
3. Find the probability density function of each variable
4. Determine the dependency relationships between the variables.

Answer

1. (X, Y, Z) has PDF f given by $f(x, y, z) = 1$ for $0 \leq x \leq 1, 0 \leq y \leq 1, 0 \leq z \leq 1$ (the uniform distribution on $[0, 1]^3$)
2. (X, Y) , (X, Z) , and (Y, Z) have common PDF g given by $g(u, v) = 1$ for $0 \leq u \leq 1, 0 \leq v \leq 1$ (the uniform distribution on $[0, 1]^2$)
3. X, Y , and Z have common PDF h given by $h(u) = 1$ for $0 \leq u \leq 1$ (the uniform distribution on $[0, 1]$)
4. X, Y, Z are independent.

Suppose that (X, Y, Z) is uniformly distributed on $\{(x, y, z) : 0 \leq x \leq y \leq z \leq 1\}$.

1. Give the joint density function of (X, Y, Z) .
2. Find the probability density function of each pair of variables.
3. Find the probability density function of each variable
4. Determine the dependency relationships between the variables.

Answer

1. (X, Y, Z) has PDF f given by $f(x, y, z) = 6$ for $0 \leq x \leq y \leq z \leq 1$
2.
 - (X, Y) has PDF $f_{1,2}$ given by $f_{1,2}(x, y) = 6(1-y)$ for $0 \leq x \leq y \leq 1$
 - (X, Z) has PDF $f_{1,3}$ given by $f_{1,3}(x, z) = 6(z-x)$ for $0 \leq x \leq z \leq 1$
 - (Y, Z) has PDF $f_{2,3}$ given by $f_{2,3}(y, z) = 6y$ for $0 \leq y \leq z \leq 1$
3.
 - X has PDF f_1 given by $f_1(x) = 3(1-x)^2$ for $0 \leq x \leq 1$
 - Y has PDF f_2 given by $f_2(y) = 6y(1-y)$ for $0 \leq y \leq 1$
 - Z has PDF f_3 given by $f_3(z) = 3z^2$ for $0 \leq z \leq 1$
4. Each pair of variables is dependent.

The Rejection Method

The following result shows how an arbitrary continuous distribution can be obtained from a uniform distribution. This result is useful for *simulating* certain continuous distributions, as we will see. To set up the basic notation, suppose that g is a probability density function for a continuous distribution on $S \subseteq \mathbb{R}^n$. Let

$$R = \{(x, y) : x \in S \text{ and } 0 \leq y \leq g(x)\} \subseteq \mathbb{R}^{n+1} \quad (3.4.27)$$

If (X, Y) is uniformly distributed on R , then X has probability density function g .

Proof

Note that since g is a probability density function on S .

$$\lambda_{n+1}(R) = \int_R 1 \, d(x, y) = \int_S \int_0^{g(x)} 1 \, dy \, dx = \int_S g(x) \, dx = 1 \quad (3.4.28)$$

Hence the probability density function f of (X, Y) is given by $f(x, y) = 1$ for $(x, y) \in R$. Thus, the probability density function of X is $x \mapsto \int_0^{g(x)} 1 \, dy = g(x)$ for $x \in S$.

A picture in the case $n = 1$ is given below:




Figure 3.4.2: If (X, Y) is uniformly distributed on R , then X has density function g .

The next result gives the *rejection method* for simulating a random variable with the probability density function g .

Suppose now that $R \subseteq T$ where $T \subseteq \mathbb{R}_{n+1}$ with $\lambda_{n+1}(T) < \infty$ and that $((X_1, Y_1), (X_2, Y_2), \dots)$ is a sequence of independent random variables with $X_k \in \mathbb{R}^n$, $Y_k \in \mathbb{R}$, and (X_k, Y_k) uniformly distributed on T for each $k \in \mathbb{N}_+$. Let

$$N = \min \{k \in \mathbb{N}_+ : (X_k, Y_k) \in R\} = \min \{k \in \mathbb{N}_+ : X_k \in S, 0 \leq Y_k \leq g(X_k)\} \quad (3.4.29)$$

1. N has the geometric distribution on \mathbb{N}_+ with success parameter $p = 1/\lambda_{n+1}(T)$.
2. (X_N, Y_N) is uniformly distributed on R .
3. X_N has probability density function g .

Proof




Figure 3.4.3: With a sequence of independent points, uniformly distributed on T , the x coordinate of the first point to land in R has probability density function g .

The point of the theorem is that if we can simulate a sequence of independent variables that are uniformly distributed on T , then we can simulate a random variable with the given probability density function g . Suppose in particular that R is bounded as a subset of \mathbb{R}^{n+1} , which would mean that the domain S is bounded as a subset of \mathbb{R}^n and that the probability density function g is bounded on S . In this case, we can find T that is the Cartesian product of $n + 1$ bounded intervals with $R \subseteq T$. It turns out to be very easy to simulate a sequence of independent variables, each uniformly distributed on such a product set, so the rejection method always works in this case. As you might guess, the rejection method works best if the size of T , namely $\lambda_{n+1}(T)$, is small, so that the success parameter p is large.

The rejection method app simulates a number of continuous distributions via the rejection method. For each of the following distributions, vary the parameters and note the shape and location of the probability density function. Then run the experiment 1000 times and observe the results.

1. The beta distribution
2. The semicircle distribution
3. The triangle distribution
4. The U-power distribution

The Multivariate Hypergeometric Distribution

Suppose that a population consists of m objects, and that each object is one of four types. There are a type 1 objects, b type 2 objects, c type 3 objects and $m - a - b - c$ type 0 objects. We sample n objects from the population at random, and without replacement. The parameters m , a , b , c , and n are nonnegative integers with $a + b + c \leq m$ and $n \leq m$. Denote the number of type 1, 2, and 3 objects in the sample by X , Y , and Z ,

respectively. Hence, the number of type 0 objects in the sample is $n - X - Y - Z$. In the problems below, the variables x , y , and z take values in \mathbb{N} .

(X, Y, Z) has a (multivariate) hypergeometric distribution with probability density function f given by

$$f(x, y, z) = \frac{\binom{a}{x} \binom{b}{y} \binom{c}{z} \binom{m-a-b-c}{n-x-y-z}}{\binom{m}{n}}, \quad x + y + z \leq n \quad (3.4.30)$$

Proof

From the basic theory of combinatorics, the numerator is the number of ways to select an unordered sample of size n from the population with x objects of type 1, y objects of type 2, z objects of type 3, and $n - x - y - z$ objects of type 0. The denominator is the total number of ways to select the unordered sample.

(X, Y) also has a (multivariate) hypergeometric distribution, with the probability density function g given by

$$g(x, y) = \frac{\binom{a}{x} \binom{b}{y} \binom{m-a-b}{n-x-y}}{\binom{m}{n}}, \quad x + y \leq n \quad (3.4.31)$$

Proof

This result could be obtained by summing the [joint PDF](#) over z for fixed (x, y) . However, there is a much nicer combinatorial argument. Note that we are selecting a random sample of size n from a population of m objects, with a objects of type 1, b objects of type 2, and $m - a - b$ objects of other types.

X has an ordinary hypergeometric distribution, with probability density function h given by

$$h(x) = \frac{\binom{a}{x} \binom{m-a}{n-x}}{\binom{m}{n}}, \quad x \leq n \quad (3.4.32)$$

Proof

Again, the result could be obtained by summing the [joint PDF](#) for (X, Y, Z) over (y, z) for fixed x , or by summing the [joint PDF](#) for (X, Y) over y for fixed x . But as before, there is a much more elegant combinatorial argument. Note that we are selecting a random sample of size n from a population of size m objects, with a objects of type 1 and $m - a$ objects of other types.

These results generalize in a straightforward way to a population with any number of types. In brief, if a random vector has a hypergeometric distribution, then any sub-vector also has a hypergeometric distribution. In other words, all of the marginal distributions of a hypergeometric distribution are themselves hypergeometric. Note however, that it's not a good idea to memorize the formulas above explicitly. It's better to just note the patterns and recall the combinatorial meaning of the binomial coefficient. The hypergeometric distribution and the multivariate hypergeometric distribution are studied in more detail in the chapter on Finite Sampling Models.

Suppose that a population of voters consists of 50 democrats, 40 republicans, and 30 independents. A sample of 10 voters is chosen at random from the population (without replacement, of course). Let X denote the number of democrats in the sample and Y the number of republicans in the sample. Find the probability density function of each of the following:

1. (X, Y)
2. X
3. Y

Answer

In the formulas for the PDFs below, the variables x and y are nonnegative integers.

1. (X, Y) has PDF f given by $f(x, y) = \frac{1}{\binom{120}{10}} \binom{50}{x} \binom{40}{y} \binom{30}{10-x-y}$ for $x + y \leq 10$
2. X has PDF g given by $g(x) = \frac{1}{\binom{120}{10}} \binom{50}{x} \binom{70}{10-x}$ for $x \leq 10$
3. Y has PDF h given by $h(y) = \frac{1}{\binom{120}{10}} \binom{40}{y} \binom{80}{10-y}$ for $y \leq 10$

Suppose that the Math Club at Enormous State University (ESU) has 50 freshmen, 40 sophomores, 30 juniors and 20 seniors. A sample of 10 club members is chosen at random to serve on the π -day committee. Let X denote the number freshmen on the committee, Y the number of sophomores, and Z the number of juniors.

1. Find the probability density function of (X, Y, Z)
2. Find the probability density function of each pair of variables.

3. Find the probability density function of each individual variable.

Answer

In the formulas for the PDFs below, the variables x , y , and z are nonnegative integers.

1. (X, Y, Z) has PDF f given by $f(x, y, z) = \frac{1}{\binom{140}{10}} \binom{50}{x} \binom{40}{y} \binom{30}{z} \binom{20}{10-x-y-z}$ for $x + y + z \leq 10$.
2.
 - (X, Y) has PDF $f_{1,2}$ given by $f_{1,2}(x, y) = \frac{1}{\binom{140}{10}} \binom{50}{x} \binom{40}{y} \binom{50}{10-x-y}$ for $x + y \leq 10$.
 - (X, Z) has PDF $f_{1,3}$ given by $f_{1,3}(y, z) = \frac{1}{\binom{140}{10}} \binom{50}{x} \binom{30}{z} \binom{60}{10-x-z}$ for $x + z \leq 10$.
 - (Y, Z) has PDF $f_{2,3}$ given by $f_{2,3}(y, z) = \frac{1}{\binom{140}{10}} \binom{40}{y} \binom{30}{z} \binom{70}{10-y-z}$ for $y + z \leq 10$.
3.
 - X has PDF f_1 given by $f_1(x) = \frac{1}{\binom{120}{10}} \binom{50}{x} \binom{90}{10-x}$ for $x \leq 10$.
 - Y has PDF f_2 given by $f_2(y) = \frac{1}{\binom{120}{10}} \binom{40}{y} \binom{100}{10-y}$ for $y \leq 10$.
 - Z has PDF f_3 given by $f_3(z) = \frac{1}{\binom{120}{10}} \binom{30}{z} \binom{110}{10-z}$ for $z \leq 10$.

Multinomial Trials

Suppose that we have a sequence of n independent trials, each with 4 possible outcomes. On each trial, outcome 1 occurs with probability p , outcome 2 with probability q , outcome 3 with probability r , and outcome 0 occurs with probability $1 - p - q - r$. The parameters p , q , and r are nonnegative numbers with $p + q + r \leq 1$, and $n \in \mathbb{N}_+$. Denote the number of times that outcome 1, outcome 2, and outcome 3 occurred in the n trials by X , Y , and Z respectively. Of course, the number of times that outcome 0 occurs is $n - X - Y - Z$. In the problems below, the variables x , y , and z take values in \mathbb{N} .

(X, Y, Z) has a *multinomial distribution* with probability density function f given by

$$f(x, y, z) = \binom{n}{x, y, z} p^x q^y r^z (1 - p - q - r)^{n-x-y-z}, \quad x + y + z \leq n \quad (3.4.33)$$

Proof

The multinomial coefficient is the number of sequences of length n with 1 occurring x times, 2 occurring y times, 3 occurring z times, and 0 occurring $n - x - y - z$ times. The result then follows by independence.

(X, Y) also has a multinomial distribution with the probability density function g given by

$$g(x, y) = \binom{n}{x, y} p^x q^y (1 - p - q)^{n-x-y}, \quad x + y \leq n \quad (3.4.34)$$

Proof

This result could be obtained from the [joint PDF](#) above, by summing over z for fixed (x, y) . However there is a much better direct argument. Note that we have n independent trials, and on each trial, outcome 1 occurs with probability p , outcome 2 with probability q , and some other outcome with probability $1 - p - q$.

X has a *binomial distribution*, with the probability density function h given by

$$h(x) = \binom{n}{x} p^x (1 - p)^{n-x}, \quad x \leq n \quad (3.4.35)$$

Proof

Again, the result could be obtained by summing the [joint PDF](#) for (X, Y, Z) over (y, z) for fixed x or by summing the [joint PDF](#) for (X, Y) over y for fixed x . But as before, there is a much better direct argument. Note that we have n independent trials, and on each trial, outcome 1 occurs with probability p and some other outcome with probability $1 - p$.

These results generalize in a completely straightforward way to multinomial trials with any number of trial outcomes. In brief, if a random vector has a multinomial distribution, then any sub-vector also has a multinomial distribution. In other terms, all of the marginal distributions of a multinomial distribution are themselves multinomial. The binomial distribution and the multinomial distribution are studied in more detail in the chapter on Bernoulli Trials.

Suppose that a system consists of 10 components that operate independently. Each component is *working* with probability $\frac{1}{2}$, *idle* with probability $\frac{1}{3}$, or *failed* with probability $\frac{1}{6}$. Let X denote the number of working components and Y the number of idle components. Give the probability density function of each of the following:

1. (X, Y)
2. X
3. Y

Answer

In the formulas below, the variables x and y are nonnegative integers.

1. (X, Y) has PDF f given by $f(x, y) = \binom{10}{x, y} \left(\frac{1}{2}\right)^x \left(\frac{1}{3}\right)^y \left(\frac{1}{6}\right)^{10-x-y}$ for $x + y \leq 10$.
2. X has PDF g given by $g(x) = \binom{10}{x} \left(\frac{1}{2}\right)^{10}$ for $x \leq 10$.
3. Y has PDF h given by $h(y) = \binom{10}{y} \left(\frac{1}{3}\right)^y \left(\frac{2}{3}\right)^{10-y}$ for $y \leq 10$.

Suppose that in a crooked, four-sided die, face i occurs with probability $\frac{i}{10}$ for $i \in \{1, 2, 3, 4\}$. The die is thrown 12 times; let X denote the number of times that score 1 occurs, Y the number of times that score 2 occurs, and Z the number of times that score 3 occurs.

1. Find the probability density function of (X, Y, Z)
2. Find the probability density function of each pair of variables.
3. Find the probability density function of each individual variable.

Answer

In the formulas for the PDFs below, the variables x , y and z are nonnegative integers.

1. (X, Y, Z) has PDF f given by $f(x, y, z) = \binom{12}{x, y, z} \left(\frac{1}{10}\right)^x \left(\frac{2}{10}\right)^y \left(\frac{3}{10}\right)^z \left(\frac{4}{10}\right)^{12-x-y-z}$, $x + y + z \leq 12$
2. (X, Y) has PDF $f_{1,2}$ given by $f_{1,2}(x, y) = \binom{12}{x, y} \left(\frac{1}{10}\right)^x \left(\frac{2}{10}\right)^y \left(\frac{7}{10}\right)^{12-x-y}$ for $x + y \leq 12$.
 - (X, Z) has PDF $f_{1,3}$ given by $f_{1,3}(x, z) = \binom{12}{x, z} \left(\frac{1}{10}\right)^x \left(\frac{3}{10}\right)^z \left(\frac{6}{10}\right)^{12-x-z}$ for $x + z \leq 12$.
 - (Y, Z) has PDF $f_{2,3}$ given by $f_{2,3}(y, z) = \binom{12}{y, z} \left(\frac{2}{10}\right)^y \left(\frac{3}{10}\right)^z \left(\frac{5}{10}\right)^{12-y-z}$ for $y + z \leq 12$.
3. X has PDF f_1 given by $f_1(x) = \binom{12}{x} \left(\frac{1}{10}\right)^x \left(\frac{9}{10}\right)^{12-x}$ for $x \leq 12$.
 - Y has PDF f_2 given by $f_2(y) = \binom{12}{y} \left(\frac{2}{10}\right)^y \left(\frac{8}{10}\right)^{12-y}$ for $y \leq 12$.
 - Z has PDF f_3 given by $f_3(z) = \binom{12}{z} \left(\frac{3}{10}\right)^z \left(\frac{7}{10}\right)^{12-z}$ for $z \leq 12$.

Bivariate Normal Distributions

Suppose that (X, Y) has probability the density function f given below:

$$f(x, y) = \frac{1}{12\pi} \exp\left[-\left(\frac{x^2}{8} + \frac{y^2}{18}\right)\right], \quad (x, y) \in \mathbb{R}^2 \quad (3.4.36)$$

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = \frac{1}{2\sqrt{2\pi}} e^{-x^2/8}$ for $x \in \mathbb{R}$.
2. Y has PDF h given by $h(y) = \frac{1}{3\sqrt{2\pi}} e^{-y^2/18}$ for $y \in \mathbb{R}$.
3. X and Y are independent.

Suppose that (X, Y) has probability density function f given below:

$$f(x, y) = \frac{1}{\sqrt{3}\pi} \exp\left[-\frac{2}{3}(x^2 - xy + y^2)\right], \quad (x, y) \in \mathbb{R}^2 \quad (3.4.37)$$

1. Find the density function of X .
2. Find the density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ for $x \in \mathbb{R}$.
2. Y has PDF h given by $h(y) = \frac{1}{\sqrt{2\pi}} e^{-y^2/2}$ for $y \in \mathbb{R}$.

3. X and Y are dependent.

The joint distributions in the last two exercises are examples of *bivariate normal distributions*. Normal distributions are widely used to model physical measurements subject to small, random errors. In both exercises, the marginal distributions of X and Y also have normal distributions, and this turns out to be true in general. The multivariate normal distribution is studied in more detail in the chapter on Special Distributions.

Exponential Distributions

Recall that the *exponential distribution* has probability density function

$$f(x) = re^{-rt}, \quad x \in [0, \infty) \quad (3.4.38)$$

where $r \in (0, \infty)$ is the *rate parameter*. The exponential distribution is widely used to model random times, and is studied in more detail in the chapter on the Poisson Process.

Suppose X and Y have exponential distributions with parameters $a \in (0, \infty)$ and $b \in (0, \infty)$, respectively, and are independent. Then $\mathbb{P}(X < Y) = \frac{a}{a+b}$.

Suppose X , Y , and Z have exponential distributions with parameters $a \in (0, \infty)$, $b \in (0, \infty)$, and $c \in (0, \infty)$, respectively, and are independent. Then

1. $\mathbb{P}(X < Y < Z) = \frac{a}{a+b+c} \frac{b}{b+c}$
2. $\mathbb{P}(X < Y, X < Z) = \frac{a}{a+b+c}$

If X , Y , and Z are the lifetimes of devices that act independently, then the results in the previous two exercises give probabilities of various failure orders. Results of this type are also very important in the study of continuous-time Markov processes. We will continue this discussion in the section on transformations of random variables.

Mixed Coordinates

Suppose X takes values in the finite set $\{1, 2, 3\}$, Y takes values in the interval $[0, 3]$, and that (X, Y) has probability density function f given by

$$f(x, y) = \begin{cases} \frac{1}{3}, & x = 1, 0 \leq y \leq 1 \\ \frac{1}{6}, & x = 2, 0 \leq y \leq 2 \\ \frac{1}{9}, & x = 3, 0 \leq y \leq 3 \end{cases} \quad (3.4.39)$$

1. Find the probability density function of X .
2. Find the probability density function of Y .
3. Are X and Y independent?

Answer

1. X has PDF g given by $g(x) = \frac{1}{3}$ for $x \in \{1, 2, 3\}$ (the uniform distribution on $\{1, 2, 3\}$).
2. Y has PDF h given by $h(y) = \begin{cases} \frac{11}{18}, & 0 < y < 1 \\ \frac{5}{18}, & 1 < y < 2 \\ \frac{2}{18}, & 2 < y < 3 \end{cases}$.
3. X and Y are dependent.

Suppose that P takes values in the interval $[0, 1]$, X takes values in the finite set $\{0, 1, 2, 3\}$ and that (P, X) has probability density function f given by

$$f(p, x) = 6 \binom{3}{x} p^{x+1} (1-p)^{4-x}, \quad (p, x) \in [0, 1] \times \{0, 1, 2, 3\} \quad (3.4.40)$$

1. Find the probability density function of P .
2. Find the probability density function of X .
3. Are P and X independent?

Answer

1. P has PDF g given by $g(p) = 6p(1-p)$ for $0 \leq p \leq 1$.
2. X has PDF h given by $h(0) = h(3) = \frac{1}{5}$, $h(1) = \frac{3}{10}$.
3. P and X are dependent.

As we will see in the section on conditional distributions, the distribution in the last exercise models the following experiment: a random probability P is selected, and then a coin with this probability of heads is tossed 3 times; X is the number of heads. Note that P has a beta distribution.

Random Samples

Recall that the *Bernoulli distribution* with parameter $p \in [0, 1]$ has probability density function g given by $g(x) = p^x(1-p)^{1-x}$ for $x \in \{0, 1\}$. Let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be a random sample of size $n \in \mathbb{N}_+$ from the distribution. Give the probability density function of \mathbf{X} in simplified form.

Answer

\mathbf{X} has PDF f given by $f(x_1, x_2, \dots, x_n) = p^y(1-p)^{n-y}$ for $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, where $y = x_1 + x_2 + \dots + x_n$.

The Bernoulli distribution is named for Jacob Bernoulli, and governs an indicator random variable. Hence if \mathbf{X} is a random sample of size n from the distribution then \mathbf{X} is a sequence of n *Bernoulli trials*. A separate chapter studies Bernoulli trials in more detail.

Recall that the *geometric distribution* on \mathbb{N}_+ with parameter $p \in (0, 1)$ has probability density function g given by $g(x) = p(1-p)^{x-1}$ for $x \in \mathbb{N}_+$. Let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be a random sample of size $n \in \mathbb{N}_+$ from the distribution. Give the probability density function of \mathbf{X} in simplified form.

Answer

\mathbf{X} has pdf f given by $f(x_1, x_2, \dots, x_n) = p^n(1-p)^{y-n}$ for $(x_1, x_2, \dots, x_n) \in \mathbb{N}_+^n$, where $y = x_1 + x_2 + \dots + x_n$.

The geometric distribution governs the trial number of the first success in a sequence of Bernoulli trials. Hence the variables in the random sample can be interpreted as the number of trials between successive successes.

Recall that the *Poisson distribution* with parameter $a \in (0, \infty)$ has probability density function g given by $g(x) = e^{-a} \frac{a^x}{x!}$ for $x \in \mathbb{N}$. Let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be a random sample of size $n \in \mathbb{N}_+$ from the distribution. Give the probability density function of \mathbf{X} in simplified form.

Answer

\mathbf{X} has PDF f given by $f(x_1, x_2, \dots, x_n) = \frac{1}{x_1!x_2!\dots x_n!} e^{-na} a^y$ for $(x_1, x_2, \dots, x_n) \in \mathbb{N}^n$, where $y = x_1 + x_2 + \dots + x_n$.

The Poisson distribution is named for Simeon Poisson, and governs the number of random points in a region of time or space under appropriate circumstances. The parameter a is proportional to the size of the region. The Poisson distribution is studied in more detail in the chapter on the Poisson process.

Recall again that the *exponential distribution* with rate parameter $r \in (0, \infty)$ has probability density function g given by $g(x) = re^{-rx}$ for $x \in (0, \infty)$. Let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be a random sample of size $n \in \mathbb{N}_+$ from the distribution. Give the probability density function of \mathbf{X} in simplified form.

Answer

\mathbf{X} has PDF f given by $f(x_1, x_2, \dots, x_n) = r^n e^{-ry}$ for $(x_1, x_2, \dots, x_n) \in [0, \infty)^n$, where $y = x_1 + x_2 + \dots + x_n$.

The exponential distribution governs failure times and other types of arrival times under appropriate circumstances. The exponential distribution is studied in more detail in the chapter on the Poisson process. The variables in the random sample can be interpreted as the times between successive arrivals in the Poisson process.

Recall that the *standard normal distribution* has probability density function ϕ given by $\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ for $z \in \mathbb{R}$. Let $\mathbf{Z} = (Z_1, Z_2, \dots, Z_n)$ be a random sample of size $n \in \mathbb{N}_+$ from the distribution. Give the probability density function of \mathbf{Z} in simplified form.

Answer

\mathbf{Z} has PDF f given by $f(z_1, z_2, \dots, z_n) = \frac{1}{(2\pi)^{n/2}} e^{-\frac{1}{2}w^2}$ for $(z_1, z_2, \dots, z_n) \in \mathbb{R}^n$, where $w^2 = z_1^2 + z_2^2 + \dots + z_n^2$.

The standard normal distribution governs physical quantities, properly scaled and centered, subject to small, random errors. The normal distribution is studied in more generality in the chapter on the Special Distributions.

Data Analysis Exercises

For the cicada data, G denotes gender and S denotes species type.

1. Find the empirical density of (G, S) .
2. Find the empirical density of G .
3. Find the empirical density of S .

4. Do you believe that S and G are independent?

Answer

The empirical joint and marginal empirical densities are given in the table below. Gender and species are probably dependent (compare the joint density with the product of the marginal densities).

$f(i, j)$	$i = 0$	1	$h(j)$
$j = 0$	$\frac{16}{104}$	$\frac{28}{104}$	$\frac{44}{104}$
1	$\frac{3}{104}$	$\frac{3}{104}$	$\frac{6}{104}$
2	$\frac{40}{104}$	$\frac{14}{104}$	$\frac{56}{104}$
$g(i)$	$\frac{59}{104}$	$\frac{45}{104}$	1

For the cicada data, let W denote body weight (in grams) and L body length (in millimeters).

1. Construct an empirical density for (W, L) .
2. Find the corresponding empirical density for W .
3. Find the corresponding empirical density for L .
4. Do you believe that W and L are independent?

Answer

The empirical joint and marginal densities, based on simple partitions of the body weight and body length ranges, are given in the table below. Body weight and body length are almost certainly dependent.

Density (W, L)	$w \in (0, 0.1]$	$(0.1, 0.2]$	$(0.2, 0.3]$	$(0.3, 0.4]$	Density L
$l \in (15, 20]$	0	0.0385	0.0192	0	0.0058
$(20, 25]$	0.1731	0.9808	0.4231	0	0.1577
$(25, 30]$	0	0.1538	0.1731	0.0192	0.0346
$(30, 35]$	0	0	0	0.0192	0.0019
Density W	0.8654	5.8654	3.0769	0.1923	

For the cicada data, let G denote gender and W body weight (in grams).

1. Construct an empirical density for (W, G) .
2. Find the empirical density for G .
3. Find the empirical density for W .
4. Do you believe that G and W are independent?

Answer

The empirical joint and marginal densities, based on a simple partition of the body weight range, are given in the table below. Body weight and gender are almost certainly dependent.

Density (W, G)	$w \in (0, 0.1]$	$(0.1, 0.2]$	$(0.2, 0.3]$	$(0.3, 0.4]$	Density G
$g = 0$	0.1923	2.5000	2.8846	0.0962	0.5673
1	0.6731	3.3654	0.1923	0.0962	0.4327
Density W	0.8654	5.8654	3.0769	0.1923	

3.5: Conditional Distributions

In this section, we study how a probability distribution changes when a given random variable has a known, specified value. So this is an essential topic that deals with how probability measures should be updated in light of new information. As usual, if you are a new student or probability, you may want to skip the technical details.

Basic Theory

Our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} is the collection of events, and \mathbb{P} is the probability measure on the underlying sample space (Ω, \mathcal{F}) .

Suppose that X is a random variable defined on the sample space (that is, defined for the experiment), taking values in a set S .

Details

Technically, the collection of events \mathcal{F} is a σ -algebra, so that the sample space (Ω, \mathcal{F}) is a measurable space. Similarly, S will have a σ -algebra of admissible subsets, so that (S, \mathcal{S}) is also a measurable space. Random variable X is measurable, so that $\{X \in A\} \in \mathcal{F}$ for every $A \in \mathcal{S}$. The *distribution* of X is the probability measure $A \mapsto \mathbb{P}(X \in A)$ for $A \in \mathcal{S}$.

The purpose of this section is to study the conditional probability measure given $X = x$ for $x \in S$. That is, if E is an event, we would like to define and study the probability of E given $X = x$, denoted $\mathbb{P}(E | X = x)$. If X has a discrete distribution, the conditioning event has positive probability, so no new concepts are involved, and the simple definition of conditional probability suffices. When X has a continuous distribution, however, the conditioning event has probability 0, so a fundamentally new approach is needed.

The Discrete Case

Suppose first that X has a discrete distribution with probability density function g . Thus S is countable and we can assume that $g(x) = \mathbb{P}(X = x) > 0$ for $x \in S$.

If E is an event and $x \in S$ then

$$\mathbb{P}(E | X = x) = \frac{\mathbb{P}(E, X = x)}{g(x)} \quad (3.5.1)$$

Proof

The meaning of *discrete distribution* is that S is countable and $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S . Technically, g is the probability density function of X with respect to counting measure $\#$ on S , the standard measure for discrete spaces. In the displayed equation above, the comma separating the events in the numerator of the fraction means *and*, and thus functions just like the intersection symbol. This result follows immediately from the definition of conditional probability:

$$\mathbb{P}(E | X = x) = \frac{\mathbb{P}(E, X = x)}{\mathbb{P}(X = x)} = \frac{\mathbb{P}(E, X = x)}{g(x)} \quad (3.5.2)$$

The next result is a special case of the law of total probability, and will be the key to the definition when X has a continuous distribution.

If E is an event then

$$\mathbb{P}(E, X \in A) = \sum_{x \in A} g(x) \mathbb{P}(E | X = x), \quad A \subseteq S \quad (3.5.3)$$

Conversely, this condition uniquely determines $\mathbb{P}(E | X = x)$.

Proof

As noted, the displayed equation is just a special case of the law of total probability. For $A \subseteq S$, the countable collection of events $\{\{X = x\} : x \in A\}$ partitions $\{X \in A\}$ so

$$\mathbb{P}(E, X \in A) = \sum_{x \in A} \mathbb{P}(E, X = x) = \sum_{x \in A} \mathbb{P}(E | X = x) \mathbb{P}(X = x) = \sum_{x \in A} \mathbb{P}(E | X = x) g(x) \quad (3.5.4)$$

Conversely, suppose that the function $Q(x, E)$, defined for $x \in S$ and $E \in \mathcal{F}$, satisfies

$$\mathbb{P}(E, X \in A) = \sum_{x \in A} g(x) Q(x, E), \quad A \subseteq S \quad (3.5.5)$$

Letting $A = \{x\}$ for $x \in S$ gives $\mathbb{P}(E, X = x) = g(x)Q(x, E)$, so $Q(x, E) = \mathbb{P}(E, X = x)/g(x) = \mathbb{P}(E | X = x)$.

The Continuous Case

Suppose now that X has a continuous distribution on $S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$, with probability density function g . We assume that $g(x) > 0$ for $x \in S$. Unlike the discrete case, we cannot use simple conditional probability to define $\mathbb{P}(E | X = x)$ for an event E and $x \in S$ because the conditioning event has probability 0. Nonetheless, the concept *should* make sense. If we actually run the experiment, X will take on some value x (even though *a priori*, this event occurs with probability 0), and surely the information that $X = x$ should in general alter the probabilities that we assign to other events. A natural approach is to use the results obtained in the discrete case as *definitions* in the continuous case.

If E is an event and $x \in S$, the *conditional probability* $\mathbb{P}(E | X = x)$ is defined by the requirement that

$$\mathbb{P}(E, X \in A) = \int_A g(x)\mathbb{P}(E | X = x) dx, \quad A \subseteq S \quad (3.5.6)$$

Details

Technically, S is a measurable subset of \mathbb{R}^n and the σ -algebra \mathcal{S} consists of the subsets of S that are also measurable as subsets of \mathbb{R}^n . The function g is also required to be measurable, and is the density function of X with respect to Lebesgue measure λ_n . Lebesgue measure is named for Henri Lebesgue and is the standard measure on \mathbb{R}^n .

We will accept the fact that $\mathbb{P}(E | X = x)$ can be defined uniquely, up to a set of measure 0, by the condition above, but we will return to this point in the section on Conditional Expectation in the chapter on Expected Value. Essentially the condition means that $\mathbb{P}(E | X = x)$ is defined so that $x \mapsto g(x)\mathbb{P}(E | X = x)$ is a density function for the finite measure $A \mapsto \mathbb{P}(E, X \in A)$.

Conditioning and Bayes' Theorem

Suppose again that X is a random variable with values in S and probability density function g , as described above. Our discussions above in the discrete and continuous cases lead to basic formulas for computing the probability of an event by *conditioning on X* .

The law of total probability. If E is an event, then $\mathbb{P}(E)$ can be computed as follows:

1. If X has a discrete distribution then

$$\mathbb{P}(E) = \sum_{x \in S} g(x)\mathbb{P}(E | X = x) \quad (3.5.7)$$

2. If X has a continuous distribution then

$$\mathbb{P}(E) = \int_S g(x)\mathbb{P}(E | X = x) dx \quad (3.5.8)$$

Proof

1. This follows from the [discrete theorem](#) with $A = S$.
2. This follows from the [fundamental definition](#) with $A = S$.

Naturally, the law of total probability is useful when $\mathbb{P}(E | X = x)$ and $g(x)$ are known for $x \in S$. Our next result is, *Bayes' Theorem*, named after Thomas Bayes.

Bayes' Theorem. Suppose that E is an event with $\mathbb{P}(E) > 0$. The conditional probability density function $x \mapsto g(x | E)$ of X given E can be computed as follows:

1. If X has a discrete distribution then

$$g(x | E) = \frac{g(x)\mathbb{P}(E | X = x)}{\sum_{s \in S} g(s)\mathbb{P}(E | X = s)}, \quad x \in S \quad (3.5.9)$$

2. If X has a continuous distribution then

$$g(x | E) = \frac{g(x)\mathbb{P}(E | X = x)}{\int_S g(s)\mathbb{P}(E | X = s) ds}, \quad x \in S \quad (3.5.10)$$

Proof

1. In the discrete case, as usual, the ordinary simple definition of conditional probability suffices. The numerator in the displayed equation is $\mathbb{P}(X = x)\mathbb{P}(E | X = x) = \mathbb{P}(E, X = x)$. The denominator is $\mathbb{P}(E)$ by part (a) of the [law of total probability](#). Hence the fraction is $\mathbb{P}(E, X = x)/\mathbb{P}(E) = \mathbb{P}(X = x | E)$.
2. In the continuous case, as usual, the argument is more subtle. We need to show that the expression in the displayed equation satisfies the defining property of a PDF for the conditional distribution. Once again, the denominator is $\mathbb{P}(E)$ by part (b) of the [law of total probability](#). If $A \subseteq S$ then using the [fundamental definition](#),

$$\int_A g(x | E) dx = \frac{1}{\mathbb{P}(E)} \int_A g(x)\mathbb{P}(E | X = x) dx = \frac{\mathbb{P}(E, X \in A)}{\mathbb{P}(E)} = \mathbb{P}(X \in A | E) \quad (3.5.11)$$

By the meaning of the term, $x \mapsto g(x | E)$ is the conditional probability density function of X given E .

In the context of Bayes' theorem, g is called the *prior probability density function* of X and $x \mapsto g(x | E)$ is the *posterior probability density function* of X given E . Note also that the conditional probability density function of X given E is proportional to the function $x \mapsto g(x)\mathbb{P}(E | X = x)$, the sum or integral of this function that occurs in the denominator is simply the normalizing constant. As with the law of total probability, Bayes' theorem is useful when $\mathbb{P}(E | X = x)$ and $g(x)$ are known for $x \in S$.

Conditional Probability Density Functions

The definitions and results above apply, of course, if E is an event defined in terms of another random variable for our experiment. Here is the setup:

Suppose that X and Y are random variables on the probability space, with values in sets S and T , respectively, so that (X, Y) is a random variable with values in $S \times T$. We assume that (X, Y) has probability density function f , as discussed in the section on Joint Distributions. Recall that X has probability density function g defined as follows:

1. If Y has a discrete distribution on the countable set T then

$$g(x) = \sum_{y \in T} f(x, y), \quad x \in S \quad (3.5.12)$$

2. If Y has a continuous distribution on $T \subseteq \mathbb{R}^k$ then

$$g(x) = \int_T f(x, y) dy, \quad x \in S \quad (3.5.13)$$

Similarly, the probability density function h of Y can be obtained by summing f over $x \in S$ if X has a discrete distribution or integrating f over S if X has a continuous distribution.

Suppose that $x \in S$ and that $g(x) > 0$. The function $y \mapsto h(y | x)$ defined below is a probability density function on T :

$$h(y | x) = \frac{f(x, y)}{g(x)}, \quad y \in T \quad (3.5.14)$$

Proof

The result is simple, since $g(x)$ is the normalizing constant for $y \mapsto h(y | x)$. Specifically, fix $x \in S$. Then $h(y | x) \geq 0$. If Y has a discrete distribution then

$$\sum_{y \in T} h(y | x) = \frac{1}{g(x)} \sum_{y \in T} f(x, y) = \frac{g(x)}{g(x)} = 1 \quad (3.5.15)$$

Similarly, if Y has a continuous distribution then

$$\int_T h(y | x) dy = \frac{1}{g(x)} \int_T f(x, y) dy = \frac{g(x)}{g(x)} = 1 \quad (3.5.16)$$

The distribution that corresponds to this probability density function is what you would expect:

For $x \in S$, the function $y \mapsto h(y | x)$ is the conditional probability density function of Y given $X = x$. That is,

1. If Y has a discrete distribution then

$$\mathbb{P}(Y \in B \mid X = x) = \sum_{y \in B} h(y \mid x), \quad B \subseteq T \quad (3.5.17)$$

2. If Y has a continuous distribution then

$$\mathbb{P}(Y \in B \mid X = x) = \int_B h(y \mid x) dy, \quad B \subseteq T \quad (3.5.18)$$

Proof

There are four cases, depending on the type of distribution of X and Y , but the computations are identical, except for sums in the discrete case and integrals in the continuous case. The main tool is the [basic theorem](#) when X has a discrete distribution and the [fundamental definition](#) when X has a continuous distribution, with the event E replaced by $\{Y \in B\}$ for $B \subseteq T$. The other main element is the fact that f is the PDF of the (joint) distribution of (X, Y) .

1. Suppose that Y has a discrete distribution on the countable set T . If X also has a discrete distribution on the countable set S then

$$\sum_{x \in A} g(x) \sum_{y \in B} h(y \mid x) = \sum_{x \in A} \sum_{y \in B} g(x) h(y \mid x) = \sum_{x \in A} \sum_{y \in B} f(x, y) = \mathbb{P}(X \in A, Y \in B), \quad A \subseteq S \quad (3.5.19)$$

In this jointly discrete case, there is a simpler argument of course:

$$h(y \mid x) = \frac{f(x, y)}{g(x)} = \frac{\mathbb{P}(X = x, Y = y)}{P(X = x)} = \mathbb{P}(Y = y \mid X = x), \quad y \in T \quad (3.5.20)$$

If X has a continuous distribution on $S \subseteq \mathbb{R}^j$ then

$$\int_A g(x) \sum_{y \in B} h(y \mid x) dx = \int_A \sum_{y \in B} g(x) h(y \mid x) dx = \int_A \sum_{y \in B} f(x, y) dx = \mathbb{P}(X \in A, Y \in B), \quad A \subseteq S \quad (3.5.21)$$

2. Suppose that Y has continuous distributions on $T \subseteq \mathbb{R}^k$. If X has a discrete distribution on the countable set S then

$$\sum_{x \in A} g(x) \int_B h(y \mid x) dy = \sum_{x \in A} \int_B g(x) h(y \mid x) dy = \sum_{x \in A} \int_B f(x, y) dy = \mathbb{P}(X \in A, Y \in B), \quad A \subseteq S \quad (3.5.22)$$

If X has a continuous distribution $S \subseteq \mathbb{R}^j$ then

$$\int_A g(x) \int_B h(y \mid x) dy dx = \int_A \int_B g(x) h(y \mid x) dy dx = \int_A \int_B f(x, y) dy dx = \mathbb{P}(X \in A, Y \in B), \quad A \subseteq S \quad (3.5.23)$$

The following theorem gives Bayes' theorem for probability density functions. We use the notation established above.

Bayes' Theorem. For $y \in T$, the conditional probability density function $x \mapsto g(x \mid y)$ of X given $y = y$ can be computed as follows:

1. If X has a discrete distribution then

$$g(x \mid y) = \frac{g(x)h(y \mid x)}{\sum_{s \in S} g(s)h(y \mid s)}, \quad x \in S \quad (3.5.24)$$

2. If X has a continuous distribution then

$$g(x \mid y) = \frac{g(x)h(y \mid x)}{\int_S g(s)h(y \mid s)ds}, \quad x \in S \quad (3.5.25)$$

Proof

In both cases the numerator is $f(x, y)$ while the denominator is $h(y)$.

In the context of Bayes' theorem, g is the *prior probability density function* of X and $x \mapsto g(x \mid y)$ is the *posterior probability density function* of X given $Y = y$ for $y \in T$. Note that the posterior probability density function $x \mapsto g(x \mid y)$ is proportional to the function $x \mapsto g(x)h(y \mid x)$. The sum or integral in the denominator is the normalizing constant.

Independence

Intuitively, X and Y should be independent if and only if the conditional distributions are the same as the corresponding unconditional distributions.

The following conditions are equivalent:

1. X and Y are independent.
2. $f(x, y) = g(x)h(y)$ for $x \in S, y \in T$
3. $h(y | x) = h(y)$ for $x \in S, y \in T$
4. $g(x | y) = g(x)$ for $x \in S, y \in T$

Proof

The equivalence of (a) and (b) was established in the section on joint distributions. Parts (c) and (d) are equivalent to (b). For a continuous distribution as described in the details in (4), a probability density function is not unique. The values of a PDF can be changed to other nonnegative values on a set of measure 0 and the resulting function is still a PDF. So if X or Y has a continuous distribution, the equations above have to be interpreted as holding for x or y , respectively, except on a set of measure 0.

Examples and Applications

In the exercises that follow, look for special models and distributions that we have studied. A special distribution may be embedded in a larger problem, as a conditional distribution, for example. In particular, a conditional distribution sometimes arises when a parameter of a standard distribution is *randomized*.

A couple of special distributions will occur frequently in the exercises. First, recall that the discrete uniform distribution on a finite, nonempty set S has probability density function f given by $f(x) = 1/\#(S)$ for $x \in S$. This distribution governs an element selected at random from S .

Recall also that *Bernoulli trials* (named for Jacob Bernoulli) are independent trials, each with two possible outcomes generically called *success* and *failure*. The probability of success $p \in [0, 1]$ is the same for each trial, and is the basic parameter of the random process. The number of successes in $n \in \mathbb{N}_+$ Bernoulli trials has the *binomial distribution* with parameters n and p . This distribution has probability density function f given by $f(x) = \binom{n}{x} p^x (1-p)^{n-x}$ for $x \in \{0, 1, \dots, n\}$. The binomial distribution is studied in more detail in the chapter on Bernoulli trials

Coins and Dice

Suppose that two standard, fair dice are rolled and the sequence of scores (X_1, X_2) is recorded. Let $U = \min\{X_1, X_2\}$ and $V = \max\{X_1, X_2\}$ denote the minimum and maximum scores, respectively.

1. Find the conditional probability density function of U given $V = v$ for each $v \in \{1, 2, 3, 4, 5, 6\}$.
2. Find the conditional probability density function of V given $U = u$ for each $u \in \{1, 2, 3, 4, 5, 6\}$

Answer

1.	$g(u v)$	$u = 1$	2	3	4	5	6
	$v = 1$	1	0	0	0	0	0
	2	$\frac{2}{3}$	$\frac{1}{3}$	0	0	0	0
	3	$\frac{2}{5}$	$\frac{2}{5}$	$\frac{1}{5}$	0	0	0
	4	$\frac{2}{7}$	$\frac{2}{7}$	$\frac{2}{7}$	$\frac{1}{7}$	0	0
	5	$\frac{2}{9}$	$\frac{2}{9}$	$\frac{2}{9}$	$\frac{2}{9}$	$\frac{1}{9}$	0
	6	$\frac{2}{11}$	$\frac{2}{11}$	$\frac{2}{11}$	$\frac{2}{11}$	$\frac{2}{11}$	$\frac{1}{11}$

2.	$h(v u)$	$u = 1$	2	3	4	5	6
	$v = 1$	$\frac{1}{11}$	0	0	0	0	0
	2	$\frac{2}{11}$	$\frac{1}{9}$	0	0	0	0
	3	$\frac{2}{11}$	$\frac{2}{9}$	$\frac{1}{7}$	0	0	0
	4	$\frac{2}{11}$	$\frac{2}{9}$	$\frac{2}{7}$	$\frac{1}{5}$	0	0
	5	$\frac{2}{11}$	$\frac{2}{9}$	$\frac{2}{7}$	$\frac{2}{5}$	$\frac{1}{3}$	0

$h(v u)$	$u = 1$	2	3	4	5	6
6	$\frac{2}{11}$	$\frac{2}{9}$	$\frac{2}{7}$	$\frac{2}{5}$	$\frac{2}{3}$	1

In the die-coin experiment, a standard, fair die is rolled and then a fair coin is tossed the number of times showing on the die. Let N denote the die score and Y the number of heads.

- Find the joint probability density function of (N, Y) .
- Find the probability density function of Y .
- Find the conditional probability density function of N given $Y = y$ for each $y \in \{0, 1, 2, 3, 4, 5, 6\}$

Answer

- and b.

$f(n, y)$	$n = 1$	2	3	4	5	6	$h(y)$
$y = 0$	$\frac{1}{12}$	$\frac{1}{24}$	$\frac{1}{48}$	$\frac{1}{96}$	$\frac{1}{102}$	$\frac{1}{384}$	$\frac{63}{384}$
1	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{16}$	$\frac{1}{24}$	$\frac{5}{192}$	$\frac{1}{64}$	$\frac{120}{384}$
2	0	$\frac{1}{24}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{5}{96}$	$\frac{5}{128}$	$\frac{99}{384}$
3	0	0	$\frac{1}{48}$	$\frac{1}{24}$	$\frac{5}{96}$	$\frac{5}{96}$	$\frac{64}{384}$
4	0	0	0	$\frac{1}{96}$	$\frac{5}{192}$	$\frac{5}{128}$	$\frac{29}{384}$
5	0	0	0	0	$\frac{1}{192}$	$\frac{1}{64}$	$\frac{8}{384}$
6	0	0	0	0	0	$\frac{1}{384}$	$\frac{1}{384}$
$g(n)$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	1

$g(n y)$	$n = 1$	2	3	4	5	6
$y = 0$	$\frac{32}{63}$	$\frac{16}{63}$	$\frac{8}{63}$	$\frac{4}{63}$	$\frac{2}{63}$	$\frac{1}{63}$
1	$\frac{16}{60}$	$\frac{16}{60}$	$\frac{12}{60}$	$\frac{8}{60}$	$\frac{5}{60}$	$\frac{3}{60}$
2	0	$\frac{16}{99}$	$\frac{24}{99}$	$\frac{24}{99}$	$\frac{20}{99}$	$\frac{15}{99}$
3	0	0	$\frac{2}{16}$	$\frac{4}{16}$	$\frac{5}{16}$	$\frac{5}{16}$
4	0	0	0	$\frac{4}{29}$	$\frac{10}{29}$	$\frac{15}{29}$
5	0	0	0	0	$\frac{1}{4}$	$\frac{3}{4}$
6	0	0	0	0	0	1

In the die-coin experiment, select the fair die and coin.

- Run the simulation of 1000 times and compare the empirical density function of Y with the true probability density function in the previous exercise
- Run the simulation 1000 times and compute the empirical conditional density function of N given $Y = 3$. Compare with the conditional probability density functions in the previous exercise.

In the coin-die experiment, a fair coin is tossed. If the coin is tails, a standard, fair die is rolled. If the coin is heads, a standard, ace-six flat die is rolled (faces 1 and 6 have probability $\frac{1}{4}$ each and faces 2, 3, 4, 5 have probability $\frac{1}{8}$ each). Let X denote the coin score (0 for tails and 1 for heads) and Y the die score.

- Find the joint probability density function of (X, Y) .
- Find the probability density function of Y .
- Find the conditional probability density function of X given $Y = y$ for each $y \in \{1, 2, 3, 4, 5, 6\}$

Answer

1. and b.

$f(x, y)$	$y = 1$	2	3	4	5	6	$g(x)$
$x = 0$	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{12}$	$\frac{1}{2}$
1	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{8}$	$\frac{1}{2}$
$h(y)$	$\frac{5}{24}$	$\frac{7}{24}$	$\frac{7}{48}$	$\frac{7}{48}$	$\frac{7}{48}$	$\frac{5}{24}$	1

3.

$g(x y)$	$y = 1$	2	3	4	5	6
$x = 0$	$\frac{2}{5}$	$\frac{4}{7}$	$\frac{4}{7}$	$\frac{4}{7}$	$\frac{4}{7}$	$\frac{2}{5}$
1	$\frac{3}{5}$	$\frac{3}{7}$	$\frac{3}{7}$	$\frac{3}{7}$	$\frac{3}{7}$	$\frac{3}{5}$

In the coin-die experiment, select the settings of the previous exercise.

1. Run the simulation 1000 times and compare the empirical density function of Y with the true probability density function in the previous exercise.
2. Run the simulation 100 times and compute the empirical conditional probability density function of X given $Y = 2$. Compare with the conditional probability density function in the previous exercise.

Suppose that a box contains 12 coins: 5 are fair, 4 are biased so that heads comes up with probability $\frac{1}{3}$, and 3 are two-headed. A coin is chosen at random and tossed 2 times. Let P denote the probability of heads of the selected coin, and X the number of heads.

1. Find the joint probability density function of (P, X) .
2. Find the probability density function of X .
3. Find the conditional probability density function of P given $X = x$ for $x \in \{0, 1, 2\}$.

Answer

1. and b.

$f(p, x)$	$x = 0$	1	2	$g(p)$
$p = \frac{1}{2}$	$\frac{5}{48}$	$\frac{10}{48}$	$\frac{5}{48}$	$\frac{5}{12}$
$\frac{1}{3}$	$\frac{4}{27}$	$\frac{4}{27}$	$\frac{1}{27}$	$\frac{4}{12}$
1	0	0	$\frac{1}{4}$	$\frac{3}{12}$
$h(x)$	$\frac{109}{432}$	$\frac{154}{432}$	$\frac{169}{432}$	1

3.

$g(p x)$	$x = 0$	1	2
$p = \frac{1}{2}$	$\frac{45}{109}$	$\frac{45}{77}$	$\frac{45}{169}$
$\frac{1}{3}$	$\frac{64}{109}$	$\frac{32}{77}$	$\frac{16}{169}$
1	0	0	$\frac{108}{169}$

Compare the [die-coin experiment](#) with the [box of coins experiment](#). In the first experiment, we toss a coin with a *fixed* probability of heads a *random* number of times. In the second experiment, we effectively toss a coin with a *random* probability of heads a *fixed* number of times.

Suppose that P has probability density function $g(p) = 6p(1 - p)$ for $p \in [0, 1]$. Given $P = p$, a coin with probability of heads p is tossed 3 times. Let X denote the number of heads.

1. Find the joint probability density function of (P, X) .
2. Find the probability density function of X .
3. Find the conditional probability density of P given $X = x$ for $x \in \{0, 1, 2, 3\}$. Graph these on the same axes.

Answer

1. $f(p, x) = 6 \binom{3}{x} p^{x+1} (1-p)^{4-x}$ for $p \in [0, 1]$ and $x \in \{0, 1, 2, 3, 4\}$
2. $h(0) = h(3) = \frac{1}{5}$, $h(1) = h(2) = \frac{3}{10}$.
3. $g(p | 0) = 30p(1-p)^4$, $g(p | 1) = 60p^2(1-p)^3$, $g(p | 2) = 60p^3(1-p)^2$, $g(p | 3) = 30p^4(1-p)$, in each case for $p \in [0, 1]$

Compare the [box of coins experiment](#) with the [last experiment](#). In the second experiment, we effectively choose a coin from a box with a continuous infinity of coin types. The prior distribution of P and each of the posterior distributions of P in part (c) are members of the family of *beta distributions*, one of the reasons for the importance of the beta family. Beta distributions are studied in more detail in the chapter on Special Distributions.

In the simulation of the beta coin experiment, set $a = b = 2$ and $n = 3$ to get the experiment studied in the previous exercise. For various “true values” of p , run the experiment in single step mode a few times and observe the posterior probability density function on each run.

Simple Mixed Distributions

Recall that the *exponential distribution* with rate parameter $r \in (0, \infty)$ has probability density function f given by $f(t) = re^{-rt}$ for $t \in [0, \infty)$. The exponential distribution is often used to model random times, under certain assumptions. The exponential distribution is studied in more detail in the chapter on the Poisson Process. Recall also that for $a, b \in \mathbb{R}$ with $a < b$, the continuous uniform distribution on the interval $[a, b]$ has probability density function f given by $f(x) = \frac{1}{b-a}$ for $x \in [a, b]$. This distribution governs a point selected at random from the interval.

Suppose that there are 5 light bulbs in a box, labeled 1 to 5. The lifetime of bulb n (in months) has the exponential distribution with rate parameter n . A bulb is selected at random from the box and tested.

1. Find the probability that the selected bulb will last more than one month.
2. Given that the bulb lasts more than one month, find the conditional probability density function of the bulb number.

Answer

Let N denote the bulb number and T the lifetime.

1. $\mathbb{P}(T > 1) = 0.1156$

2. n	1	2	3	4	5
$g(n T > 1)$	0.6364	0.2341	0.0861	0.0317	0.0117

Suppose that X is uniformly distributed on $\{1, 2, 3\}$, and given $X = x \in \{1, 2, 3\}$, random variable Y is uniformly distributed on the interval $[0, x]$.

1. Find the joint probability density function of (X, Y) .
2. Find the probability density function of Y .
3. Find the conditional probability density function of X given $Y = y$ for $y \in [0, 3]$.

Answer

1. $f(x, y) = \frac{1}{3x}$ for $y \in [0, x]$ and $x \in \{1, 2, 3\}$.
2. $h(y) = \begin{cases} \frac{11}{18}, & 0 \leq y \leq 1 \\ \frac{5}{18}, & 1 < y \leq 2 \\ \frac{2}{18}, & 2 < y \leq 3 \end{cases}$
3. For $y \in [0, 1]$, $g(1 | y) = \frac{6}{11}$, $g(2 | y) = \frac{3}{11}$, $g(3 | y) = \frac{2}{11}$
 For $y \in (1, 2]$, $g(1 | y) = 0$, $g(2 | y) = \frac{3}{5}$, $g(3 | y) = \frac{2}{5}$
 For $y \in (2, 3]$, $g(1 | y) = g(2 | y) = 0$, $g(3 | y) = 1$.

The Poisson Distribution

Recall that the Poisson distribution with parameter $a \in (0, \infty)$ has probability density function $g(n) = e^{-a} \frac{a^n}{n!}$ for $n \in \mathbb{N}$. This distribution is widely used to model the number of “random points” in a region of time or space; the parameter a is proportional to the size of the region. The Poisson distribution is named for Simeon Poisson, and is studied in more detail in the chapter on the Poisson Process.

Suppose that N is the number of elementary particles emitted by a sample of radioactive material in a specified period of time, and has the Poisson distribution with parameter a . Each particle emitted, independently of the others, is detected by a counter with probability $p \in (0, 1)$ and missed with probability $1 - p$. Let Y denote the number of particles detected by the counter.

1. For $n \in \mathbb{N}$, argue that the conditional distribution of Y given $N = n$ is binomial with parameters n and p .
2. Find the joint probability density function of (N, Y) .
3. Find the probability density function of Y .
4. For $y \in \mathbb{N}$, find the conditional probability density function of N given $Y = y$.

Answer

1. Each particle, independently, is detected (success) with probability p . This is the very definition of Bernoulli trials, so given $N = n$, the number of detected particles has the binomial distribution with parameters n and p .
2. The PDF f of (N, Y) is defined by

$$f(n, y) = e^{-a} a^n \frac{p^y (1-p)^{n-y}}{y! (n-y)!}, \quad n \in \mathbb{N}, y \in \{0, 1, \dots, n\} \quad (3.5.26)$$

3. The PDF h of Y is defined by

$$h(y) = e^{-pa} \frac{(pa)^y}{y!}, \quad y \in \mathbb{N} \quad (3.5.27)$$

This is the Poisson distribution with parameter pa .

4. The conditional PDF of N given $Y = y$ is defined by

$$g(n | y) = e^{-(1-p)a} \frac{[(1-p)a]^{n-y}}{(n-y)!}, \quad n \in \{y, y+1, \dots\} \quad (3.5.28)$$

This is the Poisson distribution with parameter $(1-p)a$, shifted to start at y .

The fact that Y also has a Poisson distribution is an interesting and characteristic property of the distribution. This property is explored in more depth in the section on thinning the Poisson process.

Simple Continuous Distributions

Suppose that (X, Y) has probability density function f defined by $f(x, y) = x + y$ for $(x, y) \in (0, 1)^2$.

1. Find the conditional probability density function of X given $Y = y$ for $y \in (0, 1)$
2. Find the conditional probability density function of Y given $X = x$ for $x \in (0, 1)$
3. Find $\mathbb{P}\left(\frac{1}{4} \leq Y \leq \frac{3}{4} \mid X = \frac{1}{3}\right)$.
4. Are X and Y independent?

Answer

1. For $y \in (0, 1)$, $g(x | y) = \frac{x+y}{y+1/2}$ for $x \in (0, 1)$
2. For $x \in (0, 1)$, $h(y | x) = \frac{x+y}{x+1/2}$ for $y \in (0, 1)$
3. $\frac{1}{2}$
4. X and Y are dependent.

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 2(x + y)$ for $0 < x < y < 1$.

1. Find the conditional probability density function of X given $Y = y$ for $y \in (0, 1)$.
2. Find the conditional probability density function of Y given $X = x$ for $x \in (0, 1)$.
3. Find $\mathbb{P}\left(Y \geq \frac{3}{4} \mid X = \frac{1}{2}\right)$.
4. Are X and Y independent?

Answer

1. For $y \in (0, 1)$, $g(x | y) = \frac{x+y}{3y^2}$ for $x \in (0, y)$.
2. For $x \in (0, 1)$, $h(y | x) = \frac{x+y}{(1+3x)(1-x)}$ for $y \in (x, 1)$.
3. $\frac{3}{10}$
4. X and Y are dependent.

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 15x^2y$ for $0 < x < y < 1$.

1. Find the conditional probability density function of X given $Y = y$ for $y \in (0, 1)$.
2. Find the conditional probability density function of Y given $X = x$ for $x \in (0, 1)$.
3. Find $\mathbb{P}\left(X \leq \frac{1}{4} \mid Y = \frac{1}{3}\right)$.
4. Are X and Y independent?

Answer

1. For $y \in (0, 1)$, $g(x \mid y) = \frac{3x^2}{y^3}$ for $x \in (0, y)$.
2. For $x \in (0, 1)$, $h(y \mid x) = \frac{2y}{1-x^2}$ for $y \in (x, 1)$.
3. $\frac{27}{64}$
4. X and Y are dependent.

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 6x^2y$ for $0 < x < 1$ and $0 < y < 1$.

1. Find the conditional probability density function of X given $Y = y$ for $y \in (0, 1)$.
2. Find the conditional probability density function of Y given $X = x$ for $x \in (0, 1)$.
3. Are X and Y independent?

Answer

1. For $y \in (0, 1)$, $g(x \mid y) = 3x^2$ for $y \in (0, 1)$.
2. For $x \in (0, 1)$, $h(y \mid x) = 2y$ for $y \in (0, 1)$.
3. X and Y are independent.

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 2e^{-x}e^{-y}$ for $0 < x < y < \infty$.

1. Find the conditional probability density function of X given $Y = y$ for $y \in (0, \infty)$.
2. Find the conditional probability density function of Y given $X = x$ for $x \in (0, \infty)$.
3. Are X and Y independent?

Answer

1. For $y \in (0, \infty)$, $g(x \mid y) = \frac{e^{-x}}{1-e^{-y}}$ for $x \in (0, y)$.
2. For $x \in (0, \infty)$, $h(y \mid x) = e^{x-y}$ for $y \in (x, \infty)$.
3. X and Y are dependent.

Suppose that X is uniformly distributed on the interval $(0, 1)$, and that given $X = x$, Y is uniformly distributed on the interval $(0, x)$.

1. Find the joint probability density function of (X, Y) .
2. Find the probability density function of Y .
3. Find the conditional probability density function of X given $Y = y$ for $y \in (0, 1)$.
4. Are X and Y independent?

Answer

1. $f(x, y) = \frac{1}{x}$ for $0 < y < x < 1$
2. $h(y) = -\ln y$ for $y \in (0, 1)$
3. For $y \in (0, 1)$, $g(x \mid y) = -\frac{1}{x \ln y}$ for $x \in (y, 1)$.
4. X and Y are dependent.

Suppose that X has probability density function g defined by $g(x) = 3x^2$ for $x \in (0, 1)$. The conditional probability density function of Y given $X = x$ is $h(y \mid x) = \frac{3y^2}{x^3}$ for $y \in (0, x)$.

1. Find the joint probability density function of (X, Y) .
2. Find the probability density function of Y .
3. Find the conditional probability density function of X given $Y = y$.
4. Are X and Y independent?

Answer

1. $f(x, y) = \frac{9y^2}{x}$ for $0 < y < x < 1$.
2. $h(y) = -9y^2 \ln y$ for $y \in (0, 1)$.

3. For $y \in (0, 1)$, $g(x | y) = -\frac{1}{x \ln y}$ for $x \in (y, 1)$.
4. X and Y are dependent.

Multivariate Uniform Distributions

Multivariate uniform distributions give a geometric interpretation of some of the concepts in this section.

Recall that For $n \in \mathbb{N}_+$, the standard measure λ_n on \mathbb{R}^n is given by

$$\lambda_n(A) = \int_A 1 \, dx, \quad A \subseteq \mathbb{R}^n \quad (3.5.29)$$

In particular, $\lambda_1(A)$ is the length of $A \subseteq \mathbb{R}$, $\lambda_2(A)$ is the area of $A \subseteq \mathbb{R}^2$ and $\lambda_3(A)$ is the volume of $A \subseteq \mathbb{R}^3$.

Details

Technically, λ_n is Lebesgue measure defined on the σ -algebra of measurable subsets of \mathbb{R}^n . In the discussion below, we assume that all sets are measurable. The integral representation is valid for the sets that occur in typical applications.

Suppose now that X takes values in \mathbb{R}^j , Y takes values in \mathbb{R}^k , and that (X, Y) is uniformly distributed on a set $R \subseteq \mathbb{R}^{j+k}$. So $0 < \lambda_{j+k}(R) < \infty$ and then the joint probability density function f of (X, Y) is given by $f(x, y) = 1/\lambda_{j+k}(R)$ for $(x, y) \in R$. Now let S and T be the projections of R onto \mathbb{R}^j and \mathbb{R}^k respectively, defined as follows:

$$S = \{x \in \mathbb{R}^j : (x, y) \in R \text{ for some } y \in \mathbb{R}^k\}, \quad T = \{y \in \mathbb{R}^k : (x, y) \in R \text{ for some } x \in \mathbb{R}^j\} \quad (3.5.30)$$

Note that $R \subseteq S \times T$. Next we denote the cross sections at $x \in S$ and at $y \in T$, respectively by

$$T_x = \{t \in T : (x, t) \in R\}, \quad S_y = \{s \in S : (s, y) \in R\} \quad (3.5.31)$$




Figure 3.5.1: The projections S and T , and the cross sections at x and y

In the last section on Joint Distributions, we saw that even though (X, Y) is uniformly distributed, the marginal distributions of X and Y are not uniform in general. However, as the next theorem shows, the conditional distributions are always uniform.

Suppose that (X, Y) is uniformly distributed on R . Then

1. The conditional distribution of Y given $X = x$ is uniformly on T_x for each $x \in S$.
2. The conditional distribution of X given $Y = y$ is uniformly on S_y for each $y \in T$.

Proof

The results are symmetric, so we will prove (a). Recall that X has PDF g given by

$$g(x) = \int_{T_x} f(x, y) \, dy = \int_{T_x} \frac{1}{\lambda_{j+k}(R)} \, dy = \frac{\lambda_k(T_x)}{\lambda_{j+k}(R)}, \quad x \in S \quad (3.5.32)$$

Hence for $x \in S$, the conditional PDF of Y given $X = x$ is

$$h(y | x) = \frac{f(x, y)}{g(x)} = \frac{1}{\lambda_k(T_x)}, \quad y \in T_x \quad (3.5.33)$$

and this is the PDF of the uniform distribution on T_x .

Find the conditional density of each variable given a value of the other, and determine if the variables are independent, in each of the following cases:

1. (X, Y) is uniformly distributed on the square $R = (-6, 6)^2$.

2. (X, Y) is uniformly distributed on the triangle $R = \{(x, y) \in \mathbb{R}^2 : -6 < y < x < 6\}$.
3. (X, Y) is uniformly distributed on the circle $R = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 < 36\}$.

Answer

The conditional PDF of X given $Y = y$ is denoted $x \mapsto g(x | y)$. The conditional PDF of Y given $X = x$ is denoted $y \mapsto h(y | x)$.

1.
 - For $y \in (-6, 6)$, $g(x | y) = \frac{1}{12}$ for $x \in (-6, 6)$.
 - For $x \in (-6, 6)$, $h(y | x) = \frac{1}{12}$ for $y \in (-6, 6)$.
 - X, Y are independent.
2.
 - For $y \in (-6, 6)$, $g(x | y) = \frac{1}{6-y}$ for $x \in (y, 6)$
 - For $x \in (-6, 6)$, $h(y | x) = \frac{1}{x+6}$ for $y \in (-6, x)$
 - X, Y are dependent.
3.
 - For $y \in (-6, 6)$, $g(x | y) = \frac{1}{2\sqrt{36-y^2}}$ for $x \in (-\sqrt{36-y^2}, \sqrt{36-y^2})$
 - For $x \in (-6, 6)$, $g(x | y) = \frac{1}{2\sqrt{36-x^2}}$ for $y \in (-\sqrt{36-x^2}, \sqrt{36-x^2})$
 - X, Y are dependent.

In the bivariate uniform experiment, run the simulation 1000 times in each of the following cases. Watch the points in the scatter plot and the graphs of the marginal distributions.

1. square
2. triangle
3. circle

Suppose that (X, Y, Z) is uniformly distributed on $R = \{(x, y, z) \in \mathbb{R}^3 : 0 < x < y < z < 1\}$.

1. Find the conditional density of each pair of variables given a value of the third variable.
2. Find the conditional density of each variable given values of the other two.

Answer

The subscripts 1, 2, and 3 correspond to the variables X, Y , and Z , respectively. Note that the conditions on (x, y, z) in each case are those in the definition of the domain R . They are stated differently to emphasize the domain of the conditional PDF as opposed to the given values, which function as parameters. Note also that each distribution is uniform on the appropriate region.

1. For $0 < z < 1$, $f_{1,2|3}(x, y | z) = \frac{2}{z^2}$ for $0 < x < y < z$
2. For $0 < y < 1$, $f_{1,3|2}(x, z | y) = \frac{1}{y(1-y)}$ for $0 < x < y$ and $y < z < 1$
3. For $0 < x < 1$, $f_{2,3|1}(y, z | x) = \frac{2}{(1-x)^2}$ for $x < y < z < 1$
4. For $0 < y < z < 1$, $f_{1|2,3}(x | y, z) = \frac{1}{y}$ for $0 < x < y$
5. For $0 < x < z < 1$, $f_{2|1,3}(y | x, z) = \frac{1}{z-x}$ for $x < y < z$
6. For $0 < x < y < 1$, $f_{3|1,2}(z | x, y) = \frac{1}{1-y}$ for $y < z < 1$

The Multivariate Hypergeometric Distribution

Recall the discussion of the (multivariate) hypergeometric distribution given in the last section on joint distributions. As in that discussion, suppose that a population consists of m objects, and that each object is one of four types. There are a objects of type 1, b objects of type 2, and c objects of type 3, and $m - a - b - c$ objects of type 0. We sample n objects from the population at random, and without replacement. The parameters a, b, c , and n are nonnegative integers with $a + b + c \leq m$ and $n \leq m$. Denote the number of type 1, 2, and 3 objects in the sample by X, Y , and Z , respectively. Hence, the number of type 0 objects in the sample is $n - X - Y - Z$. In the following exercises, $x, y, z \in \mathbb{N}$.

Suppose that $z \leq c$ and $n - m + c \leq z \leq n$. Then the conditional distribution of (X, Y) given $Z = z$ is hypergeometric, and has the probability density function defined by

$$g(x, y | z) = \frac{\binom{a}{x} \binom{b}{y} \binom{m-a-b-c}{n-x-y-z}}{\binom{m-c}{n-z}}, \quad x + y \leq n - z \quad (3.5.34)$$

Proof

This result can be proved analytically but a combinatorial argument is better. The essence of the argument is that we are selecting a random sample of size $n - z$ without replacement from a population of size $m - c$, with a objects of type 1, b objects of type 2, and $m - a - b - c$ objects of type 0. The conditions on z ensure that $\mathbb{P}(Z = z) > 0$, or equivalently, that the new parameters make sense.

Suppose that $y \leq b$, $z \leq c$, and $n - m + b \leq y + z \leq n$. Then the conditional distribution of X given $Y = y$ and $Z = z$ is hypergeometric, and has the probability density function defined by

$$g(x | y, z) = \frac{\binom{a}{x} \binom{m-a-b-c}{n-x-y-z}}{\binom{m-b-c}{n-y-z}}, \quad x \leq n - y - z \quad (3.5.35)$$

Proof

Again, this result can be proved analytically, but a combinatorial argument is better. The essence of the argument is that we are selecting a random sample of size $n - y - z$ from a population of size $m - b - c$, with a objects of type 1 and $m - a - b - c$ objects type 0. The conditions on y and z ensure that $\mathbb{P}(Y = y, Z = z) > 0$, or equivalently that the new parameters make sense.

These results generalize in a completely straightforward way to a population with any number of types. In brief, if a random vector has a hypergeometric distribution, then the conditional distribution of some of the variables, given values of the other variables, is also hypergeometric. Moreover, it is clearly not necessary to remember the hideous formulas in the previous two theorems. You just need to recognize the problem as sampling without replacement from a multi-type population, and then identify the number of objects of each type and the sample size. The hypergeometric distribution and the multivariate hypergeometric distribution are studied in more detail in the chapter on Finite Sampling Models.

In a population of 150 voters, 60 are democrats and 50 are republicans and 40 are independents. A sample of 15 voters is selected at random, without replacement. Let X denote the number of democrats in the sample and Y the number of republicans in the sample. Give the probability density function of each of the following:

1. (X, Y)
2. X
3. Y given $X = 5$

Answer

1. $f(x, y) = \frac{1}{\binom{150}{15}} \binom{60}{x} \binom{50}{y} \binom{40}{15-x-y}$ for $x + y \leq 15$
2. $g(x) = \frac{1}{\binom{150}{15}} \binom{60}{x} \binom{90}{15-x}$ for $x \leq 15$
3. $h(y | 5) = \frac{1}{\binom{90}{10}} \binom{50}{y} \binom{40}{10-y}$ for $y \leq 10$

Recall that a *bridge hand* consists of 13 cards selected at random and without replacement from a standard deck of 52 cards. Let X , Y , and Z denote the number of spades, hearts, and diamonds, respectively, in the hand. Find the probability density function of each of the following:

1. (X, Y, Z)
2. (X, Y)
3. X
4. (X, Y) given $Z = 3$
5. X given $Y = 3$ and $Z = 2$

Answer

1. $f(x, y, z) = \frac{1}{\binom{52}{13}} \binom{13}{x} \binom{13}{y} \binom{13}{z} \binom{13}{13-x-y-z}$ for $x + y + z \leq 13$.
2. $g(x, y) = \frac{1}{\binom{52}{13}} \binom{13}{x} \binom{13}{y} \binom{26}{13-x-y}$ for $x + y \leq 13$
3. $h(x) = \frac{1}{\binom{52}{13}} \binom{13}{x} \binom{39}{13-x}$ for $x \leq 13$
4. $g(x, y | 3) = \frac{1}{\binom{39}{10}} \binom{13}{x} \binom{13}{y} \binom{13}{10-x-y}$ for $x + y \leq 10$
5. $h(x | 3, 2) = \frac{1}{\binom{26}{8}} \binom{13}{x} \binom{13}{8-x}$ for $x \leq 8$

Multinomial Trials

Recall the discussion of *multinomial trials* in the last section on joint distributions. As in that discussion, suppose that we have a sequence of n independent trials, each with 4 possible outcomes. On each trial, outcome 1 occurs with probability p , outcome 2 with probability q , outcome 3 with probability r , and outcome 0 with probability $1 - p - q - r$. The parameters $p, q, r \in (0, 1)$, with $p + q + r < 1$, and $n \in \mathbb{N}_+$. Denote the number of times that outcome 1, outcome 2, and outcome 3 occurs in the n trials by X, Y , and Z respectively. Of course, the number of times that outcome 0 occurs is $n - X - Y - Z$. In the following exercises, $x, y, z \in \mathbb{N}$.

For $z \leq n$, the conditional distribution of (X, Y) given $Z = z$ is also multinomial, and has the probability density function.

$$g(x, y | z) = \binom{n-z}{x, y} \left(\frac{p}{1-r} \right)^x \left(\frac{q}{1-r} \right)^y \left(1 - \frac{p}{1-r} - \frac{q}{1-r} \right)^{n-x-y-z}, \quad x + y \leq n - z \quad (3.5.36)$$

Proof

This result can be proved analytically, but a probability argument is better. First, let I denote the outcome of a generic trial. Then $\mathbb{P}(I = 1 | I \neq 3) = \mathbb{P}(I = 1) / \mathbb{P}(I \neq 3) = p / (1 - r)$. Similarly, $\mathbb{P}(I = 2 | I \neq 3) = q / (1 - r)$ and $\mathbb{P}(I = 0 | I \neq 3) = (1 - p - q - r) / (1 - r)$. Now, the essence of the argument is that effectively, we have $n - z$ independent trials, and on each trial, outcome 1 occurs with probability $p / (1 - r)$ and outcome 2 with probability $q / (1 - r)$.

For $y + z \leq n$, the conditional distribution of X given $Y = y$ and $Z = z$ is binomial, with the probability density function

$$h(x | y, z) = \binom{n-y-z}{x} \left(\frac{p}{1-q-r} \right)^x \left(1 - \frac{p}{1-q-r} \right)^{n-x-y-z}, \quad x \leq n - y - z \quad (3.5.37)$$

Proof

Again, this result can be proved analytically, but a probability argument is better. As before, let I denote the outcome of a generic trial. Then $\mathbb{P}(I = 1 | I \notin \{2, 3\}) = p / (1 - q - r)$ and $\mathbb{P}(I = 0 | I \notin \{2, 3\}) = (1 - p - q - r) / (1 - q - r)$. Thus, the essence of the argument is that effectively, we have $n - y - z$ independent trials, and on each trial, outcome 1 occurs with probability $p / (1 - q - r)$.

These results generalize in a completely straightforward way to multinomial trials with any number of trial outcomes. In brief, if a random vector has a multinomial distribution, then the conditional distribution of some of the variables, given values of the other variables, is also multinomial. Moreover, it is clearly not necessary to remember the specific formulas in the previous two exercises. You just need to recognize a problem as one involving independent trials, and then identify the probability of each outcome and the number of trials. The binomial distribution and the multinomial distribution are studied in more detail in the chapter on Bernoulli Trials.

Suppose that peaches from an orchard are classified as *small*, *medium*, or *large*. Each peach, independently of the others is small with probability $\frac{3}{10}$, medium with probability $\frac{1}{2}$, and large with probability $\frac{1}{5}$. In a sample of 20 peaches from the orchard, let X denote the number of small peaches and Y the number of medium peaches. Give the probability density function of each of the following:

1. (X, Y)
2. X
3. Y given $X = 5$

Answer

1. $f(x, y) = \binom{20}{x, y} \left(\frac{3}{10} \right)^x \left(\frac{1}{2} \right)^y \left(\frac{1}{5} \right)^{20-x-y}$ for $x + y \leq 20$
2. $g(x) = \binom{20}{x} \left(\frac{3}{10} \right)^x \left(\frac{7}{10} \right)^{20-x}$ for $x \leq 20$
3. $h(y | 5) = \binom{15}{y} \left(\frac{5}{7} \right)^y \left(\frac{2}{7} \right)^{15-y}$ for $y \leq 15$

For a certain crooked, 4-sided die, face 1 has probability $\frac{2}{5}$, face 2 has probability $\frac{3}{10}$, face 3 has probability $\frac{1}{5}$, and face 4 has probability $\frac{1}{10}$. Suppose that the die is thrown 50 times. Let X, Y , and Z denote the number of times that scores 1, 2, and 3 occur, respectively. Find the probability density function of each of the following:

1. (X, Y, Z)
2. (X, Y)
3. X
4. (X, Y) given $Z = 5$

5. X given $Y = 10$ and $Z = 5$

Answer

1. $f(x, y, z) = \binom{50}{x, y, z} \left(\frac{2}{5}\right)^x \left(\frac{3}{10}\right)^y \left(\frac{1}{5}\right)^z \left(\frac{1}{10}\right)^{50-x-y-z}$ for $x + y + z \leq 50$
2. $g(x, y) = \binom{50}{x, y} \left(\frac{2}{5}\right)^x \left(\frac{3}{10}\right)^y \left(\frac{3}{10}\right)^{50-x-y}$ for $x + y \leq 50$
3. $h(x) = \binom{50}{x} \left(\frac{2}{5}\right)^x \left(\frac{3}{5}\right)^{50-x}$ for $x \leq 50$
4. $g(x, y | 5) = \binom{45}{x, y} \left(\frac{1}{2}\right)^x \left(\frac{3}{8}\right)^y \left(\frac{1}{8}\right)^{45-x-y}$ for $x + y \leq 45$
5. $h(x | 10, 5) = \binom{35}{x} \left(\frac{4}{5}\right)^x \left(\frac{1}{4}\right)^{10-x}$ for $x \leq 35$

Bivariate Normal Distributions

The joint distributions in the next two exercises are examples of *bivariate normal distributions*. The conditional distributions are also normal, an important property of the bivariate normal distribution. In general, normal distributions are widely used to model physical measurements subject to small, random errors. The bivariate normal distribution is studied in more detail in the chapter on Special Distributions.

Suppose that (X, Y) has the bivariate normal distribution with probability density function f defined by

$$f(x, y) = \frac{1}{12\pi} \exp\left[-\left(\frac{x^2}{8} + \frac{y^2}{18}\right)\right], \quad (x, y) \in \mathbb{R}^2 \quad (3.5.38)$$

1. Find the conditional probability density function of X given $Y = y$ for $y \in \mathbb{R}$.
2. Find the conditional probability density function of Y given $X = x$ for $x \in \mathbb{R}$.
3. Are X and Y independent?

Answer

1. For $y \in \mathbb{R}$, $g(x | y) = \frac{1}{2\sqrt{2\pi}} e^{-x^2/8}$ for $x \in \mathbb{R}$. This is the PDF of the normal distribution with mean 0 and variance 4.
2. For $x \in \mathbb{R}$, $h(y | x) = \frac{1}{3\sqrt{2\pi}} e^{-y^2/18}$ for $y \in \mathbb{R}$. This is the PDF of the normal distribution with mean 0 and variance 9.
3. X and Y are independent.

Suppose that (X, Y) has the bivariate normal distribution with probability density function f defined by

$$f(x, y) = \frac{1}{\sqrt{3}\pi} \exp\left[-\frac{2}{3}(x^2 - xy + y^2)\right], \quad (x, y) \in \mathbb{R}^2 \quad (3.5.39)$$

1. Find the conditional probability density function of X given $Y = y$ for $y \in \mathbb{R}$.
2. Find the conditional probability density function of Y given $X = x$ for $x \in \mathbb{R}$.
3. Are X and Y independent?

Answer

1. For $y \in \mathbb{R}$, $g(x | y) = \sqrt{\frac{2}{3\pi}} e^{-\frac{2}{3}(x-y/2)^2}$ for $x \in \mathbb{R}$. This is the PDF of the normal distribution with mean $y/2$ and variance $3/4$.
2. For $x \in \mathbb{R}$, $h(y | x) = \sqrt{\frac{2}{3\pi}} e^{-\frac{2}{3}(y-x/2)^2}$ for $y \in \mathbb{R}$. This is the PDF of the normal distribution with mean $x/2$ and variance $3/4$.
3. X and Y are dependent.

Mixtures of Distributions

With our usual sets S and T , as above, suppose that P_x is a probability measure on T for each $x \in S$. Suppose also that g is a probability density function on S . We can obtain a new probability measure on T by averaging (or *mixing*) the given distributions according to g .

First suppose that g is the probability density function of a discrete distribution on the countable set S . Then the function \mathbb{P} defined below is a probability measure on T :

$$\mathbb{P}(B) = \sum_{x \in S} g(x) P_x(B), \quad B \subseteq T \quad (3.5.40)$$

Proof

Clearly $\mathbb{P}(B) \geq 0$ for $B \subseteq T$ and $\mathbb{P}(T) = \sum_{x \in S} g(x) 1 = 1$. Suppose that $\{B_i : i \in I\}$ is a countable, disjoint collection of subsets of T . Then

$$\mathbb{P}\left(\bigcup_{i \in I} B_i\right) = \sum_{x \in S} g(x) P_x\left(\bigcup_{i \in I} B_i\right) = \sum_{x \in S} g(x) \sum_{i \in I} P_x(B_i) = \sum_{i \in I} \sum_{x \in S} g(x) P_x(B_i) = \sum_{i \in I} \mathbb{P}(B_i) \quad (3.5.41)$$

Reversing the order of summation is justified since the terms are nonnegative.

In the setting of the previous theorem, suppose that P_x has probability density function h_x for each $x \in S$. Then \mathbb{P} has probability density function h given by

$$h(y) = \sum_{x \in S} g(x) h_x(y), \quad y \in T \quad (3.5.42)$$

Proof

As usual, we will consider the discrete and continuous cases for the distributions on T separately.

1. Suppose that T is countable so that P_x is a discrete probability measure for each $x \in S$. By definition, for each $x \in S$, $h_x(y) = \mathbb{P}_x(\{y\})$ for $y \in T$. So the probability density function h of P is given by

$$h(y) = P(\{y\}) = \sum_{x \in S} g(x) P_x(\{y\}) = \sum_{x \in S} g(x) h_x(y), \quad y \in T \quad (3.5.43)$$

2. Suppose now that P_x has a continuous distribution on $T \subseteq \mathbb{R}^k$, with PDF g_x for each $x \in S$. For $B \subseteq T$,

$$\mathbb{P}(B) = \sum_{x \in S} g(x) P_x(B) = \sum_{x \in S} g(x) \int_B h_x(y) dy = \int_B \sum_{x \in S} g(x) h_x(y) dy = \int_B h(y) dy \quad (3.5.44)$$

So by definition, h is the PDF of \mathbb{P} . Again, the interchange of sum and integral is justified because the functions are nonnegative. Technically, we also need $y \mapsto h_x(y)$ to be measurable for $x \in S$ so that the integral makes sense.

Conversely, given a probability density function g on S and a probability density function h_x on T for each $x \in S$, the function h defined in the previous theorem is a probability density function on T .

Suppose now that g is the probability density function of a continuous distribution on $S \subseteq \mathbb{R}^j$. Then the function \mathbb{P} defined below is a probability measure on T :

$$\mathbb{P}(B) = \int_S g(x) P_x(B) dx, \quad B \subseteq T \quad (3.5.45)$$

Proof

The proof is just like the proof of Theorem (45) with integrals over S replacing the sums over S . Clearly $\mathbb{P}(B) \geq 0$ for $B \subseteq T$ and $\mathbb{P}(T) = \int_S g(x) P_x(T) dx = \int_S g(x) dx = 1$. Suppose that $\{B_i : i \in I\}$ is a countable, disjoint collection of subsets of T . Then

$$\mathbb{P}\left(\bigcup_{i \in I} B_i\right) = \int_S g(x) P_x\left(\bigcup_{i \in I} B_i\right) dx = \int_S g(x) \sum_{i \in I} P_x(B_i) dx = \sum_{i \in I} \int_S g(x) P_x(B_i) dx = \sum_{i \in I} \mathbb{P}(B_i) \quad (3.5.46)$$

Reversing the integral and the sum is justified since the terms are nonnegative. Technically, we need the subsets of T and the mapping $x \mapsto P_x(B)$ to be measurable.

In the setting of the previous theorem, suppose that P_x is a discrete (respectively continuous) distribution with probability density function h_x for each $x \in S$. Then \mathbb{P} is also discrete (respectively continuous) with probability density function h given by

$$h(y) = \int_S g(x) h_x(y) dx, \quad y \in T \quad (3.5.47)$$

Proof

The proof is just like the proof of Theorem (46) with integrals over S replacing the sums over S .

1. Suppose that T is countable so that P_x is a discrete probability measure for each $x \in S$. By definition, for each $x \in S$, $h_x(y) = \mathbb{P}_x(\{y\})$ for $y \in T$. So the probability density function h of P is given by

$$h(y) = P(\{y\}) = \int_S g(x)P_x(\{y\}) dx = \int_S g(x)h_x(y) dx, \quad y \in T \quad (3.5.48)$$

Technically, we need $x \mapsto P_x(\{y\}) = h_x(y)$ to be measurable for $y \in T$.

2. Suppose now that P_x has a continuous distribution on $T \subseteq \mathbb{R}^k$, with PDF g_x for each $x \in S$. For $B \subseteq T$,

$$\mathbb{P}(B) = \int_S g(x)P_x(B) dx = \int_S g(x) \int_B h_x(y) dy dx = \int_B \int_S g(x)h_x(y) dx dy = \int_B h(y) dy \quad (3.5.49)$$

So by definition, h is the PDF of \mathbb{P} . Again, the interchange of sum and integral is justified because the functions are nonnegative.

Technically, we also need $(x, y) \mapsto h_x(y)$ to be measurable so that the integral makes sense.

In both cases, the distribution \mathbb{P} is said to be a *mixture* of the set of distributions $\{P_x : x \in S\}$, with mixing density g .

One can have a mixture of distributions, without having random variables defined on a common probability space. However, mixtures are intimately related to conditional distributions. Returning to our usual setup, suppose that X and Y are random variables for an experiment, taking values in S and T respectively and that X probability density function g . The following result is simply a restatement of the law of total probability.

The distribution of Y is a mixture of the conditional distributions of Y given $X = x$, over $x \in S$, with mixing density g .

Proof

Only the notation is different.

1. If X has a discrete distribuion on the countable set S then

$$\mathbb{P}(Y \in B) = \sum_{x \in S} g(x)\mathbb{P}(Y \in B \mid X = x), \quad B \subseteq T \quad (3.5.50)$$

2. If X has a continuous distribution $S \subseteq \mathbb{R}^j$ then

$$\mathbb{P}(Y \in B) = \int_S g(x)\mathbb{P}(Y \in B \mid X = x) dx, \quad B \subseteq T \quad (3.5.51)$$

Finally we note that a mixed distribution (with discrete and continuous parts) really is a mixture, in the sense of this discussion.

Suppose that \mathbb{P} is a mixed distribution on a set T . Then \mathbb{P} is a mixture of a discrete distribution and a continuous distribution.

Proof

Recall that *mixed distribution* means that T can be partitioned into a countable set D and a set $C \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$ with the properties that $\mathbb{P}(\{x\}) > 0$ for $x \in D$, $\mathbb{P}(\{x\}) = 0$ for $x \in C$, and $p = \mathbb{P}(D) \in (0, 1)$. Let $S = \{d, c\}$ and define the PDF g on S by $g(d) = p$ and $g(c) = 1 - p$. Recall that the conditional distribution P_d defined by $P_d(A) = \mathbb{P}(A \cap D)/\mathbb{P}(D)$ for $A \subseteq T$ is a discrete distribution on T and similarly the conditional distribution P_c defined by $P_c(A) = \mathbb{P}(A \cap C)/\mathbb{P}(C)$ for $A \subseteq T$ is a continuous distribution on T . Clearly with this setup,

$$\mathbb{P}(A) = g(c)P_c(A) + g(d)P_d(A), \quad A \subseteq T \quad (3.5.52)$$

This page titled [3.5: Conditional Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.6: Distribution and Quantile Functions

As usual, our starting point is a random experiment modeled by a with probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} is the collection of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . In this section, we will study two types of functions that can be used to specify the distribution of a real-valued random variable.

Distribution Functions

Definition

Suppose that X is a random variable with values in \mathbb{R} . The (cumulative) *distribution function* of X is the function $F : \mathbb{R} \rightarrow [0, 1]$ defined by

$$F(x) = \mathbb{P}(X \leq x), \quad x \in \mathbb{R} \quad (3.6.1)$$

The distribution function is important because it makes sense for any type of random variable, regardless of whether the distribution is discrete, continuous, or even mixed, and because it completely determines the distribution of X . In the picture below, the light shading is intended to represent a continuous distribution of probability, while the darker dots represents points of positive probability; $F(x)$ is the total probability mass to the left of (and including) x .




Figure 3.6.1: $F(x)$ is the total probability to the left of (and including) x

Basic Properties

A few basic properties completely characterize distribution functions. Notationally, it will be helpful to abbreviate the limits of F from the left and right at $x \in \mathbb{R}$, and at ∞ and $-\infty$ as follows:

$$F(x^+) = \lim_{t \downarrow x} F(t), \quad F(x^-) = \lim_{t \uparrow x} F(t), \quad F(\infty) = \lim_{t \rightarrow \infty} F(t), \quad F(-\infty) = \lim_{t \rightarrow -\infty} F(t) \quad (3.6.2)$$

Suppose that F is the distribution function of a real-valued random variable X .

1. F is increasing: if $x \leq y$ then $F(x) \leq F(y)$.
2. $F(x^+) = F(x)$ for $x \in \mathbb{R}$. Thus, F is *continuous from the right*.
3. $F(x^-) = \mathbb{P}(X < x)$ for $x \in \mathbb{R}$. Thus, F has *limits from the left*.
4. $F(-\infty) = 0$.
5. $F(\infty) = 1$.

Proof




Figure 3.6.2: The graph of a distribution function

The following result shows how the distribution function can be used to compute the probability that X is in an interval. Recall that a probability distribution on \mathbb{R} is completely determined by the probabilities of *intervals*; thus, the *distribution function* determines

the *distribution* of X .

Suppose again that F is the distribution function of a real-valued random variable X . If $a, b \in \mathbb{R}$ with $a < b$ then

1. $\mathbb{P}(X = a) = F(a) - F(a^-)$
2. $\mathbb{P}(a < X \leq b) = F(b) - F(a)$
3. $\mathbb{P}(a < X < b) = F(b^-) - F(a)$
4. $\mathbb{P}(a \leq X \leq b) = F(b) - F(a^-)$
5. $\mathbb{P}(a \leq X < b) = F(b^-) - F(a^-)$

Proof

These results follow from the [definition](#), the [basic properties](#), and the difference rule: $\mathbb{P}(B \setminus A) = \mathbb{P}(B) - \mathbb{P}(A)$ if A, B are events and $A \subseteq B$.

1. $\{X = a\} = \{X \leq a\} \setminus \{X < a\}$, so $\mathbb{P}(X = a) = \mathbb{P}(X \leq a) - \mathbb{P}(X < a) = F(a) - F(a^-)$.
2. $\{a < X \leq b\} = \{X \leq b\} \setminus \{X \leq a\}$, so $\mathbb{P}(a < X \leq b) = \mathbb{P}(X \leq b) - \mathbb{P}(X \leq a) = F(b) - F(a)$.
3. $\{a < X < b\} = \{X < b\} \setminus \{X \leq a\}$, so $\mathbb{P}(a < X < b) = \mathbb{P}(X < b) - \mathbb{P}(X \leq a) = F(b^-) - F(a)$.
4. $\{a \leq X \leq b\} = \{X \leq b\} \setminus \{X < a\}$, so $\mathbb{P}(a \leq X \leq b) = \mathbb{P}(X \leq b) - \mathbb{P}(X < a) = F(b) - F(a^-)$.
5. $\{a \leq X < b\} = \{X < b\} \setminus \{X < a\}$, so $\mathbb{P}(a \leq X < b) = \mathbb{P}(X < b) - \mathbb{P}(X < a) = F(b^-) - F(a^-)$.

Conversely, if a Function $F : \mathbb{R} \rightarrow [0, 1]$ satisfies the [basic properties](#), then the [formulas above](#) define a probability distribution on \mathbb{R} , with F as the distribution function. For more on this point, read the section on Existence and Uniqueness.

If X has a continuous distribution, then the distribution function F is continuous.

Proof

If X has a continuous distribution, then by definition, $\mathbb{P}(X = x) = 0$ so $\mathbb{P}(X < x) = \mathbb{P}(X \leq x)$ for $x \in \mathbb{R}$. Hence from part (a) of the [previous theorem](#), $F(x^-) = F(x^+) = F(x)$.

Thus, the two meanings of *continuous* come together: continuous distribution and continuous function in the calculus sense. Next recall that the distribution of a real-valued random variable X is *symmetric* about a point $a \in \mathbb{R}$ if the distribution of $X - a$ is the same as the distribution of $a - X$.

Suppose that X has a continuous distribution on \mathbb{R} that is symmetric about a point a . Then the distribution function F satisfies $F(a - t) = 1 - F(a + t)$ for $t \in \mathbb{R}$.

Proof

Since $X - a$ and $a - X$ have the same distribution,

$$F(a - t) = \mathbb{P}(X \leq a - t) = \mathbb{P}(X - a \leq -t) = \mathbb{P}(a - X \leq -t) = \mathbb{P}(X \geq a + t) = 1 - F(a + t) \quad (3.6.3)$$

Relation to Density Functions

There are simple relationships between the distribution function and the probability density function. Recall that if X takes value in $S \subseteq \mathbb{R}$ and has probability density function f , we can extend f to all of \mathbb{R} by the convention that $f(x) = 0$ for $x \in S^c$. As in Definition (1), it's customary to define the distribution function F on all of \mathbb{R} , even if the random variable takes values in a subset.

Suppose that X has discrete distribution on a countable subset $S \subseteq \mathbb{R}$. Let f denote the probability density function and F the distribution function.

1. $F(x) = \sum_{t \in S, t \leq x} f(t)$ for $x \in \mathbb{R}$
2. $f(x) = F(x) - F(x^-)$ for $x \in S$

Proof

1. This follows from the definition of the PDF of X , $f(t) = \mathbb{P}(X = t)$ for $t \in S$, and the additivity of probability.
2. This is a restatement of part (a) of the [theorem above](#).

Thus, F is a step function with jumps at the points in S ; the size of the jump at x is $f(x)$.




Figure 3.6.3: The distribution function of a discrete distribution

There is an analogous result for a continuous distribution with a probability density function.

Suppose that X has a continuous distribution on \mathbb{R} with probability density function f and distribution function F .

1. $F(x) = \int_{-\infty}^x f(t)dt$ for $x \in \mathbb{R}$.
2. $f(x) = F'(x)$ if f is continuous at x .

Proof




Figure 3.6.4: The distribution function of a continuous distribution

The last result is the basic probabilistic version of the fundamental theorem of calculus. For mixed distributions, we have a combination of the results in the last two theorems.

Suppose that X has a mixed distribution, with discrete part on a countable subset $D \subseteq \mathbb{R}$, and continuous part on $\mathbb{R} \setminus D$. Let g denote the partial probability density function of the discrete part and assume that the continuous part has partial probability density function h . Let F denote the distribution function.

1. $F(x) = \sum_{t \in D, t \leq x} g(t) + \int_{-\infty}^x h(t)dt$ for $x \in \mathbb{R}$
2. $g(x) = F(x) - F(x^-)$ for $x \in D$
3. $h(x) = F'(x)$ if $x \notin D$ and h is continuous at x

Go back to the [graph](#) of a general distribution function. At a point of positive probability, the probability is the size of the jump. At a smooth point of the graph, the continuous probability density is the slope.

Recall that the existence of a probability density function is not guaranteed for a continuous distribution, but of course the distribution function always makes perfect sense. The advanced section on absolute continuity and density functions has an example of a continuous distribution on the interval $(0, 1)$ that has no probability density function. The distribution function is continuous and strictly increases from 0 to 1 on the interval, but has derivative 0 at almost every point!

Naturally, the distribution function can be defined relative to any of the conditional distributions we have discussed. No new concepts are involved, and all of the results above hold.

Reliability

Suppose again that X is a real-valued random variable with distribution function F . The function in the following definition clearly gives the same information as F .

The function F^c defined by

$$F^c(x) = 1 - F(x) = \mathbb{P}(X > x), \quad x \in \mathbb{R} \quad (3.6.4)$$

is the *right-tail distribution function* of X . Give the mathematical properties of F^c analogous to the properties of F in (2).

Answer

1. F^c is decreasing.
2. $F^c(t) \rightarrow F^c(x)$ as $t \downarrow x$ for $x \in \mathbb{R}$, so F^c is continuous from the right.
3. $F^c(t) \rightarrow \mathbb{P}(X \geq x)$ as $t \uparrow x$ for $x \in \mathbb{R}$, so F^c has left limits.
4. $F^c(x) \rightarrow 0$ as $x \rightarrow \infty$.
5. $F^c(x) \rightarrow 1$ as $x \rightarrow -\infty$.

So F might be called the *left-tail distribution function*. But why have two distribution functions that give essentially the same information? The right-tail distribution function, and related functions, arise naturally in the context of *reliability theory*. For the remainder of this subsection, suppose that T is a random variable with values in $[0, \infty)$ and that T has a continuous distribution with probability density function f . Here are the important definitions:

Suppose that T represents the lifetime of a device.

1. The right tail distribution function F^c is the *reliability function* of T .
2. The function h defined by $h(t) = f(t)/F^c(t)$ for $t \geq 0$ is the *failure rate function* of T .

To interpret the reliability function, note that $F^c(t) = \mathbb{P}(T > t)$ is the probability that the device lasts at least t time units. To interpret the failure rate function, note that if dt is “small” then

$$\mathbb{P}(t < T < t + dt \mid T > t) = \frac{\mathbb{P}(t < T < t + dt)}{\mathbb{P}(T > t)} \approx \frac{f(t) dt}{F^c(t)} = h(t) dt \quad (3.6.5)$$

So $h(t) dt$ is the approximate probability that the device will fail in the interval $(t, t + dt)$, given survival up to time t . Moreover, like the distribution function and the reliability function, the failure rate function also completely determines the distribution of T .

The reliability function can be expressed in terms of the failure rate function by

$$F^c(t) = \exp\left(-\int_0^t h(s) ds\right), \quad t \geq 0 \quad (3.6.6)$$

Proof

At the points of continuity of f we have $\frac{d}{dt} F^c(t) = -f(t)$. Hence

$$\int_0^t h(s) ds = \int_0^t \frac{f(s)}{F^c(s)} ds = \int_0^t -\frac{\frac{d}{ds} F^c(s)}{F^c(s)} ds = -\ln[F^c(t)] \quad (3.6.7)$$

The failure rate function h satisfies the following properties:

1. $h(t) \geq 0$ for $t \geq 0$
2. $\int_0^\infty h(t) dt = \infty$

Proof

1. This follows from the definition.
2. This follows from the [previous result](#) and the fact that $F^c(t) \rightarrow 0$ as $t \rightarrow \infty$.

Conversely, a function that satisfies these properties is the failure rate function for a continuous distribution on $[0, \infty)$:

Suppose that $h : [0, \infty) \rightarrow [0, \infty)$ is piecewise continuous and $\int_0^\infty h(t) dt = \infty$. Then the function G defined by

$$F^c(t) = \exp\left(-\int_0^t h(s) ds\right), \quad t \geq 0 \quad (3.6.8)$$

is a reliability function for a continuous distribution on $[0, \infty)$

Proof

The function F^c is continuous, decreasing, and satisfies $F^c(0) = 1$ and $F^c(t) \rightarrow 0$ as $t \rightarrow \infty$. Hence $F = 1 - F^c$ is the distribution function for a continuous distribution on $[0, \infty)$.

Multivariate Distribution Functions

Suppose now that X and Y are real-valued random variables for an experiment (that is, defined on the same probability space), so that (X, Y) is random vector taking values in a subset of \mathbb{R}^2 .

The *distribution function* of (X, Y) is the function F defined by

$$F(x, y) = \mathbb{P}(X \leq x, Y \leq y), \quad (x, y) \in \mathbb{R}^2 \quad (3.6.9)$$




Figure 3.6.5: $F(x, y)$ is the total probability below and to the left of (x, y) .

In the graph above, the light shading is intended to suggest a continuous distribution of probability, while the darker dots represent points of positive probability. Thus, $F(x, y)$ is the total probability mass below and to the left (that is, southwest) of the point (x, y) . As in the single variable case, the distribution function of (X, Y) completely determines the distribution of (X, Y) .

Suppose that $a, b, c, d \in \mathbb{R}$ with $a < b$ and $c < d$. Then

$$\mathbb{P}(a < X \leq b, c < Y \leq d) = F(b, d) - F(a, d) - F(b, c) + F(a, c) \quad (3.6.10)$$

Proof

Note that $\{X \leq a, Y \leq d\} \cup \{X \leq b, Y \leq c\} \cup \{a < X \leq b, c < Y \leq d\} = \{X \leq b, Y \leq d\}$. The intersection of the first two events is $\{X \leq a, Y \leq c\}$ while the first and third events and the second and third events are disjoint. Thus, from the inclusion-exclusion rule we have

$$F(a, d) + F(b, c) + \mathbb{P}(a < X \leq b, c < Y \leq d) - F(a, c) = F(b, d) \quad (3.6.11)$$

A probability distribution on \mathbb{R}^2 is completely determined by its values on rectangles of the form $(a, b] \times (c, d]$, so just as in the single variable case, it follows that the distribution function of (X, Y) completely determines the distribution of (X, Y) . See the advanced section on existence and uniqueness of positive measures in the chapter on Probability Measures for more details.

In the setting of the previous result, give the appropriate formula on the right for all possible combinations of weak and strong inequalities on the left.

The joint distribution function determines the individual (marginal) distribution functions.

Let F denote the distribution function of (X, Y) , and let G and H denote the distribution functions of X and Y , respectively. Then

1. $G(x) = F(x, \infty)$ for $x \in \mathbb{R}$

2. $H(y) = F(\infty, y)$ for $y \in \mathbb{R}$

Proof

These results follow from the continuity theorem for increasing events. For example, in (a)

$$\mathbb{P}(X \leq x) = \mathbb{P}(X \leq x, Y < \infty) = \lim_{y \rightarrow \infty} \mathbb{P}(X \leq x, Y \leq y) = \lim_{y \rightarrow \infty} F(x, y) \quad (3.6.12)$$

On the other hand, we cannot recover the distribution function of (X, Y) from the individual distribution functions, except when the variables are independent.

Random variables X and Y are independent if and only if

$$F(x, y) = G(x)H(y), \quad (x, y) \in \mathbb{R}^2 \quad (3.6.13)$$

Proof

If X and Y are independent then $F(x, y) = \mathbb{P}(X \leq x, Y \leq y) = \mathbb{P}(X \leq x)\mathbb{P}(Y \leq y) = G(x)H(y)$ for $(x, y) \in \mathbb{R}^2$. Conversely, suppose $F(x, y) = G(x)H(y)$ for $(x, y) \in \mathbb{R}^2$. If $a, b, c, d \in \mathbb{R}$ with $a < b$ and $c < d$ then from (15),

$$\mathbb{P}(a < X \leq b, c < Y \leq d) = G(b)H(d) - G(a)H(d) - G(b)H(c) + G(a)H(c) \quad (3.6.14)$$

$$= [G(b) - G(a)][H(d) - H(c)] = \mathbb{P}(a < X \leq b)\mathbb{P}(c < Y \leq d) \quad (3.6.15)$$

so it follows that X and Y are independent. (Recall again that a probability distribution on \mathbb{R}^2 is completely determined by its values on rectangles.)

All of the results of this subsection generalize in a straightforward way to n -dimensional random vectors. Only the notation is more complicated.

The Empirical Distribution Function

Suppose now that X is a real-valued random variable for a basic random experiment and that we repeat the experiment n times independently. This generates (for the new compound experiment) a sequence of independent variables (X_1, X_2, \dots, X_n) each with the same distribution as X . In statistical terms, this sequence is a *random sample* of size n from the distribution of X . In statistical inference, the observed values (x_1, x_2, \dots, x_n) of the random sample form our *data*.

The *empirical distribution function*, based on the data (x_1, x_2, \dots, x_n) , is defined by

$$F_n(x) = \frac{1}{n} \# \{i \in \{1, 2, \dots, n\} : x_i \leq x\} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i \leq x), \quad x \in \mathbb{R} \quad (3.6.16)$$

Thus, $F_n(x)$ gives the proportion of values in the data set that are less than or equal to x . The function F_n is a statistical estimator of F , based on the given data set. This concept is explored in more detail in the section on the sample mean in the chapter on Random Samples. In addition, the empirical distribution function is related to the Brownian bridge stochastic process which is studied in the chapter on Brownian motion.

Quantile Functions

Definitions

Suppose again that X is a real-valued random variable with distribution function F .

For $p \in (0, 1)$, a value of x such that $F(x^-) = \mathbb{P}(X < x) \leq p$ and $F(x) = \mathbb{P}(X \leq x) \geq p$ is called a *quantile* of order p for the distribution.

Roughly speaking, a quantile of order p is a value where the graph of the distribution function crosses (or jumps over) p . For example, in the picture below, a is the unique quantile of order p and b is the unique quantile of order q . On the other hand, the quantiles of order r form the interval $[c, d]$, and moreover, d is a quantile for all orders in the interval $[r, s]$. Note also that if X has a continuous distribution (so that F is continuous) and x is a quantile of order $p \in (0, 1)$, then $F(x) = p$.




Figure 3.6.6: Quantiles of various orders

Note that there is an inverse relation of sorts between the quantiles and the cumulative distribution values, but the relation is more complicated than that of a function and its ordinary inverse function, because the distribution function is not one-to-one in general. For many purposes, it is helpful to select a specific quantile for each order; to do this requires defining a *generalized inverse* of the distribution function F .

The *quantile function* F^{-1} of X is defined by

$$F^{-1}(p) = \min\{x \in \mathbb{R} : F(x) \geq p\}, \quad p \in (0, 1) \quad (3.6.17)$$

F^{-1} is well defined

Since F is right continuous and increasing, $\{x \in \mathbb{R} : F(x) \geq p\}$ is an interval of the form $[a, \infty)$. Thus, the minimum of the set is a .

Note that if F strictly increases from 0 to 1 on an interval S (so that the underlying distribution is continuous and is supported on S), then F^{-1} is the ordinary inverse of F . We do not usually define the quantile function at the endpoints 0 and 1. If we did, note that $F^{-1}(0)$ would always be $-\infty$.

Properties

The following exercise justifies the name: $F^{-1}(p)$ is the minimum of the quantiles of order p .

Let $p \in (0, 1)$.

1. $F^{-1}(p)$ is a quantile of order p .
2. If x is a quantile of order p then $F^{-1}(p) \leq x$.

Proof

Let $y = F^{-1}(p)$.

1. Note that $F(y) \geq p$ by definition, and if $x < y$ then $F(x) < p$. Hence $F(y^-) \leq p$. Therefore y is a quantile of order p .
2. Suppose that x is a quantile of order p . Then $F(x) \geq p$ so by definition, $y \leq x$.

Other basic properties of the quantile function are given in the following theorem.

F^{-1} satisfies the following properties:

1. F^{-1} is increasing on $(0, 1)$.
2. $F^{-1}[F(x)] \leq x$ for any $x \in \mathbb{R}$ with $F(x) < 1$.
3. $F[F^{-1}(p)] \geq p$ for any $p \in (0, 1)$.
4. $F^{-1}(p^-) = F^{-1}(p)$ for $p \in (0, 1)$. Thus F^{-1} is continuous from the left.
5. $F^{-1}(p^+) = \inf\{x \in \mathbb{R} : F(x) > p\}$ for $p \in (0, 1)$. Thus F^{-1} has limits from the right.

Proof

1. Note that if $p, q \in (0, 1)$ with $p \leq q$, then $\{x \in \mathbb{R} : F(x) \geq q\} \subseteq \{x \in \mathbb{R} : F(x) \geq p\}$.
2. This follows from the definition: $F^{-1}[F(x)]$ is the smallest $y \in \mathbb{R}$ with $F(y) \geq F(x)$.

3. This also follows from the definition: $F^{-1}(p)$ is a value $y \in \mathbb{R}$ satisfying $F(y) \geq p$.
4. This follows from the fact that F is continuous from the right
5. This follows from the fact that F has limits from the left.

As always, the inverse of a function is obtained essentially by reversing the roles of independent and dependent variables. In the graphs below, note that jumps of F become flat portions of F^{-1} while flat portions of F become jumps of F^{-1} . For $p \in (0, 1)$, the set of quantiles of order p is the closed, bounded interval $[F^{-1}(p), F^{-1}(p^+)]$. Thus, $F^{-1}(p)$ is the smallest quantile of order p , as we noted earlier, while $F^{-1}(p^+)$ is the largest quantile of order p .




Figure 3.6.7: Graph of the distribution function




Figure 3.6.8: Graph of the quantile function

The following basic property will be useful in simulating random variables, a topic explored in the section on transformations of random variables.

For $x \in \mathbb{R}$ and $p \in (0, 1)$, $F^{-1}(p) \leq x$ if and only if $p \leq F(x)$.

Proof

Suppose that $F^{-1}(p) \leq x$. Then, since F is increasing, $F[F^{-1}(p)] \leq F(x)$. But $p \leq F[F^{-1}(p)]$ by part (c) of the previous result, so $p \leq F(x)$. Conversely, suppose that $p \leq F(x)$. Then, since F^{-1} is increasing, $F^{-1}(p) \leq F^{-1}[F(x)]$. But $F^{-1}[F(x)] \leq x$ by part (b) of the previous result, so $F^{-1}(p) \leq x$.

Special Quantiles

Certain quantiles are important enough to deserve special names.

Suppose that X is a real-valued random variable.

1. A quantile of order $\frac{1}{4}$ is a *first quartile* of the distribution.
2. A quantile of order $\frac{1}{2}$ is a *median* or *second quartile* of the distribution.
3. A quantile of order $\frac{3}{4}$ is a *third quartile* of the distribution.

When there is only one median, it is frequently used as a measure of the *center* of the distribution, since it divides the set of values of X in half, by probability. More generally, the quartiles can be used to divide the set of values into fourths, by probability.

Assuming uniqueness, let q_1 , q_2 , and q_3 denote the first, second, and third quartiles of X , respectively, and let $a = F^{-1}(0^+)$ and $b = F^{-1}(1)$.

1. The *interquartile range* is defined to be $q_3 - q_1$.
2. The five parameters (a, q_1, q_2, q_3, b) are referred to as the *five number summary* of the distribution.

Note that the interval $[q_1, q_3]$ roughly gives the middle half of the distribution, so the interquartile range, the length of the interval, is a natural measure of the dispersion of the distribution about the median. Note also that a and b are essentially the minimum and maximum values of X , respectively, although of course, it's possible that $a = -\infty$ or $b = \infty$ (or both). Collectively, the five parameters give a great deal of information about the distribution in terms of the center, spread, and skewness. Graphically, the five numbers are often displayed as a *boxplot* or *box and whisker plot*, which consists of a line extending from the minimum value a to the maximum value b , with a rectangular box from q_1 to q_3 , and “whiskers” at a , the median q_2 , and b . Roughly speaking, the five numbers separate the set of values of X into 4 intervals of approximate probability $\frac{1}{4}$ each.




Figure 3.6.9: The probability density function and boxplot for a continuous distribution

Suppose that X has a continuous distribution that is symmetric about a point $a \in \mathbb{R}$. If $a + t$ is a quantile of order $p \in (0, 1)$ then $a - t$ is a quantile of order $1 - p$.

Proof

Note that this is the quantile function version of [symmetry result](#) for the distribution function. If $a + t$ is a quantile of order p then (since X has a continuous distribution) $F(a + t) = p$. But then $F(a - t) = 1 - F(a + t) = 1 - p$ so $a - t$ is a quantile of order $1 - p$.

Examples and Applications

Distributions of Different Types

Let F be the function defined by

$$F(x) = \begin{cases} 0, & x < 1 \\ \frac{1}{10}, & 1 \leq x < \frac{3}{2} \\ \frac{3}{10}, & \frac{3}{2} \leq x < 2 \\ \frac{6}{10}, & 2 \leq x < \frac{5}{2} \\ \frac{9}{10}, & \frac{5}{2} \leq x < 3 \\ 1, & x \geq 3; \end{cases} \quad (3.6.18)$$

1. Sketch the graph of F and show that F is the distribution function for a discrete distribution.
2. Find the corresponding probability density function f and sketch the graph.
3. Find $\mathbb{P}(2 \leq X < 3)$ where X has this distribution.
4. Find the quantile function and sketch the graph.
5. Find the five number summary and sketch the boxplot.

Answer

1. Note that F increases from 0 to 1, is a step function, and is right continuous.

2. $f(x) = \begin{cases} \frac{1}{10}, & x = 1 \\ \frac{1}{5}, & x = \frac{3}{2} \\ \frac{3}{10}, & x = 2 \\ \frac{3}{10}, & x = \frac{5}{2} \\ \frac{1}{10}, & x = 3 \end{cases}$
3. $\mathbb{P}(2 \leq X < 3) = \frac{3}{5}$
4. $F^{-1}(p) = \begin{cases} 1, & 0 < p \leq \frac{1}{10} \\ \frac{3}{2}, & \frac{1}{10} < p \leq \frac{3}{10} \\ 2, & \frac{3}{10} < p \leq \frac{6}{10} \\ \frac{5}{2}, & \frac{6}{10} < p \leq \frac{9}{10} \\ 3, & \frac{9}{10} < p \leq 1 \end{cases}$
5. $(1, \frac{3}{2}, 2, \frac{5}{2}, 3)$

Let F be the function defined by

$$F(x) = \begin{cases} 0, & x < 0 \\ \frac{x}{x+1}, & x \geq 0 \end{cases} \quad (3.6.19)$$

1. Sketch the graph of F and show that F is the distribution function for a continuous distribution.
2. Find the corresponding probability density function f and sketch the graph.
3. Find $\mathbb{P}(2 \leq X < 3)$ where X has this distribution.
4. Find the quantile function and sketch the graph.
5. Find the five number summary and sketch the boxplot.

Answer

1. Note that F is continuous and increases from 0 to 1.
2. $f(x) = \frac{1}{(x+1)^2}, \quad x > 0$
3. $\mathbb{P}(2 \leq X < 3) = \frac{1}{12}$
4. $F^{-1}(p) = \frac{p}{1-p}, \quad 0 < p < 1$
5. $(0, \frac{1}{3}, 1, 3, \infty)$

The expression $\frac{p}{1-p}$ that occurs in the quantile function in the last exercise is known as the *odds ratio* associated with p , particularly in the context of gambling.

Let F be the function defined by

$$F(x) = \begin{cases} 0, & x < 0 \\ \frac{1}{4}x, & 0 \leq x < 1 \\ \frac{1}{3} + \frac{1}{4}(x-1)^2, & 1 \leq x < 2 \\ \frac{2}{3} + \frac{1}{4}(x-2)^3, & 2 \leq x < 3 \\ 1, & x \geq 3 \end{cases} \quad (3.6.20)$$

1. Sketch the graph of F and show that F is the distribution function of a mixed distribution.
2. Find the partial probability density function of the discrete part and sketch the graph.
3. Find the partial probability density function of the continuous part and sketch the graph.
4. Find $\mathbb{P}(2 \leq X < 3)$ where X has this distribution.
5. Find the quantile function and sketch the graph.
6. Find the five number summary and sketch the boxplot.

Answer

1. Note that F is piece-wise continuous, increases from 0 to 1, and is right continuous.

2. $g(1) = g(2) = g(3) = \frac{1}{12}$
3. $h(x) = \begin{cases} \frac{1}{4}, & 0 < x < 1 \\ \frac{1}{2}(x-1), & 1 < x < 2 \\ \frac{3}{4}(x-2)^2, & 2 < x < 3 \end{cases}$
4. $\mathbb{P}(2 \leq X < 3) = \frac{1}{3}$
5. $F^{-1}(p) = \begin{cases} 4p, & 0 < p \leq \frac{1}{4} \\ 1, & \frac{1}{4} < p \leq \frac{1}{3} \\ 1 + \sqrt{4(p - \frac{1}{3})}, & \frac{1}{3} < p \leq \frac{7}{12} \\ 2, & \frac{7}{12} < p \leq \frac{2}{3} \\ 2 + \sqrt[3]{4(p - \frac{2}{3})}, & \frac{2}{3} < p \leq \frac{11}{12} \\ 3, & \frac{11}{12} < p \leq 1 \end{cases}$
6. $(0, 1, 1 + \sqrt{\frac{2}{3}}, 2 + \sqrt[3]{\frac{1}{3}}, 3)$

The Uniform Distribution

Suppose that X has probability density function $f(x) = \frac{1}{b-a}$ for $x \in [a, b]$, where $a, b \in \mathbb{R}$ and $a < b$.

1. Find the distribution function and sketch the graph.
2. Find the quantile function and sketch the graph.
3. Compute the five-number summary.
4. Sketch the graph of the probability density function with the boxplot on the horizontal axis.

Answer

1. $F(x) = \frac{x-a}{b-a}, \quad a \leq x < b$
2. $F^{-1}(p) = a + (b-a)p, \quad 0 \leq p \leq 1$
3. $(a, \frac{3a+b}{4}, \frac{a+b}{2}, \frac{a+3b}{4}, b)$

The distribution in the last exercise is the *uniform distribution* on the interval $[a, b]$. The left endpoint a is the *location parameter* and the length of the interval $w = b - a$ is the *scale parameter*. The uniform distribution models a point chose “at random” from the interval, and is studied in more detail in the chapter on Special Distributions.

In the special distribution calculator, select the continuous uniform distribution. Vary the location and scale parameters and note the shape of the probability density function and the distribution function.

The Exponential Distribution

Suppose that T has probability density function $f(t) = re^{-rt}$ for $0 \leq t < \infty$, where $r > 0$ is a parameter.

1. Find the distribution function and sketch the graph.
2. Find the reliability function and sketch the graph.
3. Find the failure rate function and sketch the graph.
4. Find the quantile function and sketch the graph.
5. Compute the five-number summary.
6. Sketch the graph of the probability density function with the boxplot on the horizontal axis.

Answer

1. $F(t) = 1 - e^{-rt}, \quad 0 \leq t < \infty$
2. $F^c(t) = e^{-rt}, \quad 0 \leq t < \infty$
3. $h(t) = r, \quad 0 \leq t < \infty$
4. $F^{-1}(p) = -\frac{1}{r} \ln(1-p), \quad 0 \leq p < 1$
5. $(0, \frac{1}{r}[\ln 4 - \ln 3], \frac{1}{r} \ln 2, \frac{1}{r} \ln 4, \infty)$

The distribution in the last exercise is the *exponential distribution* with rate parameter r . Note that this distribution is characterized by the fact that it has *constant failure rate* (and this is the reason for referring to r as the rate parameter). The reciprocal of the rate parameter is the *scale parameter*. The exponential distribution is used to model failure times and other random times under certain conditions, and is studied in detail in the chapter on The Poisson Process.

In the special distribution calculator, select the exponential distribution. Vary the scale parameter b and note the shape of the probability density function and the distribution function.

The Pareto Distribution

Suppose that X has probability density function $f(x) = \frac{a}{x^{a+1}}$ for $1 \leq x < \infty$ where $a > 0$ is a parameter.

1. Find the distribution function.
2. Find the reliability function.
3. Find the failure rate function.
4. Find the quantile function.
5. Compute the five-number summary.
6. In the case $a = 2$, sketch the graph of the probability density function with the boxplot on the horizontal axis.

Answer

1. $F(x) = 1 - \frac{1}{x^a}, \quad 1 \leq x < \infty$
2. $F^c(x) = \frac{1}{x^a}, \quad 1 \leq x < \infty$
3. $h(x) = \frac{a}{x}, \quad 1 \leq x < \infty$
4. $F^{-1}(p) = (1 - p)^{-1/a}, \quad 0 \leq p < 1$
5. $\left(1, \left(\frac{3}{4}\right)^{-1/a}, \left(\frac{1}{2}\right)^{-1/a}, \left(\frac{1}{4}\right)^{-1/a}, \infty\right)$

The distribution in the last exercise is the *Pareto distribution* with shape parameter a , named after Vilfredo Pareto. The Pareto distribution is a heavy-tailed distribution that is sometimes used to model income and certain other economic variables. It is studied in detail in the chapter on Special Distributions.

In the special distribution calculator, select the Pareto distribution. Keep the default value for the scale parameter, but vary the shape parameter and note the shape of the density function and the distribution function.

The Cauchy Distribution

Suppose that X has probability density function $f(x) = \frac{1}{\pi(1+x^2)}$ for $x \in \mathbb{R}$.

1. Find the distribution function and sketch the graph.
2. Find the quantile function and sketch the graph.
3. Compute the five-number summary and the interquartile range.
4. Sketch the graph of the probability density function with the boxplot on the horizontal axis.

Answer

1. $F(x) = \frac{1}{2} + \frac{1}{\pi} \arctan x, \quad x \in \mathbb{R}$
2. $F^{-1}(p) = \tan\left[\pi\left(p - \frac{1}{2}\right)\right], \quad 0 < p < 1$
3. $(-\infty, -1, 0, 1, \infty), \text{IQR} = 2$

The distribution in the last exercise is the *Cauchy distribution*, named after Augustin Cauchy. The Cauchy distribution is studied in more generality in the chapter on Special Distributions.

In the special distribution calculator, select the Cauchy distribution and keep the default parameter values. Note the shape of the density function and the distribution function.

The Weibull Distribution

Let $h(t) = kt^{k-1}$ for $0 < t < \infty$ where $k > 0$ is a parameter.

1. Sketch the graph of h in the cases $0 < k < 1$, $k = 1$, $1 < k < 2$, $k = 2$, and $k > 2$.
2. Show that h is a failure rate function.
3. Find the reliability function and sketch the graph.
4. Find the distribution function and sketch the graph.
5. Find the probability density function and sketch the graph.
6. Find the quantile function and sketch the graph.
7. Compute the five-number summary.

Answer

1. h is decreasing and concave upward if $0 < k < 1$; $h = 1$ (constant) if $k = 1$; h is increasing and concave downward if $1 < k < 2$; $h(t) = t$ (linear) if $k = 2$; h is increasing and concave upward if $k > 2$;
2. $h(t) > 0$ for $0 < t < \infty$ and $\int_0^\infty h(t) dt = \infty$
3. $F^c(t) = \exp(-t^k)$, $0 \leq t < \infty$
4. $F(t) = 1 - \exp(-t^k)$, $0 \leq t < \infty$
5. $f(t) = kt^{k-1} \exp(-t^k)$, $0 \leq t < \infty$
6. $F^{-1}(p) = [-\ln(1-p)]^{1/k}$, $0 \leq p < 1$
7. $(0, [\ln 4 - \ln 3]^{1/k}, [\ln 2]^{1/k}, [\ln 4]^{1/k}, \infty)$

The distribution in the previous exercise is the *Weibull distributions* with shape parameter k , named after Walodi Weibull. The Weibull distribution is studied in detail in the chapter on Special Distributions. Since this family includes increasing, decreasing, and constant failure rates, it is widely used to model the lifetimes of various types of devices.

In the special distribution calculator, select the Weibull distribution. Keep the default scale parameter, but vary the shape parameter and note the shape of the density function and the distribution function.

Beta Distributions

Suppose that X has probability density function $f(x) = 12x^2(1-x)$ for $0 \leq x \leq 1$.

1. Find the distribution function of X and sketch the graph.
2. Find $\mathbb{P}(\frac{1}{4} \leq X \leq \frac{1}{2})$.
3. Compute the five number summary and the interquartile range. You will have to approximate the quantiles.
4. Sketch the graph of the density function with the boxplot on the horizontal axis.

Answer

1. $F(x) = 4x^3 - 3x^4$, $0 \leq x \leq 1$
2. $\mathbb{P}(\frac{1}{4} \leq X \leq \frac{1}{2}) = \frac{67}{256}$
3. $(0, 0.4563, 0.6413, 0.7570, 1)$, IQR = 0.3007

Suppose that X has probability density function $f(x) = \frac{1}{\pi\sqrt{x(1-x)}}$ for $0 < x < 1$.

1. Find the distribution function of X and sketch the graph.
2. Compute $\mathbb{P}(\frac{1}{3} \leq X \leq \frac{2}{3})$.
3. Find the quantile function and sketch the graph.
4. Compute the five number summary and the interquartile range.
5. Sketch the graph of the probability density function with the boxplot on the horizontal axis.

Answer

1. $F(x) = \frac{2}{\pi} \arcsin(\sqrt{x})$, $0 \leq x \leq 1$
2. $\mathbb{P}(\frac{1}{3} \leq X \leq \frac{2}{3}) = 0.2163$
3. $F^{-1}(p) = \sin^2(\frac{\pi p}{2})$, $0 < p < 1$

$$4. \left(0, \frac{1}{2} - \frac{\sqrt{2}}{4}, \frac{1}{2}, \frac{1}{2} + \frac{\sqrt{2}}{4}, 1\right), \text{IQR} = \frac{\sqrt{2}}{2}$$

The distributions in the last two exercises are examples of *beta distributions*. The particular beta distribution in the last exercise is also known as the *arcsine distribution*; the distribution function explains the name. Beta distributions are used to model random proportions and probabilities, and certain other types of random variables, and are studied in detail in the chapter on Special Distributions.

In the special distribution calculator, select the beta distribution. For each of the following parameter values, note the location and shape of the density function and the distribution function.

1. $a = 3, b = 2$. This gives the [first beta distribution](#) above.
2. $a = b = \frac{1}{2}$. This gives the [arcsine distribution](#) above

Logistic Distribution

Let $F(x) = \frac{e^x}{1+e^x}$ for $x \in \mathbb{R}$.

1. Show that F is a distribution function for a continuous distribution, and sketch the graph.
2. Compute $\mathbb{P}(-1 \leq X \leq 1)$ where X is a random variable with distribution function F .
3. Find the quantile function and sketch the graph.
4. Compute the five-number summary and the interquartile range.
5. Find the probability density function and sketch the graph with the boxplot on the horizontal axis.

Answer

1. Note that F is continuous, and increases from 0 to 1.
2. $\mathbb{P}(-1 \leq X \leq 1) = 0.4621$
3. $F^{-1}(p) = \ln\left(\frac{p}{1-p}\right), \quad 0 < p < 1$
4. $(-\infty, -\ln 3, 0, \ln 3, \infty)$
5. $f(x) = \frac{e^x}{(1+e^x)^2}, \quad x \in \mathbb{R}$

The distribution in the last exercise is an *logistic distribution* and the quantile function is known as the *logit function*. The logistic distribution is studied in detail in the chapter on Special Distributions.

In the special distribution calculator, select the logistic distribution and keep the default parameter values. Note the shape of the probability density function and the distribution function.

Extreme Value Distribution

Let $F(x) = e^{-e^{-x}}$ for $x \in \mathbb{R}$.

1. Show that F is a distribution function for a continuous distribution, and sketch the graph.
2. Compute $\mathbb{P}(-1 \leq X \leq 1)$ where X is a random variable with distribution function F .
3. Find the quantile function and sketch the graph.
4. Compute the five-number summary.
5. Find the probability density function and sketch the graph with the boxplot on the horizontal axis.

Answer

1. Note that F is continuous, and increases from 0 to 1.
2. $\mathbb{P}(-1 \leq X \leq 1) = 0.6262$
3. $F^{-1}(p) = -\ln(-\ln p), \quad 0 < p < 1$
4. $(-\infty, -\ln(\ln 4), -\ln(\ln 2), -\ln(\ln 4 - \ln 3), \infty)$
5. $f(x) = e^{-e^{-x}} e^{-x}, \quad x \in \mathbb{R}$

The distribution in the last exercise is the *type 1 extreme value distribution*, also known as the *Gumbel distribution* in honor of Emil Gumbel. Extreme value distributions are studied in detail in the chapter on Special Distributions.

In the special distribution calculator, select the extreme value distribution and keep the default parameter values. Note the shape and location of the probability density function and the distribution function.

The Standard Normal Distribution

Recall that the *standard normal distribution* has probability density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (3.6.21)$$

This distribution models physical measurements of all sorts subject to small, random errors, and is one of the most important distributions in probability. The normal distribution is studied in more detail in the chapter on Special Distributions. The distribution function Φ , of course, can be expressed as

$$\Phi(z) = \int_{-\infty}^z \phi(x) dx, \quad z \in \mathbb{R} \quad (3.6.22)$$

but Φ and the quantile function Φ^{-1} cannot be expressed, in closed form, in terms of elementary functions. Because of the importance of the normal distribution Φ and Φ^{-1} are themselves considered *special functions*, like \sin , \ln , and many others. Approximate values of these functions can be computed using most mathematical and statistical software packages. Because the distribution is symmetric about 0, $\Phi(-z) = 1 - \Phi(z)$ for $z \in \mathbb{R}$, and equivalently, $\Phi^{-1}(1-p) = -\Phi^{-1}(p)$. In particular, the median is 0.

Open the special distribution calculator and choose the normal distribution. Keep the default parameter values and select CDF view. Note the shape and location of the distribution/quantile function. Compute each of the following:

1. The first and third quartiles
2. The quantiles of order 0.9 and 0.1
3. The quantiles of order 0.95 and 0.05

Miscellaneous Exercises

Suppose that X has probability density function $f(x) = -\ln x$ for $0 < x \leq 1$.

1. Sketch the graph of f .
2. Find the distribution function F and sketch the graph.
3. Find $\mathbb{P}(\frac{1}{3} \leq X \leq \frac{1}{2})$.

Answer

2. $F(x) = x - x \ln x, \quad 0 < x < 1$
3. $\mathbb{P}(\frac{1}{3} \leq X \leq \frac{1}{2}) = \frac{1}{6} + \frac{1}{2} \ln 2 - \frac{1}{3} \ln 3$

Suppose that a pair of fair dice are rolled and the sequence of scores (X_1, X_2) is recorded.

1. Find the distribution function of $Y = X_1 + X_2$, the sum of the scores.
2. Find the distribution function of $V = \max\{X_1, X_2\}$, the maximum score.
3. Find the conditional distribution function of Y given $V = 5$.

Answer

The random variables are discrete, so the CDFs are step functions, with jumps at the values of the variables. The following tables give the values of the CDFs at the values of the random variables.

1.	y	2	3	4	5	6	7	8	9	10	11	12
	$\mathbb{P}(Y \leq y)$	$\frac{1}{36}$	$\frac{3}{36}$	$\frac{6}{36}$	$\frac{10}{36}$	$\frac{15}{36}$	$\frac{21}{36}$	$\frac{26}{36}$	$\frac{30}{36}$	$\frac{33}{36}$	$\frac{35}{36}$	1

2.	v	1	2	3	4	5	6
	$\mathbb{P}(V \leq v)$	$\frac{1}{36}$	$\frac{4}{36}$	$\frac{9}{36}$	$\frac{16}{36}$	$\frac{25}{36}$	1

3.	y	6	7	8	9	10
	$\mathbb{P}(Y \leq y \mid V = 5)$	$\frac{2}{9}$	$\frac{4}{9}$	$\frac{6}{9}$	$\frac{8}{9}$	1

Suppose that (X, Y) has probability density function $f(x, y) = x + y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

- Find the distribution function of X, Y .
- Compute $\mathbb{P}\left(\frac{1}{4} \leq X \leq \frac{1}{2}, \frac{1}{3} \leq Y \leq \frac{2}{3}\right)$.
- Find the distribution function of X .
- Find the distribution function of Y .
- Find the conditional distribution function of X given $Y = y$ for $0 < y < 1$.
- Find the conditional distribution function of Y given $X = x$ for $0 < x < 1$.
- Are X and Y independent?

Answer

- $F(x, y) = \frac{1}{2}(xy^2 + yx^2); \quad 0 < x < 1, 0 < y < 1$
- $\mathbb{P}\left(\frac{1}{4} \leq X \leq \frac{1}{2}, \frac{1}{3} \leq Y \leq \frac{2}{3}\right) = \frac{7}{96}$
- $G(x) = \frac{1}{2}(x + x^2), \quad 0 < x < 1$
- $H(y) = \frac{1}{2}(y + y^2), \quad 0 < y < 1$
- $G(x \mid y) = \frac{x^2/2 + xy}{y + 1/2}; \quad 0 < x < 1, 0 < y < 1$
- $H(y \mid x) = \frac{y^2/2 + xy}{x + 1/2}; \quad 0 < x < 1, 0 < y < 1$

Statistical Exercises

For the M&M data, compute the empirical distribution function of the total number of candies.

Answer

Let N denote the total number of candies. The empirical distribution function of N is a step function; the following table gives the values of the function at the jump points.

n	50	53	54	55	56	57	58	59	60	61
$\mathbb{P}(N \leq n)$	$\frac{1}{30}$	$\frac{2}{30}$	$\frac{3}{30}$	$\frac{7}{30}$	$\frac{11}{30}$	$\frac{14}{30}$	$\frac{23}{30}$	$\frac{26}{30}$	$\frac{28}{30}$	1

For the cicada data, let BL denotes body length and let G denote gender. Compute the empirical distribution function of the following variables:

- BL
- BL given $G = 1$ (male)
- BL given $G = 0$ (female).
- Do you believe that BL and G are independent?

For statistical versions of some of the topics in this section, see the chapter on Random Samples, and in particular, the sections on empirical distributions and order statistics.

This page titled [3.6: Distribution and Quantile Functions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.7: Transformations of Random Variables

This section studies how the distribution of a random variable changes when the variable is transformed in a deterministic way. If you are a new student of probability, you should skip the technical details.

Basic Theory

The Problem

As usual, we start with a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} is the collection of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . Suppose now that we have a random variable X for the experiment, taking values in a set S , and a function r from S into another set T . Then $Y = r(X)$ is a new random variable taking values in T . If the distribution of X is known, how do we find the distribution of Y ? This is a very basic and important question, and in a superficial sense, the solution is easy. But first recall that for $B \subseteq T$, $r^{-1}(B) = \{x \in S : r(x) \in B\}$ is the inverse image of B under r .

$$\mathbb{P}(Y \in B) = \mathbb{P}[X \in r^{-1}(B)] \quad \text{for } B \subseteq T.$$

Proof




Figure 3.7.1: A function $r : S \rightarrow T$. How is a probability distribution on S transformed by r to a distribution on T ?

However, frequently the distribution of X is known either through its distribution function F or its probability density function f , and we would similarly like to find the distribution function or probability density function of Y . This is a difficult problem in general, because as we will see, even simple transformations of variables with simple distributions can lead to variables with complex distributions. We will solve the problem in various special cases.

Transformed Variables with Discrete Distributions

When the transformed variable Y has a discrete distribution, the probability density function of Y can be computed using basic rules of probability.

Suppose that X has a discrete distribution on a countable set S , with probability density function f . Then Y has a discrete distribution with probability density function g given by

$$g(y) = \sum_{x \in r^{-1}\{y\}} f(x), \quad y \in T \quad (3.7.1)$$

Proof




Figure 3.7.2: A transformation of a discrete probability distribution.

Suppose that X has a continuous distribution on a subset $S \subseteq \mathbb{R}^n$ with probability density function f , and that T is countable. Then Y has a discrete distribution with probability density function g given by

$$g(y) = \int_{r^{-1}\{y\}} f(x) dx, \quad y \in T \quad (3.7.2)$$

Proof




Figure 3.7.3: A continuous distribution on S transformed by a discrete function $r : S \rightarrow T$

So the main problem is often computing the inverse images $r^{-1}\{y\}$ for $y \in T$. The formulas above in the [discrete](#) and [continuous](#) cases are not worth memorizing explicitly; it's usually better to just work each problem from scratch. The main step is to write the event $\{Y = y\}$ in terms of X , and then find the probability of this event using the probability density function of X .

Transformed Variables with Continuous Distributions

Suppose that X has a continuous distribution on a subset $S \subseteq \mathbb{R}^n$ and that $Y = r(X)$ has a continuous distributions on a subset $T \subseteq \mathbb{R}^m$. Suppose also that X has a known probability density function f . In many cases, the probability density function of Y can be found by first finding the distribution function of Y (using basic rules of probability) and then computing the appropriate derivatives of the distribution function. This general method is referred to, appropriately enough, as the *distribution function method*.

Suppose that Y is real valued. The distribution function G of Y is given by

$$G(y) = \int_{r^{-1}(-\infty, y]} f(x) dx, \quad y \in \mathbb{R} \quad (3.7.3)$$

Proof

Again, this follows from the definition of f as a PDF of X . For $y \in \mathbb{R}$,

$$G(y) = \mathbb{P}(Y \leq y) = \mathbb{P}[r(X) \in (-\infty, y]] = \mathbb{P}[X \in r^{-1}(-\infty, y]] = \int_{r^{-1}(-\infty, y]} f(x) dx \quad (3.7.4)$$

As in the discrete case, the formula in (4) not much help, and it's usually better to work each problem from scratch. The main step is to write the event $\{Y \leq y\}$ in terms of X , and then find the probability of this event using the probability density function of X .

The Change of Variables Formula

When the transformation r is one-to-one and smooth, there is a formula for the probability density function of Y directly in terms of the probability density function of X . This is known as the *change of variables* formula. Note that since r is one-to-one, it has an inverse function r^{-1} .

We will explore the one-dimensional case first, where the concepts and formulas are simplest. Thus, suppose that random variable X has a continuous distribution on an interval $S \subseteq \mathbb{R}$, with distribution function F and probability density function f . Suppose that $Y = r(X)$ where r is a differentiable function from S onto an interval T . As usual, we will let G denote the distribution function of Y and g the probability density function of Y .

Suppose that r is strictly increasing on S . For $y \in T$,

1. $G(y) = F[r^{-1}(y)]$
2. $g(y) = f[r^{-1}(y)] \frac{d}{dy} r^{-1}(y)$

Proof

1. $G(y) = \mathbb{P}(Y \leq y) = \mathbb{P}[r(X) \leq y] = \mathbb{P}[X \leq r^{-1}(y)] = F[r^{-1}(y)]$ for $y \in T$. Note that the inequality is preserved since r is increasing.
2. This follows from part (a) by taking derivatives with respect to y and using the chain rule. Recall that $F' = f$.

Suppose that r is strictly decreasing on S . For $y \in T$,

1. $G(y) = 1 - F[r^{-1}(y)]$
2. $g(y) = -f[r^{-1}(y)] \frac{d}{dy} r^{-1}(y)$

Proof

1. $G(y) = \mathbb{P}(Y \leq y) = \mathbb{P}[r(X) \leq y] = \mathbb{P}[X \geq r^{-1}(y)] = 1 - F[r^{-1}(y)]$ for $y \in T$. Note that the inequality is reversed since r is decreasing.
2. This follows from part (a) by taking derivatives with respect to y and using the chain rule. Recall again that $F' = f$.

The formulas for the probability density functions in the [increasing case](#) and the [decreasing case](#) can be combined:

If r is strictly increasing or strictly decreasing on S then the probability density function g of Y is given by

$$g(y) = f[r^{-1}(y)] \left| \frac{d}{dy} r^{-1}(y) \right| \quad (3.7.5)$$

Letting $x = r^{-1}(y)$, the change of variables formula can be written more compactly as

$$g(y) = f(x) \left| \frac{dx}{dy} \right| \quad (3.7.6)$$

Although succinct and easy to remember, the formula is a bit less clear. It must be understood that x on the right should be written in terms of y via the inverse function. The images below give a graphical interpretation of the formula in the two cases where r is increasing and where r is decreasing.




Figure 3.7.4: The change of variables theorems in the increasing and decreasing cases

The generalization of this result from \mathbb{R} to \mathbb{R}^n is basically a theorem in multivariate calculus. First we need some notation. Suppose that r is a one-to-one differentiable function from $S \subseteq \mathbb{R}^n$ onto $T \subseteq \mathbb{R}^n$. The *first derivative* of the inverse function $\mathbf{x} = r^{-1}(\mathbf{y})$ is the $n \times n$ matrix of first partial derivatives:

$$\left(\frac{d\mathbf{x}}{d\mathbf{y}} \right)_{ij} = \frac{\partial x_i}{\partial y_j} \quad (3.7.7)$$

The *Jacobian* (named in honor of Karl Gustav Jacobi) of the inverse function is the determinant of the first derivative matrix

$$\det \left(\frac{d\mathbf{x}}{d\mathbf{y}} \right) \quad (3.7.8)$$

With this compact notation, the multivariate change of variables formula is easy to state.

Suppose that \mathbf{X} is a random variable taking values in $S \subseteq \mathbb{R}^n$, and that \mathbf{X} has a continuous distribution with probability density function f . Suppose also $\mathbf{Y} = r(\mathbf{X})$ where r is a differentiable function from S onto $T \subseteq \mathbb{R}^n$. Then the probability density function g of \mathbf{Y} is given by

$$g(\mathbf{y}) = f(\mathbf{x}) \left| \det \left(\frac{d\mathbf{x}}{d\mathbf{y}} \right) \right|, \quad \mathbf{y} \in T \quad (3.7.9)$$

Proof

The result follows from the multivariate change of variables formula in calculus. If $B \subseteq T$ then

$$\mathbb{P}(\mathbf{Y} \in B) = \mathbb{P}[r(\mathbf{X}) \in B] = \mathbb{P}[\mathbf{X} \in r^{-1}(B)] = \int_{r^{-1}(B)} f(\mathbf{x}) d\mathbf{x} \quad (3.7.10)$$

Using the change of variables $\mathbf{x} = r^{-1}(\mathbf{y})$, $d\mathbf{x} = \left| \det \left(\frac{d\mathbf{x}}{d\mathbf{y}} \right) \right| d\mathbf{y}$ we have

$$\mathbb{P}(\mathbf{Y} \in B) = \int_B f[r^{-1}(\mathbf{y})] \left| \det \left(\frac{d\mathbf{x}}{d\mathbf{y}} \right) \right| d\mathbf{y} \quad (3.7.11)$$

So it follows that g defined in the theorem is a PDF for \mathbf{Y} .

The Jacobian is the infinitesimal scale factor that describes how n -dimensional volume changes under the transformation.




Figure 3.7.5: The multivariate change of variables theorem

Special Transformations

Linear Transformations

Linear transformations (or more technically *affine transformations*) are among the most common and important transformations. Moreover, this type of transformation leads to simple applications of the change of variable theorems. Suppose first that X is a random variable taking values in an interval $S \subseteq \mathbb{R}$ and that X has a continuous distribution on S with probability density function f . Let $Y = a + bX$ where $a \in \mathbb{R}$ and $b \in \mathbb{R} \setminus \{0\}$. Note that Y takes values in $T = \{y = a + bx : x \in S\}$, which is also an interval.

Y has probability density function g given by

$$g(y) = \frac{1}{|b|} f\left(\frac{y-a}{b}\right), \quad y \in T \quad (3.7.12)$$

Proof

The transformation is $y = a + bx$. Hence the inverse transformation is $x = (y - a)/b$ and $dx/dy = 1/b$. The result now follows from the [change of variables theorem](#).

When $b > 0$ (which is often the case in applications), this transformation is known as a *location-scale transformation*; a is the *location parameter* and b is the *scale parameter*. Scale transformations arise naturally when physical units are changed (from feet to meters, for example). Location transformations arise naturally when the physical reference point is changed (measuring time relative to 9:00 AM as opposed to 8:00 AM, for example). The change of temperature measurement from Fahrenheit to Celsius is a location and scale transformation. Location-scale transformations are studied in more detail in the chapter on Special Distributions.

The multivariate version of this result has a simple and elegant form when the linear transformation is expressed in matrix-vector form. Thus suppose that \mathbf{X} is a random variable taking values in $S \subseteq \mathbb{R}^n$ and that \mathbf{X} has a continuous distribution on S with probability density function f . Let $\mathbf{Y} = \mathbf{a} + \mathbf{B}\mathbf{X}$ where $\mathbf{a} \in \mathbb{R}^n$ and \mathbf{B} is an invertible $n \times n$ matrix. Note that \mathbf{Y} takes values in $T = \{\mathbf{a} + \mathbf{B}\mathbf{x} : \mathbf{x} \in S\} \subseteq \mathbb{R}^n$.

\mathbf{Y} has probability density function g given by

$$g(\mathbf{y}) = \frac{1}{|\det(\mathbf{B})|} f[\mathbf{B}^{-1}(\mathbf{y} - \mathbf{a})], \quad \mathbf{y} \in T \quad (3.7.13)$$

Proof

The transformation $\mathbf{y} = \mathbf{a} + \mathbf{B}\mathbf{x}$ maps \mathbb{R}^n one-to-one and onto \mathbb{R}^n . The inverse transformation is $\mathbf{x} = \mathbf{B}^{-1}(\mathbf{y} - \mathbf{a})$. The Jacobian of the inverse transformation is the constant function $\det(\mathbf{B}^{-1}) = 1/\det(\mathbf{B})$. The result now follows from the [multivariate change of variables theorem](#).

Sums and Convolution

Simple addition of random variables is perhaps the most important of all transformations. Suppose that X and Y are random variables on a probability space, taking values in $R \subseteq \mathbb{R}$ and $S \subseteq \mathbb{R}$, respectively, so that (X, Y) takes values in a subset of $R \times S$. Our goal is to find the distribution of $Z = X + Y$. Note that Z takes values in $T = \{z \in \mathbb{R} : z = x + y \text{ for some } x \in R, y \in S\}$. For $z \in T$, let $D_z = \{x \in R : z - x \in S\}$.

Suppose that (X, Y) probability density function f .

1. If (X, Y) has a discrete distribution then $Z = X + Y$ has a discrete distribution with probability density function u given by

$$u(z) = \sum_{x \in D_z} f(x, z - x), \quad z \in T \quad (3.7.14)$$

2. If (X, Y) has a continuous distribution then $Z = X + Y$ has a continuous distribution with probability density function u given by

$$u(z) = \int_{D_z} f(x, z - x) dx, \quad z \in T \quad (3.7.15)$$

Proof

1. $\mathbb{P}(Z = z) = \mathbb{P}(X = x, Y = z - x \text{ for some } x \in D_z) = \sum_{x \in D_z} f(x, z - x)$
2. For $A \subseteq T$, let $C = \{(u, v) \in R \times S : u + v \in A\}$. Then

$$\mathbb{P}(Z \in A) = \mathbb{P}(X + Y \in A) = \int_C f(u, v) d(u, v) \quad (3.7.16)$$

Now use the change of variables $x = u$, $z = u + v$. Then the inverse transformation is $u = x$, $v = z - x$ and the Jacobian is 1. Using the change of variables theorem (8) we have

$$\mathbb{P}(Z \in A) = \int_{D_z \times A} f(x, z - x) d(x, z) = \int_A \int_{D_z} f(x, z - x) dx dz \quad (3.7.17)$$

It follows that Z has probability density function $z \mapsto \int_{D_z} f(x, z - x) dx$.

In the discrete case, R and S are countable, so T is also countable as is D_z for each $z \in T$. In the continuous case, R and S are typically intervals, so T is also an interval as is D_z for $z \in T$. In both cases, determining D_z is often the most difficult step. By far the most important special case occurs when X and Y are independent.

Suppose that X and Y are independent and have probability density functions g and h respectively.

1. If X and Y have discrete distributions then $Z = X + Y$ has a discrete distribution with probability density function $g * h$ given by

$$(g * h)(z) = \sum_{x \in D_z} g(x)h(z - x), \quad z \in T \quad (3.7.18)$$

2. If X and Y have continuous distributions then $Z = X + Y$ has a continuous distribution with probability density function $g * h$ given by

$$(g * h)(z) = \int_{D_z} g(x)h(z - x) dx, \quad z \in T \quad (3.7.19)$$

In both cases, the probability density function $g * h$ is called the *convolution* of g and h .

Proof

Both results follows from the previous result above since $f(x, y) = g(x)h(y)$ is the probability density function of (X, Y) .

As before, determining this set D_z is often the most challenging step in finding the probability density function of Z . However, there is one case where the computations simplify significantly.

Suppose again that X and Y are independent random variables with probability density functions g and h , respectively.

1. In the discrete case, suppose X and Y take values in \mathbb{N} . Then Z has probability density function

$$(g * h)(z) = \sum_{x=0}^z g(x)h(z-x), \quad z \in \mathbb{N} \quad (3.7.20)$$

2. In the continuous case, suppose that X and Y take values in $[0, \infty)$. Then Z has probability density function

$$(g * h)(z) = \int_0^z g(x)h(z-x) dx, \quad z \in [0, \infty) \quad (3.7.21)$$

Proof

1. In this case, $D_z = \{0, 1, \dots, z\}$ for $z \in \mathbb{N}$.
2. In this case, $D_z = [0, z]$ for $z \in [0, \infty)$.

Convolution is a very important mathematical operation that occurs in areas of mathematics outside of probability, and so involving functions that are not necessarily probability density functions. The following result gives some simple properties of convolution.

Convolution (either discrete or continuous) satisfies the following properties, where f , g , and h are probability density functions of the same type.

1. $f * g = g * f$ (the *commutative property*)
2. $(f * g) * h = f * (g * h)$ (the *associative property*)

Proof

An analytic proof is possible, based on the definition of convolution, but a probabilistic proof, based on sums of independent random variables is much better. Thus, suppose that X , Y , and Z are independent random variables with PDFs f , g , and h , respectively.

1. The commutative property of convolution follows from the commutative property of addition: $X + Y = Y + X$.
2. The associative property of convolution follows from the associative property of addition: $(X + Y) + Z = X + (Y + Z)$.

Thus, in part (b) we can write $f * g * h$ without ambiguity. Of course, the constant 0 is the additive identity so $X + 0 = 0 + X = X$ for every random variable X . Also, a constant is independent of every other random variable. It follows that the probability density function δ of 0 (given by $\delta(0) = 1$) is the identity with respect to convolution (at least for discrete PDFs). That is, $f * \delta = \delta * f = f$. The next result is a simple corollary of the [convolution theorem](#), but is important enough to be highlighted.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent and identically distributed real-valued random variables, with common probability density function f . Then $Y_n = X_1 + X_2 + \dots + X_n$ has probability density function $f^{*n} = f * f * \dots * f$, the n -fold *convolution power* of f , for $n \in \mathbb{N}$.

In statistical terms, \mathbf{X} corresponds to *sampling* from the common distribution. By convention, $Y_0 = 0$, so naturally we take $f^{*0} = \delta$. When appropriately scaled and centered, the distribution of Y_n converges to the standard normal distribution as $n \rightarrow \infty$. The precise statement of this result is the *central limit theorem*, one of the fundamental theorems of probability. The central limit theorem is studied in detail in the chapter on Random Samples. Clearly convolution power satisfies the *law of exponents*: $f^{*n} * f^{*m} = f^{*(n+m)}$ for $m, n \in \mathbb{N}$.

Convolution can be generalized to sums of independent variables that are not of the same type, but this generalization is usually done in terms of distribution functions rather than probability density functions.

Products and Quotients

While not as important as sums, products and quotients of real-valued random variables also occur frequently. We will limit our discussion to continuous distributions.

Suppose that (X, Y) has a continuous distribution on \mathbb{R}^2 with probability density function f .

1. Random variable $V = XY$ has probability density function

$$v \mapsto \int_{-\infty}^{\infty} f(x, v/x) \frac{1}{|x|} dx \quad (3.7.22)$$

2. Random variable $W = Y/X$ has probability density function

$$w \mapsto \int_{-\infty}^{\infty} f(x, wx) |x| dx \quad (3.7.23)$$

Proof

We introduce the auxiliary variable $U = X$ so that we have bivariate transformations and can use our [change of variables formula](#).

1. We have the transformation $u = x, v = xy$ and so the inverse transformation is $x = u, y = v/u$. Hence

$$\frac{\partial(x, y)}{\partial(u, v)} = \begin{bmatrix} 1 & 0 \\ -v/u^2 & 1/u \end{bmatrix} \quad (3.7.24)$$

and so the Jacobian is $1/u$. Using the change of variables theorem, the joint PDF of (U, V) is $(u, v) \mapsto f(u, v/u) |1/u|$. Hence the PDF of V is

$$v \mapsto \int_{-\infty}^{\infty} f(u, v/u) \frac{1}{|u|} du \quad (3.7.25)$$

2. We have the transformation $u = x, w = y/x$ and so the inverse transformation is $x = u, y = uw$. Hence

$$\frac{\partial(x, y)}{\partial(u, w)} = \begin{bmatrix} 1 & 0 \\ w & u \end{bmatrix} \quad (3.7.26)$$

and so the Jacobian is u . Using the change of variables formula, the joint PDF of (U, W) is $(u, w) \mapsto f(u, uw) |u|$. Hence the PDF of W is

$$w \mapsto \int_{-\infty}^{\infty} f(u, uw) |u| du \quad (3.7.27)$$

If (X, Y) takes values in a subset $D \subseteq \mathbb{R}^2$, then for a given $v \in \mathbb{R}$, the integral in (a) is over $\{x \in \mathbb{R} : (x, v/x) \in D\}$, and for a given $w \in \mathbb{R}$, the integral in (b) is over $\{x \in \mathbb{R} : (x, wx) \in D\}$. As usual, the most important special case of this result is when X and Y are independent.

Suppose that X and Y are independent random variables with continuous distributions on \mathbb{R} having probability density functions g and h , respectively.

1. Random variable $V = XY$ has probability density function

$$v \mapsto \int_{-\infty}^{\infty} g(x) h(v/x) \frac{1}{|x|} dx \quad (3.7.28)$$

2. Random variable $W = Y/X$ has probability density function

$$w \mapsto \int_{-\infty}^{\infty} g(x) h(wx) |x| dx \quad (3.7.29)$$

Proof

These results follow immediately from the previous theorem, since $f(x, y) = g(x)h(y)$ for $(x, y) \in \mathbb{R}^2$.

If X takes values in $S \subseteq \mathbb{R}$ and Y takes values in $T \subseteq \mathbb{R}$, then for a given $v \in \mathbb{R}$, the integral in (a) is over $\{x \in S : v/x \in T\}$, and for a given $w \in \mathbb{R}$, the integral in (b) is over $\{x \in S : wx \in T\}$. As with convolution, determining the domain of integration is often the most challenging step.

Minimum and Maximum

Suppose that (X_1, X_2, \dots, X_n) is a sequence of independent real-valued random variables. The *minimum* and *maximum* transformations

$$U = \min\{X_1, X_2, \dots, X_n\}, \quad V = \max\{X_1, X_2, \dots, X_n\} \quad (3.7.30)$$

are very important in a number of applications. For example, recall that in the standard model of *structural reliability*, a system consists of n components that operate independently. Suppose that X_i represents the lifetime of component $i \in \{1, 2, \dots, n\}$. Then U is the lifetime of the *series system* which operates if and only if each component is operating. Similarly, V is the lifetime of the *parallel system* which operates if and only if at least one component is operating.

A particularly important special case occurs when the random variables are *identically distributed*, in addition to being independent. In this case, the sequence of variables is a *random sample* of size n from the common distribution. The minimum and maximum variables are the extreme examples of *order statistics*. Order statistics are studied in detail in the chapter on Random Samples.

Suppose that (X_1, X_2, \dots, X_n) is a sequence of independent real-valued random variables and that X_i has distribution function F_i for $i \in \{1, 2, \dots, n\}$.

1. $V = \max\{X_1, X_2, \dots, X_n\}$ has distribution function H given by $H(x) = F_1(x)F_2(x) \cdots F_n(x)$ for $x \in \mathbb{R}$.
2. $U = \min\{X_1, X_2, \dots, X_n\}$ has distribution function G given by $G(x) = 1 - [1 - F_1(x)][1 - F_2(x)] \cdots [1 - F_n(x)]$ for $x \in \mathbb{R}$.

Proof

1. Note that since V is the maximum of the variables, $\{V \leq x\} = \{X_1 \leq x, X_2 \leq x, \dots, X_n \leq x\}$. Hence by independence,

$$H(x) = \mathbb{P}(V \leq x) = \mathbb{P}(X_1 \leq x)\mathbb{P}(X_2 \leq x) \cdots \mathbb{P}(X_n \leq x) = F_1(x)F_2(x) \cdots F_n(x), \quad x \in \mathbb{R} \quad (3.7.31)$$

2. Note that since U is the minimum of the variables, $\{U > x\} = \{X_1 > x, X_2 > x, \dots, X_n > x\}$. Hence by independence,

$$\begin{aligned} G(x) &= \mathbb{P}(U > x) = 1 - \mathbb{P}(U \leq x) = 1 - \mathbb{P}(X_1 \leq x)\mathbb{P}(X_2 \leq x) \cdots \mathbb{P}(X_n \leq x) \\ &= 1 - [1 - F_1(x)][1 - F_2(x)] \cdots [1 - F_n(x)], \quad x \in \mathbb{R} \end{aligned}$$

From part (a), note that the product of n distribution functions is another distribution function. From part (b), the product of n right-tail distribution functions is a right-tail distribution function. In the reliability setting, where the random variables are nonnegative, the last statement means that the product of n reliability functions is another reliability function. If X_i has a continuous distribution with probability density function f_i for each $i \in \{1, 2, \dots, n\}$, then U and V also have continuous distributions, and their probability density functions can be obtained by differentiating the distribution functions in parts (a) and (b) of [last theorem](#). The computations are straightforward using the product rule for derivatives, but the results are a bit of a mess.

The formulas in [last theorem](#) are particularly nice when the random variables are identically distributed, in addition to being independent

Suppose that (X_1, X_2, \dots, X_n) is a sequence of independent real-valued random variables, with common distribution function F .

1. $V = \max\{X_1, X_2, \dots, X_n\}$ has distribution function H given by $H(x) = F^n(x)$ for $x \in \mathbb{R}$.
2. $U = \min\{X_1, X_2, \dots, X_n\}$ has distribution function G given by $G(x) = 1 - [1 - F(x)]^n$ for $x \in \mathbb{R}$.

In particular, it follows that a positive integer power of a distribution function is a distribution function. More generally, it's easy to see that every positive power of a distribution function is a distribution function. How could we construct a non-integer power of a distribution function in a probabilistic way?

Suppose that (X_1, X_2, \dots, X_n) is a sequence of independent real-valued random variables, with a common continuous distribution that has probability density function f .

1. $V = \max\{X_1, X_2, \dots, X_n\}$ has probability density function h given by $h(x) = nF^{n-1}(x)f(x)$ for $x \in \mathbb{R}$.
2. $U = \min\{X_1, X_2, \dots, X_n\}$ has probability density function g given by $g(x) = n[1 - F(x)]^{n-1}f(x)$ for $x \in \mathbb{R}$.

Coordinate Systems

For our next discussion, we will consider transformations that correspond to common distance-angle based coordinate systems—polar coordinates in the plane, and cylindrical and spherical coordinates in 3-dimensional space. First, for $(x, y) \in \mathbb{R}^2$, let (r, θ)

denote the standard *polar coordinates* corresponding to the Cartesian coordinates (x, y) , so that $r \in [0, \infty)$ is the *radial distance* and $\theta \in [0, 2\pi)$ is the *polar angle*.




Figure 3.7.6: Polar coordinates. Stover, Christopher and Weisstein, Eric W. "Polar Coordinates." From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/PolarCoordinates.html>

It's best to give the *inverse transformation*: $x = r \cos \theta$, $y = r \sin \theta$. As we all know from calculus, the Jacobian of the transformation is r . Hence the following result is an immediate consequence of our [change of variables theorem](#):

Suppose that (X, Y) has a continuous distribution on \mathbb{R}^2 with probability density function f , and that (R, Θ) are the polar coordinates of (X, Y) . Then (R, Θ) has probability density function g given by

$$g(r, \theta) = f(r \cos \theta, r \sin \theta)r, \quad (r, \theta) \in [0, \infty) \times [0, 2\pi) \quad (3.7.32)$$

Next, for $(x, y, z) \in \mathbb{R}^3$, let (r, θ, z) denote the standard *cylindrical coordinates*, so that (r, θ) are the standard polar coordinates of (x, y) as above, and coordinate z is left unchanged. Given our previous result, the one for cylindrical coordinates should come as no surprise.

Suppose that (X, Y, Z) has a continuous distribution on \mathbb{R}^3 with probability density function f , and that (R, Θ, Z) are the cylindrical coordinates of (X, Y, Z) . Then (R, Θ, Z) has probability density function g given by

$$g(r, \theta, z) = f(r \cos \theta, r \sin \theta, z)r, \quad (r, \theta, z) \in [0, \infty) \times [0, 2\pi) \times \mathbb{R} \quad (3.7.33)$$

Finally, for $(x, y, z) \in \mathbb{R}^3$, let (r, θ, ϕ) denote the standard *spherical coordinates* corresponding to the Cartesian coordinates (x, y, z) , so that $r \in [0, \infty)$ is the *radial distance*, $\theta \in [0, 2\pi)$ is the *azimuth angle*, and $\phi \in [0, \pi]$ is the *polar angle*. (In spite of our use of the word *standard*, different notations and conventions are used in different subjects.)




Figure 3.7.7: Spherical coordinates, By Dmcq—Own work, CC BY-SA 3.0, [Wikipedia](#)

Once again, it's best to give the *inverse transformation*: $x = r \sin \phi \cos \theta$, $y = r \sin \phi \sin \theta$, $z = r \cos \phi$. As we remember from calculus, the absolute value of the Jacobian is $r^2 \sin \phi$. Hence the following result is an immediate consequence of the change of variables theorem (8):

Suppose that (X, Y, Z) has a continuous distribution on \mathbb{R}^3 with probability density function f , and that (R, Θ, Φ) are the spherical coordinates of (X, Y, Z) . Then (R, Θ, Φ) has probability density function g given by

$$g(r, \theta, \phi) = f(r \sin \phi \cos \theta, r \sin \phi \sin \theta, r \cos \phi)r^2 \sin \phi, \quad (r, \theta, \phi) \in [0, \infty) \times [0, 2\pi) \times [0, \pi] \quad (3.7.34)$$

Sign and Absolute Value

Our next discussion concerns the sign and absolute value of a real-valued random variable.

Suppose that X has a continuous distribution on \mathbb{R} with distribution function F and probability density function f .

1. $|X|$ has distribution function G given by $G(y) = F(y) - F(-y)$ for $y \in [0, \infty)$.
2. $|X|$ has probability density function g given by $g(y) = f(y) + f(-y)$ for $y \in [0, \infty)$.

Proof

1. $\mathbb{P}(|X| \leq y) = \mathbb{P}(-y \leq X \leq y) = F(y) - F(-y)$ for $y \in [0, \infty)$.
2. This follows from part (a) by taking derivatives with respect to y .

Recall that the *sign* function on \mathbb{R} (not to be confused, of course, with the *sine* function) is defined as follows:

$$\text{sgn}(x) = \begin{cases} -1, & x < 0 \\ 0, & x = 0 \\ 1, & x > 0 \end{cases} \quad (3.7.35)$$

Suppose again that X has a continuous distribution on \mathbb{R} with distribution function F and probability density function f , and suppose in addition that the distribution of X is symmetric about 0. Then

1. $|X|$ has distribution function G given by $G(y) = 2F(y) - 1$ for $y \in [0, \infty)$.
2. $|X|$ has probability density function g given by $g(y) = 2f(y)$ for $y \in [0, \infty)$.
3. $\text{sgn}(X)$ is uniformly distributed on $\{-1, 1\}$.
4. $|X|$ and $\text{sgn}(X)$ are independent.

Proof

1. This follows from the previous theorem, since $F(-y) = 1 - F(y)$ for $y > 0$ by symmetry.
2. This follows from part (a) by taking derivatives.
3. Note that $\mathbb{P}[\text{sgn}(X) = 1] = \mathbb{P}(X > 0) = \frac{1}{2}$ and so $\mathbb{P}[\text{sgn}(X) = -1] = \frac{1}{2}$ also.
4. If $A \subseteq (0, \infty)$ then

$$\mathbb{P}[|X| \in A, \text{sgn}(X) = 1] = \mathbb{P}(X \in A) = \int_A f(x) dx = \frac{1}{2} \int_A 2f(x) dx = \mathbb{P}[\text{sgn}(X) = 1] \mathbb{P}(|X| \in A) \quad (3.7.36)$$

Examples and Applications

This subsection contains computational exercises, many of which involve special parametric families of distributions. It is always interesting when a random variable from one parametric family can be transformed into a variable from another family. It is also interesting when a parametric family is *closed* or *invariant* under some transformation on the variables in the family. Often, such properties are what make the parametric families *special* in the first place. Please note these properties when they occur.

Dice

Recall that a *standard die* is an ordinary 6-sided die, with faces labeled from 1 to 6 (usually in the form of dots). A *fair* die is one in which the faces are equally likely. An *ace-six flat die* is a standard die in which faces 1 and 6 occur with probability $\frac{1}{4}$ each and the other faces with probability $\frac{1}{8}$ each.

Suppose that two six-sided dice are rolled and the sequence of scores (X_1, X_2) is recorded. Find the probability density function of $Y = X_1 + X_2$, the sum of the scores, in each of the following cases:

1. Both dice are standard and fair.
2. Both dice are ace-six flat.
3. The first die is standard and fair, and the second is ace-six flat
4. The dice are both fair, but the first die has faces labeled 1, 2, 2, 3, 3, 4 and the second die has faces labeled 1, 3, 4, 5, 6, 8.

Answer

Let $Y = X_1 + X_2$ denote the sum of the scores.

1.	y	2	3	4	5	6	7	8	9	10	11	12
	$\mathbb{P}(Y = y)$	$\frac{1}{36}$	$\frac{2}{36}$	$\frac{3}{36}$	$\frac{4}{36}$	$\frac{5}{36}$	$\frac{6}{36}$	$\frac{5}{36}$	$\frac{4}{36}$	$\frac{3}{36}$	$\frac{2}{36}$	$\frac{1}{36}$

2.	y	2	3	4	5	6	7	8	9	10	11	12
	$\mathbb{P}(Y = y)$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{5}{64}$	$\frac{3}{32}$	$\frac{7}{64}$	$\frac{3}{16}$	$\frac{7}{64}$	$\frac{3}{32}$	$\frac{3}{32}$	$\frac{1}{16}$	$\frac{1}{16}$

3.	y	2	3	4	5	6	7	8	9	10	11	12
	$\mathbb{P}(Y = y)$	$\frac{2}{48}$	$\frac{3}{48}$	$\frac{4}{48}$	$\frac{5}{48}$	$\frac{6}{48}$	$\frac{8}{48}$	$\frac{6}{48}$	$\frac{5}{48}$	$\frac{4}{48}$	$\frac{3}{48}$	$\frac{2}{48}$

4. The distribution is the same as for two standard, fair dice in (a).

In the dice experiment, select two dice and select the sum random variable. Run the simulation 1000 times and compare the empirical density function to the probability density function for each of the following cases:

1. fair dice
2. ace-six flat dice

Suppose that n standard, fair dice are rolled. Find the probability density function of the following variables:

1. the minimum score
2. the maximum score.

Answer

Let U denote the minimum score and V the maximum score.

1. $f(u) = \left(1 - \frac{u-1}{6}\right)^n - \left(1 - \frac{u}{6}\right)^n, \quad u \in \{1, 2, 3, 4, 5, 6\}$
2. $g(v) = \left(\frac{v}{6}\right)^n - \left(\frac{v-1}{6}\right)^n, \quad v \in \{1, 2, 3, 4, 5, 6\}$

In the dice experiment, select fair dice and select each of the following random variables. Vary n with the scroll bar and note the shape of the density function. With $n = 4$, run the simulation 1000 times and note the agreement between the empirical density function and the probability density function.

1. minimum score
2. maximum score.

Uniform Distributions

Recall that for $n \in \mathbb{N}_+$, the standard measure of the size of a set $A \subseteq \mathbb{R}^n$ is

$$\lambda_n(A) = \int_A 1 \, dx \quad (3.7.37)$$

In particular, $\lambda_1(A)$ is the length of A for $A \subseteq \mathbb{R}$, $\lambda_2(A)$ is the area of A for $A \subseteq \mathbb{R}^2$, and $\lambda_3(A)$ is the volume of A for $A \subseteq \mathbb{R}^3$. See the technical details in (1) for more advanced information.

Now if $S \subseteq \mathbb{R}^n$ with $0 < \lambda_n(S) < \infty$, recall that the *uniform distribution* on S is the continuous distribution with constant probability density function f defined by $f(x) = 1/\lambda_n(S)$ for $x \in S$. Uniform distributions are studied in more detail in the chapter on Special Distributions.

Let $Y = X^2$. Find the probability density function of Y and sketch the graph in each of the following cases:

1. X is uniformly distributed on the interval $[0, 4]$.
2. X is uniformly distributed on the interval $[-2, 2]$.
3. X is uniformly distributed on the interval $[-1, 3]$.

Answer

1. $g(y) = \frac{1}{8\sqrt{y}}, \quad 0 < y < 16$
2. $g(y) = \frac{1}{4\sqrt{y}}, \quad 0 < y < 4$
3. $g(y) = \begin{cases} \frac{1}{4\sqrt{y}}, & 0 < y < 1 \\ \frac{1}{8\sqrt{y}}, & 1 < y < 9 \end{cases}$

Compare the distributions in the last exercise. In part (c), note that even a simple transformation of a simple distribution can produce a complicated distribution. In this particular case, the complexity is caused by the fact that $x \mapsto x^2$ is one-to-one on part of the domain $\{0\} \cup (1, 3]$ and two-to-one on the other part $[-1, 1] \setminus \{0\}$.

On the other hand, the uniform distribution is preserved under a linear transformation of the random variable.

Suppose that \mathbf{X} has the continuous uniform distribution on $S \subseteq \mathbb{R}^n$. Let $\mathbf{Y} = \mathbf{a} + \mathbf{B}\mathbf{X}$, where $\mathbf{a} \in \mathbb{R}^n$ and \mathbf{B} is an invertible $n \times n$ matrix. Then \mathbf{Y} is uniformly distributed on $T = \{\mathbf{a} + \mathbf{B}\mathbf{x} : \mathbf{x} \in S\}$.

Proof

This follows directly from the general result on linear transformations in (10). Note that the PDF g of \mathbf{Y} is constant on T .

For the following three exercises, recall that the *standard uniform distribution* is the uniform distribution on the interval $[0, 1]$.

Suppose that X and Y are independent and that each has the standard uniform distribution. Let $U = X + Y$, $V = X - Y$, $W = XY$, $Z = Y/X$. Find the probability density function of each of the follow:

1. (U, V)
2. U
3. V
4. W
5. Z

Answer

1. $g(u, v) = \frac{1}{2}$ for (u, v) in the square region $T \subset \mathbb{R}^2$ with vertices $\{(0, 0), (1, 1), (2, 0), (1, -1)\}$ So (U, V) is uniformly distributed on T .
2. $g_1(u) = \begin{cases} u, & 0 < u < 1 \\ 2 - u, & 1 < u < 2 \end{cases}$
3. $g_2(v) = \begin{cases} 1 - v, & 0 < v < 1 \\ 1 + v, & -1 < v < 0 \end{cases}$
4. $h_1(w) = -\ln w$ for $0 < w \leq 1$
5. $h_2(z) = \begin{cases} \frac{1}{2}, & 0 \leq z \leq 1 \\ \frac{1}{2z^2}, & 1 \leq z < \infty \end{cases}$

Suppose that X , Y , and Z are independent, and that each has the standard uniform distribution. Find the probability density function of $(U, V, W) = (X + Y, Y + Z, X + Z)$.

Answer

$g(u, v, w) = \frac{1}{2}$ for (u, v, w) in the rectangular region $T \subset \mathbb{R}^3$ with vertices $\{(0, 0, 0), (1, 0, 1), (1, 1, 0), (0, 1, 1), (2, 1, 1), (1, 1, 2), (1, 2, 1), (2, 2, 2)\}$ So (U, V, W) is uniformly distributed on T .

Suppose that (X_1, X_2, \dots, X_n) is a sequence of independent random variables, each with the standard uniform distribution. Find the distribution function and probability density function of the following variables.

1. $U = \min\{X_1, X_2, \dots, X_n\}$
2. $V = \max\{X_1, X_2, \dots, X_n\}$

Answer

1. $G(t) = 1 - (1 - t)^n$ and $g(t) = n(1 - t)^{n-1}$, both for $t \in [0, 1]$
2. $H(t) = t^n$ and $h(t) = nt^{n-1}$, both for $t \in [0, 1]$

Both distributions in the last exercise are *beta distributions*. More generally, all of the order statistics from a random sample of standard uniform variables have beta distributions, one of the reasons for the importance of this family of distributions. Beta distributions are studied in more detail in the chapter on Special Distributions.

In the order statistic experiment, select the uniform distribution.

1. Set $k = 1$ (this gives the minimum U). Vary n with the scroll bar and note the shape of the probability density function. With $n = 5$, run the simulation 1000 times and note the agreement between the empirical density function and the true probability density function.
2. Vary n with the scroll bar, set $k = n$ each time (this gives the maximum V), and note the shape of the probability density function. With $n = 5$ run the simulation 1000 times and compare the empirical density function and the probability density function.

Let f denote the probability density function of the standard uniform distribution.

1. Compute f^{*2}
2. Compute f^{*3}
3. Graph f , f^{*2} , and f^{*3} on the same set of axes.

Answer

$$1. f^{*2}(z) = \begin{cases} z, & 0 < z < 1 \\ 2 - z, & 1 < z < 2 \end{cases}$$

$$2. f^{*3}(z) = \begin{cases} \frac{1}{2}z^2, & 0 < z < 1 \\ 1 - \frac{1}{2}(z-1)^2 - \frac{1}{2}(2-z)^2, & 1 < z < 2 \\ \frac{1}{2}(3-z)^2, & 2 < z < 3 \end{cases}$$

In the last exercise, you can see the behavior predicted by the central limit theorem beginning to emerge. Recall that if (X_1, X_2, X_3) is a sequence of independent random variables, each with the standard uniform distribution, then f , f^{*2} , and f^{*3} are the probability density functions of X_1 , $X_1 + X_2$, and $X_1 + X_2 + X_3$, respectively. More generally, if (X_1, X_2, \dots, X_n) is a sequence of independent random variables, each with the standard uniform distribution, then the distribution of $\sum_{i=1}^n X_i$ (which has probability density function f^{*n}) is known as the *Irwin-Hall* distribution with parameter n . The Irwin-Hall distributions are studied in more detail in the chapter on Special Distributions.

Open the Special Distribution Simulator and select the Irwin-Hall distribution. Vary the parameter n from 1 to 3 and note the shape of the probability density function. (These are the density functions in the previous exercise). For each value of n , run the simulation 1000 times and compare the empirical density function and the probability density function.

Simulations

A remarkable fact is that the standard uniform distribution can be transformed into almost any other distribution on \mathbb{R} . This is particularly important for simulations, since many computer languages have an algorithm for generating *random numbers*, which are simulations of independent variables, each with the standard uniform distribution. Conversely, any continuous distribution supported on an interval of \mathbb{R} can be transformed into the standard uniform distribution.

Suppose first that F is a distribution function for a distribution on \mathbb{R} (which may be discrete, continuous, or mixed), and let F^{-1} denote the quantile function.

Suppose that U has the standard uniform distribution. Then $X = F^{-1}(U)$ has distribution function F .

Proof

The critical property satisfied by the quantile function (regardless of the type of distribution) is $F^{-1}(p) \leq x$ if and only if $p \leq F(x)$ for $p \in (0, 1)$ and $x \in \mathbb{R}$. Hence for $x \in \mathbb{R}$, $\mathbb{P}(X \leq x) = \mathbb{P}[F^{-1}(U) \leq x] = \mathbb{P}[U \leq F(x)] = F(x)$.

Assuming that we can compute F^{-1} , the previous exercise shows how we can simulate a distribution with distribution function F . To rephrase the result, we can simulate a variable with distribution function F by simply computing a *random quantile*. Most of the apps in this project use this method of simulation. The first image below shows the graph of the distribution function of a rather complicated mixed distribution, represented in blue on the horizontal axis. In the second image, note how the uniform distribution on $[0, 1]$, represented by the thick red line, is transformed, via the quantile function, into the given distribution.




Figure 3.7.8: The random quantile method of simulation

There is a partial converse to the previous result, for continuous distributions.

Suppose that X has a continuous distribution on an interval $S \subseteq \mathbb{R}$. Then $U = F(X)$ has the standard uniform distribution.

Proof

For $u \in (0, 1)$ recall that $F^{-1}(u)$ is a quantile of order u . Since X has a continuous distribution,

$$\mathbb{P}(U \geq u) = \mathbb{P}[F(X) \geq u] = \mathbb{P}[X \geq F^{-1}(u)] = 1 - F[F^{-1}(u)] = 1 - u \quad (3.7.38)$$

Hence U is uniformly distributed on $(0, 1)$.

Show how to simulate the uniform distribution on the interval $[a, b]$ with a random number. Using your calculator, simulate 5 values from the uniform distribution on the interval $[2, 10]$.

Answer

$X = a + U(b - a)$ where U is a random number.

Beta Distributions

Suppose that X has the probability density function f given by $f(x) = 3x^2$ for $0 \leq x \leq 1$. Find the probability density function of each of the following:

1. $U = X^2$
2. $V = \sqrt{X}$
3. $W = \frac{1}{X}$

Proof

1. $g(u) = \frac{3}{2}u^{1/2}$, for $0 < u \leq 1$
2. $h(v) = 6v^5$ for $0 \leq v \leq 1$
3. $k(w) = \frac{3}{w^4}$ for $1 \leq w < \infty$

Random variables X , U , and V in the previous exercise have *beta distributions*, the same family of distributions that we saw in the exercise [above](#) for the minimum and maximum of independent standard uniform variables. In general, beta distributions are widely used to model random proportions and probabilities, as well as physical quantities that take values in closed bounded intervals (which after a change of units can be taken to be $[0, 1]$). On the other hand, W has a Pareto distribution, named for Vilfredo Pareto. The family of beta distributions and the family of Pareto distributions are studied in more detail in the chapter on Special Distributions.

Suppose that the radius R of a sphere has a beta distribution probability density function f given by $f(r) = 12r^2(1 - r)$ for $0 \leq r \leq 1$. Find the probability density function of each of the following:

1. The circumference $C = 2\pi R$
2. The surface area $A = 4\pi R^2$
3. The volume $V = \frac{4}{3}\pi R^3$

Answer

1. $g(c) = \frac{3}{4\pi^4} c^2 (2\pi - c)$ for $0 \leq c \leq 2\pi$
2. $h(a) = \frac{3}{8\pi^2} \sqrt{a} (2\sqrt{\pi} - \sqrt{a})$ for $0 \leq a \leq 4\pi$
3. $k(v) = \frac{3}{\pi} \left[1 - \left(\frac{3}{4\pi} \right)^{1/3} v^{1/3} \right]$ for $0 \leq v \leq \frac{4}{3}\pi$

Suppose that the grades on a test are described by the random variable $Y = 100X$ where X has the beta distribution with probability density function f given by $f(x) = 12x(1-x)^2$ for $0 \leq x \leq 1$. The grades are generally low, so the teacher decides to “curve” the grades using the transformation $Z = 10\sqrt{Y} = 100\sqrt{X}$. Find the probability density function of

1. Y
2. Z

Answer

1. $g(y) = \frac{3}{25} \left(\frac{y}{100} \right) \left(1 - \frac{y}{100} \right)^2$ for $0 \leq y \leq 100$.
2. $h(z) = \frac{3}{1250} z \left(\frac{z^2}{10\,000} \right) \left(1 - \frac{z^2}{10\,000} \right)^2$ for $0 \leq z \leq 100$

Bernoulli Trials

Recall that a *Bernoulli trials sequence* is a sequence (X_1, X_2, \dots) of independent, identically distributed indicator random variables. In the usual terminology of reliability theory, $X_i = 0$ means *failure* on trial i , while $X_i = 1$ means *success* on trial i . The basic parameter of the process is the probability of success $p = \mathbb{P}(X_i = 1)$, so $p \in [0, 1]$. The random process is named for Jacob Bernoulli and is studied in detail in the chapter on Bernoulli trials.

For $i \in \mathbb{N}_+$, the probability density function f of the trial variable X_i is $f(x) = p^x(1-p)^{1-x}$ for $x \in \{0, 1\}$.

Proof

By definition, $f(0) = 1 - p$ and $f(1) = p$. These can be combined succinctly with the formula $f(x) = p^x(1-p)^{1-x}$ for $x \in \{0, 1\}$.

Now let Y_n denote the number of successes in the first n trials, so that $Y_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}$.

Y_n has the probability density function f_n given by

$$f_n(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (3.7.39)$$

Proof

We have seen this derivation before. The number of bit strings of length n with 1 occurring exactly y times is $\binom{n}{y}$ for $y \in \{0, 1, \dots, n\}$. By the Bernoulli trials assumptions, the probability of each such bit string is $p^y(1-p)^{n-y}$.

The distribution of Y_n is the *binomial distribution* with parameters n and p . The binomial distribution is studied in more detail in the chapter on Bernoulli trials

For $m, n \in \mathbb{N}$

1. $f_n = f^{*n}$.
2. $f_m * f_n = f_{m+n}$.

Proof

Part (a) can be proved directly from the definition of convolution, but the result also follows simply from the fact that $Y_n = X_1 + X_2 + \dots + X_n$.

From part (b) it follows that if Y and Z are independent variables, and that Y has the binomial distribution with parameters $n \in \mathbb{N}$ and $p \in [0, 1]$ while Z has the binomial distribution with parameter $m \in \mathbb{N}$ and p , then $Y + Z$ has the binomial distribution with parameter $m + n$ and p .

Find the probability density function of the difference between the number of successes and the number of failures in $n \in \mathbb{N}$ Bernoulli trials with success parameter $p \in [0, 1]$

Answer

$$f(k) = \binom{n}{(n+k)/2} p^{(n+k)/2} (1-p)^{(n-k)/2} \text{ for } k \in \{-n, -n+2, \dots, n-2, n\}$$

The Poisson Distribution

Recall that the *Poisson distribution* with parameter $t \in (0, \infty)$ has probability density function f given by

$$f_t(n) = e^{-t} \frac{t^n}{n!}, \quad n \in \mathbb{N} \quad (3.7.40)$$

This distribution is named for Simeon Poisson and is widely used to model the number of random points in a region of time or space; the parameter t is proportional to the size of the region. The Poisson distribution is studied in detail in the chapter on The Poisson Process.

If $a, b \in (0, \infty)$ then $f_a * f_b = f_{a+b}$.

Proof

Let $z \in \mathbb{N}$. Using the definition of convolution and the binomial theorem we have

$$(f_a * f_b)(z) = \sum_{x=0}^z f_a(x) f_b(z-x) = \sum_{x=0}^z e^{-a} \frac{a^x}{x!} e^{-b} \frac{b^{z-x}}{(z-x)!} = e^{-(a+b)} \frac{1}{z!} \sum_{x=0}^z \frac{z!}{x!(z-x)!} a^x b^{z-x} \quad (3.7.41)$$

$$= e^{-(a+b)} \frac{1}{z!} \sum_{x=0}^z \binom{z}{x} a^x b^{z-x} = e^{-(a+b)} \frac{(a+b)^z}{z!} = f_{a+b}(z) \quad (3.7.42)$$

The last result means that if X and Y are independent variables, and X has the Poisson distribution with parameter $a > 0$ while Y has the Poisson distribution with parameter $b > 0$, then $X + Y$ has the Poisson distribution with parameter $a + b$. In terms of the Poisson model, X could represent the number of points in a region A and Y the number of points in a region B (of the appropriate sizes so that the parameters are a and b respectively). The independence of X and Y corresponds to the regions A and B being disjoint. Then $X + Y$ is the number of points in $A \cup B$.

The Exponential Distribution

Recall that the *exponential distribution* with rate parameter $r \in (0, \infty)$ has probability density function f given by $f(t) = re^{-rt}$ for $t \in [0, \infty)$. This distribution is often used to model random times such as failure times and lifetimes. In particular, the times between arrivals in the Poisson model of random points in time have independent, identically distributed exponential distributions. The Exponential distribution is studied in more detail in the chapter on Poisson Processes.

Show how to simulate, with a random number, the exponential distribution with rate parameter r . Using your calculator, simulate 5 values from the exponential distribution with parameter $r = 3$.

Answer

$$X = -\frac{1}{r} \ln(1 - U) \text{ where } U \text{ is a random number. Since } 1 - U \text{ is also a random number, a simpler solution is } X = -\frac{1}{r} \ln U.$$

For the next exercise, recall that the *floor* and *ceiling* functions on \mathbb{R} are defined by

$$\lfloor x \rfloor = \max\{n \in \mathbb{Z} : n \leq x\}, \quad \lceil x \rceil = \min\{n \in \mathbb{Z} : n \geq x\}, \quad x \in \mathbb{R} \quad (3.7.43)$$

Suppose that T has the exponential distribution with rate parameter $r \in (0, \infty)$. Find the probability density function of each of the following random variables:

1. $Y = \lfloor T \rfloor$
2. $Z = \lceil T \rceil$

Answer

1. $\mathbb{P}(Y = n) = e^{-rn} (1 - e^{-r}) \text{ for } n \in \mathbb{N}$
2. $\mathbb{P}(Z = n) = e^{-r(n-1)} (1 - e^{-r}) \text{ for } n \in \mathbb{N}$

Note that the distributions in the previous exercise are geometric distributions on \mathbb{N} and on \mathbb{N}_+ , respectively. In many respects, the geometric distribution is a discrete version of the exponential distribution.

Suppose that T has the exponential distribution with rate parameter $r \in (0, \infty)$. Find the probability density function of each of the following random variables:

1. $X = T^2$
2. $Y = e^T$
3. $Z = \ln T$

Answer

1. $g(x) = re^{-r\sqrt{x}} / 2\sqrt{x}$ for $0 < x < \infty$
2. $h(y) = ry^{-(r+1)}$ for $1 < y < \infty$
3. $k(z) = r \exp(-re^z)e^z$ for $z \in \mathbb{R}$

In the previous exercise, Y has a Pareto distribution while Z has an extreme value distribution. Both of these are studied in more detail in the chapter on Special Distributions.

Suppose that X and Y are independent random variables, each having the exponential distribution with parameter 1. Let $Z = \frac{Y}{X}$.

1. Find the distribution function of Z .
2. Find the probability density function of Z .

Answer

1. $G(z) = 1 - \frac{1}{1+z}$, $0 < z < \infty$
2. $g(z) = \frac{1}{(1+z)^2}$, $0 < z < \infty$

Suppose that X has the exponential distribution with rate parameter $a > 0$, Y has the exponential distribution with rate parameter $b > 0$, and that X and Y are independent. Find the probability density function of $Z = X + Y$ in each of the following cases.

1. $a = b$
2. $a \neq b$

Answer

1. $h(z) = a^2 ze^{-az}$ for $0 < z < \infty$
2. $h(z) = \frac{ab}{b-a} (e^{-az} - e^{-bz})$ for $0 < z < \infty$

Suppose that (T_1, T_2, \dots, T_n) is a sequence of independent random variables, and that T_i has the exponential distribution with rate parameter $r_i > 0$ for each $i \in \{1, 2, \dots, n\}$.

1. Find the probability density function of $U = \min\{T_1, T_2, \dots, T_n\}$.
2. Find the distribution function of $V = \max\{T_1, T_2, \dots, T_n\}$.
3. Find the probability density function of V in the special case that $r_i = r$ for each $i \in \{1, 2, \dots, n\}$.

Answer

1. $g(t) = ae^{-at}$ for $0 \leq t < \infty$ where $a = r_1 + r_2 + \dots + r_n$
2. $H(t) = (1 - e^{-r_1 t})(1 - e^{-r_2 t}) \dots (1 - e^{-r_n t})$ for $0 \leq t < \infty$
3. $h(t) = nre^{-rt}(1 - e^{-rt})^{n-1}$ for $0 \leq t < \infty$

Note that the minimum U in part (a) has the exponential distribution with parameter $r_1 + r_2 + \dots + r_n$. In particular, suppose that a series system has independent components, each with an exponentially distributed lifetime. Then the lifetime of the system is also exponentially distributed, and the failure rate of the system is the sum of the component failure rates.

In the order statistic experiment, select the exponential distribution.

1. Set $k = 1$ (this gives the minimum U). Vary n with the scroll bar and note the shape of the probability density function. With $n = 5$, run the simulation 1000 times and compare the empirical density function and the probability density function.
2. Vary n with the scroll bar and set $k = n$ each time (this gives the maximum V). Note the shape of the density function. With $n = 5$, run the simulation 1000 times and compare the empirical density function and the probability density function.

Suppose again that (T_1, T_2, \dots, T_n) is a sequence of independent random variables, and that T_i has the exponential distribution with rate parameter $r_i > 0$ for each $i \in \{1, 2, \dots, n\}$. Then

$$\mathbb{P}(T_i < T_j \text{ for all } j \neq i) = \frac{r_i}{\sum_{j=1}^n r_j} \quad (3.7.44)$$

Proof

When $n = 2$, the result was shown in the section on joint distributions. Returning to the case of general n , note that $T_i < T_j$ for all $j \neq i$ if and only if $T_i < \min\{T_j : j \neq i\}$. Note that the minimum on the right is independent of T_i and by the [result above](#), has an exponential distribution with parameter $\sum_{j \neq i} r_j$.

The result in the previous exercise is very important in the theory of continuous-time Markov chains. If we have a bunch of independent alarm clocks, with exponentially distributed alarm times, then the probability that clock i is the first one to sound is $r_i / \sum_{j=1}^n r_j$.

The Gamma Distribution

Recall that the (standard) *gamma distribution* with shape parameter $n \in \mathbb{N}_+$ has probability density function

$$g_n(t) = e^{-t} \frac{t^{n-1}}{(n-1)!}, \quad 0 \leq t < \infty \quad (3.7.45)$$

With a positive integer shape parameter, as we have here, it is also referred to as the *Erlang distribution*, named for Agner Erlang. This distribution is widely used to model random times under certain basic assumptions. In particular, the n th arrival times in the Poisson model of random points in time has the gamma distribution with parameter n . The Erlang distribution is studied in more detail in the chapter on the Poisson Process, and in greater generality, the gamma distribution is studied in the chapter on Special Distributions.

Let $g = g_1$, and note that this is the probability density function of the exponential distribution with parameter 1, which was the topic of our last discussion.

If $m, n \in \mathbb{N}_+$ then

1. $g_n = g^{*n}$
2. $g_m * g_n = g_{m+n}$

Proof

Part (a) hold trivially when $n = 1$. Also, for $t \in [0, \infty)$,

$$g_n * g(t) = \int_0^t g_n(s)g(t-s) ds = \int_0^t e^{-s} \frac{s^{n-1}}{(n-1)!} e^{-(t-s)} ds = e^{-t} \int_0^t \frac{s^{n-1}}{(n-1)!} ds = e^{-t} \frac{t^n}{n!} = g_{n+1}(t) \quad (3.7.46)$$

Part (b) follows from (a).

Part (b) means that if X has the gamma distribution with shape parameter m and Y has the gamma distribution with shape parameter n , and if X and Y are independent, then $X + Y$ has the gamma distribution with shape parameter $m + n$. In the context of the Poisson model, part (a) means that the n th arrival time is the sum of the n independent interarrival times, which have a common exponential distribution.

Suppose that T has the gamma distribution with shape parameter $n \in \mathbb{N}_+$. Find the probability density function of $X = \ln T$.

Answer

$$h(x) = \frac{1}{(n-1)!} \exp(-e^x) e^{nx} \text{ for } x \in \mathbb{R}$$

The Pareto Distribution

Recall that the *Pareto distribution* with shape parameter $a \in (0, \infty)$ has probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad 1 \leq x < \infty \quad (3.7.47)$$

Members of this family have already come up in several of the previous exercises. The Pareto distribution, named for Vilfredo Pareto, is a heavy-tailed distribution often used for modeling income and other financial variables. The Pareto distribution is studied in more detail in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with shape parameter a . Find the probability density function of each of the following random variables:

1. $U = X^2$
2. $V = \frac{1}{X}$
3. $Y = \ln X$

Answer

1. $g(u) = \frac{a/2}{u^{a/2+1}}$ for $1 \leq u < \infty$
2. $h(v) = av^{a-1}$ for $0 < v < 1$
3. $k(y) = ae^{-ay}$ for $0 \leq y < \infty$

In the previous exercise, V also has a Pareto distribution but with parameter $\frac{a}{2}$; Y has the beta distribution with parameters a and $b = 1$; and Z has the exponential distribution with rate parameter a .

Show how to simulate, with a random number, the Pareto distribution with shape parameter a . Using your calculator, simulate 5 values from the Pareto distribution with shape parameter $a = 2$.

Answer

Using the [random quantile method](#), $X = \frac{1}{(1-U)^{1/a}}$ where U is a random number. More simply, $X = \frac{1}{U^{1/a}}$, since $1 - U$ is also a random number.

The Normal Distribution

Recall that the *standard normal distribution* has probability density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (3.7.48)$$

Suppose that Z has the standard normal distribution, and that $\mu \in (-\infty, \infty)$ and $\sigma \in (0, \infty)$.

1. Find the probability density function f of $X = \mu + \sigma Z$
2. Sketch the graph of f , noting the important qualitative features.

Answer

1. $f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right]$ for $x \in \mathbb{R}$
2. f is symmetric about $x = \mu$. f increases and then decreases, with mode $x = \mu$. f is concave upward, then downward, then upward again, with inflection points at $x = \mu \pm \sigma$. $f(x) \rightarrow 0$ as $x \rightarrow \infty$ and as $x \rightarrow -\infty$

Random variable X has the *normal distribution* with *location parameter* μ and *scale parameter* σ . The normal distribution is perhaps the most important distribution in probability and mathematical statistics, primarily because of the central limit theorem, one of the fundamental theorems. It is widely used to model physical measurements of all types that are subject to small, random errors. The normal distribution is studied in detail in the chapter on Special Distributions.

Suppose that Z has the standard normal distribution. Find the probability density function of Z^2 and sketch the graph.

Answer

$$g(v) = \frac{1}{\sqrt{2\pi v}} e^{-\frac{1}{2}v} \text{ for } 0 < v < \infty$$

Random variable V has the *chi-square distribution* with 1 degree of freedom. Chi-square distributions are studied in detail in the chapter on Special Distributions.

Suppose that X and Y are independent random variables, each with the standard normal distribution, and let (R, Θ) be the standard polar coordinates (X, Y) . Find the probability density function of

1. (R, Θ)
2. R
3. Θ

Answer

Note that the joint PDF of (X, Y) is

$$f(x, y) = \phi(x)\phi(y) = \frac{1}{2\pi} e^{-\frac{1}{2}(x^2+y^2)}, \quad (x, y) \in \mathbb{R}^2 \quad (3.7.49)$$

From the result above [polar coordinates](#), the PDF of (R, Θ) is

$$g(r, \theta) = f(r \cos \theta, r \sin \theta) r = \frac{1}{2\pi} r e^{-\frac{1}{2}r^2}, \quad (r, \theta) \in [0, \infty) \times [0, 2\pi) \quad (3.7.50)$$

From the factorization theorem for joint PDFs, it follows that R has probability density function $h(r) = r e^{-\frac{1}{2}r^2}$ for $0 \leq r < \infty$, Θ is uniformly distributed on $[0, 2\pi)$, and that R and Θ are independent.

The distribution of R is the (standard) *Rayleigh distribution*, and is named for John William Strutt, Lord Rayleigh. The Rayleigh distribution is studied in more detail in the chapter on Special Distributions.

The standard normal distribution does not have a simple, closed form quantile function, so the random quantile method of simulation does not work well. However, the last exercise points the way to an alternative method of simulation.

Show how to simulate a pair of independent, standard normal variables with a pair of random numbers. Using your calculator, simulate 6 values from the standard normal distribution.

Answer

The Rayleigh distribution in the last exercise has CDF $H(r) = 1 - e^{-\frac{1}{2}r^2}$ for $0 \leq r < \infty$, and hence quantile function $H^{-1}(p) = \sqrt{-2 \ln(1-p)}$ for $0 \leq p < 1$. Thus we can simulate the polar radius R with a random number U by $R = \sqrt{-2 \ln(1-U)}$, or a bit more simply by $R = \sqrt{-2 \ln U}$, since $1-U$ is also a random number. We can simulate the polar angle Θ with a random number V by $\Theta = 2\pi V$. Then, a pair of independent, standard normal variables can be simulated by $X = R \cos \Theta$, $Y = R \sin \Theta$.

The Cauchy Distribution

Suppose that X and Y are independent random variables, each with the standard normal distribution. Find the probability density function of $T = X/Y$.

Answer

As usual, let ϕ denote the standard normal PDF, so that $\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ for $z \in \mathbb{R}$. Using the [theorem on quotient](#) above, the PDF f of T is given by

$$f(t) = \int_{-\infty}^{\infty} \phi(x)\phi(tx)|x|dx = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-(1+t^2)x^2/2}|x|dx, \quad t \in \mathbb{R} \quad (3.7.51)$$

Using symmetry and a simple substitution,

$$f(t) = \frac{1}{\pi} \int_0^{\infty} x e^{-(1+t^2)x^2/2} dx = \frac{1}{\pi(1+t^2)}, \quad t \in \mathbb{R} \quad (3.7.52)$$

Random variable T has the (standard) *Cauchy distribution*, named after Augustin Cauchy. The Cauchy distribution is studied in detail in the chapter on Special Distributions.

Suppose that a light source is 1 unit away from position 0 on an infinite straight wall. We shine the light at the wall an angle Θ to the perpendicular, where Θ is uniformly distributed on $(-\frac{\pi}{2}, \frac{\pi}{2})$. Find the probability density function of the position of the light beam $X = \tan \Theta$ on the wall.

Answer

The PDF of Θ is $f(\theta) = \frac{1}{\pi}$ for $-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}$. The transformation is $x = \tan \theta$ so the inverse transformation is $\theta = \arctan x$. Recall that $\frac{d\theta}{dx} = \frac{1}{1+x^2}$, so by the [change of variables formula](#), X has PDF g given by

$$g(x) = \frac{1}{\pi(1+x^2)}, \quad x \in \mathbb{R} \quad (3.7.53)$$

Thus, X also has the standard Cauchy distribution. Clearly we can simulate a value of the Cauchy distribution by $X = \tan(-\frac{\pi}{2} + \pi U)$ where U is a random number. This is the random quantile method.

Open the Cauchy experiment, which is a simulation of the light problem in the previous exercise. Keep the default parameter values and run the experiment in single step mode a few times. Then run the experiment 1000 times and compare the empirical density function and the probability density function.

This page titled [3.7: Transformations of Random Variables](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.8: Convergence in Distribution

This section is concerned with the convergence of probability distributions, a topic of basic importance in probability theory. Since we will be almost exclusively concerned with the convergences of sequences of various kinds, it's helpful to introduce the notation $\mathbb{N}_+^* = \mathbb{N}_+ \cup \{\infty\} = \{1, 2, \dots\} \cup \{\infty\}$.

Distributions on $(\mathbb{R}, \mathcal{R})$

Definition

We start with the most important and basic setting, the measurable space $(\mathbb{R}, \mathcal{R})$, where \mathbb{R} is the set of real numbers of course, and \mathcal{R} is the Borel σ -algebra of subsets of \mathbb{R} . Recall that if P is a probability measure on $(\mathbb{R}, \mathcal{R})$, then the function $F : \mathbb{R} \rightarrow [0, 1]$ defined by $F(x) = P(-\infty, x]$ for $x \in \mathbb{R}$ is the (cumulative) distribution function of P . Recall also that F completely determines P . Here is the definition for convergence of probability measures in this setting:

Suppose P_n is a probability measure on $(\mathbb{R}, \mathcal{R})$ with distribution function F_n for each $n \in \mathbb{N}_+^*$. Then P_n converges (weakly) to P_∞ as $n \rightarrow \infty$ if $F_n(x) \rightarrow F_\infty(x)$ as $n \rightarrow \infty$ for every $x \in \mathbb{R}$ where F_∞ is continuous. We write $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$.

Recall that a distribution function F is continuous at $x \in \mathbb{R}$ if and only if $\mathbb{P}(X = x) = 0$, so that x is not an *atom* of the distribution (a point of positive probability). We will see shortly why this condition on F_∞ is appropriate. Of course, a probability measure on $(\mathbb{R}, \mathcal{R})$ is usually associated with a real-valued random variable for some random experiment that is modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . If X is a real-valued random variable defined on the probability space, then the *distribution* of X is the probability measure P on $(\mathbb{R}, \mathcal{R})$ defined by $P(A) = \mathbb{P}(X \in A)$ for $A \in \mathcal{R}$, and then of course, the distribution function of X is the function F defined by $F(x) = \mathbb{P}(X \leq x)$ for $x \in \mathbb{R}$. Here is the convergence terminology used in this setting:

Suppose that X_n is a real-valued random variable with distribution P_n for each $n \in \mathbb{N}_+^*$. If $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$ then we say that X_n converges in distribution to X_∞ as $n \rightarrow \infty$. We write $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.

So if F_n is the distribution function of X_n for $n \in \mathbb{N}_+^*$, then $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution if $F_n(x) \rightarrow F_\infty(x)$ at every point $x \in \mathbb{R}$ where F_∞ is continuous. On the one hand, the terminology and notation are helpful, since again most probability measures are associated with random variables (and every probability measure can be). On the other hand, the terminology and notation can be a bit misleading since the random variables, as functions, do not converge in any sense, and indeed the random variables need not be defined on the same probability spaces. It is only the *distributions* that converge. However, often the random variables *are* defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$, in which case we can compare convergence in distribution with the other modes of convergence we have or will study:

- Convergence with probability 1
- Convergence in probability
- Convergence in mean

We will show, in fact, that convergence in distribution is the weakest of all of these modes of convergence. However, strength of convergence should not be confused with importance. Convergence in distribution is one of the most important modes of convergence; the central limit theorem, one of the two fundamental theorems of probability, is a theorem about convergence in distribution.

Preliminary Examples

The examples below show why the definition is given in terms of distribution functions, rather than probability density functions, and why convergence is only required at the points of continuity of the limiting distribution function. Note that the distributions considered are probability measures on $(\mathbb{R}, \mathcal{R})$, even though the support of the distribution may be a much smaller subset. For the first example, note that if a deterministic sequence converges in the ordinary calculus sense, then naturally we want the sequence (thought of as random variables) to converge in distribution. Expand the proof to understand the example fully.

Suppose that $x_n \in \mathbb{R}$ for $n \in \mathbb{N}_+^*$. Define random variable $X_n = x_n$ with probability 1 for each $n \in \mathbb{N}_+^*$. Then $x_n \rightarrow x_\infty$ as $n \rightarrow \infty$ if and only if $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.

Proof

For $n \in \mathbb{N}_+^*$, the CDF F_n of X_n is given by $F_n(x) = 0$ for $x < x_n$ and $F_n(x) = 1$ for $x \geq x_n$.

1. Suppose that $x_n \rightarrow x_\infty$ as $n \rightarrow \infty$. If $x < x_\infty$ then $x < x_n$, and hence $F_n(x) = 0$, for all but finitely many $n \in \mathbb{N}_+$, and so $F_n(x) \rightarrow 0$ as $n \rightarrow \infty$. If $x > x_\infty$ then $x > x_n$, and hence $F_n(x) = 1$, for all but finitely many $n \in \mathbb{N}_+$, and so $F_n(x) \rightarrow 1$ as $n \rightarrow \infty$. Nothing can be said about the limiting behavior of $F_n(x_\infty)$ as $n \rightarrow \infty$ without more information. For example, if $x_n \leq x_\infty$ for all but finitely many $n \in \mathbb{N}_+$ then $F_n(x_\infty) \rightarrow 1$ as $n \rightarrow \infty$. If $x_n > x_\infty$ for all but finitely many $n \in \mathbb{N}_+$ then $F_n(x_\infty) \rightarrow 0$ as $n \rightarrow \infty$. If $x_n < x_\infty$ for infinitely many $n \in \mathbb{N}_+$ and $x_n > x_\infty$ for infinitely many $n \in \mathbb{N}_+$ then $F_n(x_\infty)$ does not have a limit as $n \rightarrow \infty$. But regardless, we have $F_n(x) \rightarrow F_\infty(x)$ as $n \rightarrow \infty$ for every $x \in \mathbb{R}$ except perhaps x_∞ , the one point of discontinuity of F_∞ . Hence $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.
2. Conversely, suppose that $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution. If $x < x_\infty$ then $F_n(x) \rightarrow 0$ as $n \rightarrow \infty$ and hence $x < x_n$ for all but finitely many $n \in \mathbb{N}_+$. If $x > x_\infty$ then $F_n(x) \rightarrow 1$ as $n \rightarrow \infty$ and hence $x \geq x_n$ for all but finitely many $n \in \mathbb{N}_+$. So, for every $\epsilon > 0$, $x_n \in (x_\infty - \epsilon, x_\infty + \epsilon)$ for all but finitely many $n \in \mathbb{N}_+$, and hence $x_n \rightarrow x_\infty$ as $n \rightarrow \infty$.

The proof is finished, but let's look at the probability density functions to see that these are not the proper objects of study. For $n \in \mathbb{N}_+^*$, the PDF f_n of X_n is given by $f_n(x_n) = 1$ and $f_n(x) = 0$ for $x \in \mathbb{R} \setminus \{x_n\}$. Only when $x_n = x_\infty$ for all but finitely many $n \in \mathbb{N}_+$ do we have $f_n(x) \rightarrow f(x)$ for $x \in \mathbb{R}$.

For the example below, recall that \mathbb{Q} denotes the set of rational numbers. Once again, expand the proof to understand the example fully

For $n \in \mathbb{N}_+$, let P_n denote the discrete uniform distribution on $\left\{\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\right\}$ and let P_∞ denote the continuous uniform distribution on the interval $[0, 1]$. Then

1. $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$
2. $P_n(\mathbb{Q}) = 1$ for each $n \in \mathbb{N}_+$ but $P_\infty(\mathbb{Q}) = 0$.

Proof

As usual, let F_n denote the CDF of P_n for $n \in \mathbb{N}_+^*$.

1. For $n \in \mathbb{N}_+$ note that F_n is given by $F_n(x) = \lfloor nx \rfloor / n$ for $x \in [0, 1]$. But $nx - 1 \leq \lfloor nx \rfloor \leq nx$ so $\lfloor nx \rfloor / n \rightarrow x$ as $n \rightarrow \infty$ for $x \in [0, 1]$. Of course, $F_n(x) = 0$ for $x < 0$ and $F_n(x) = 1$ for $x > 1$. So $F_n(x) \rightarrow F_\infty(x)$ as $n \rightarrow \infty$ for all $x \in \mathbb{R}$.
2. Note that by definition, so $P_n(\mathbb{Q}) = 1$ for $n \in \mathbb{N}_+$. On the other hand, P_∞ is a continuous distribution and \mathbb{Q} is countable, so $P_\infty(\mathbb{Q}) = 0$.

The proof is finished, but let's look at the probability density functions. For $n \in \mathbb{N}_+$, the PDF f_n of P_n is given by $f_n(x) = \frac{1}{n}$ for $x \in \left\{\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\right\}$ and $f_n(x) = 0$ otherwise. Hence $0 \leq f_n(x) \leq \frac{1}{n}$ for $n \in \mathbb{N}_+$ and $x \in \mathbb{R}$, so $f_n(x) \rightarrow 0$ as $n \rightarrow \infty$ for every $x \in \mathbb{R}$.

The point of the example is that it's reasonable for the discrete uniform distribution on $\left\{\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\right\}$ to converge to the continuous uniform distribution on $[0, 1]$, but once again, the probability density functions are evidently not the correct objects of study.

Probability Density Functions

As the [previous example](#) shows, it is quite possible to have a sequence of discrete distributions converge to a continuous distribution (or the other way around). Recall that probability density functions have very different meanings in the discrete and continuous cases: density with respect to counting measure in the first case, and density with respect to Lebesgue measure in the second case. This is another indication that distribution functions, rather than density functions, are the correct objects of study. However, if probability density functions of a fixed type converge then the distributions converge. Recall again that we are thinking of our probability distributions as measures on $(\mathbb{R}, \mathcal{B})$ even when supported on a smaller subset.

Convergence in distribution in terms of probability density functions.

1. Suppose that f_n is a probability density function for a discrete distribution P_n on a countable set $S \subseteq \mathbb{R}$ for each $n \in \mathbb{N}_+^*$. If $f_n(x) \rightarrow f_\infty(x)$ as $n \rightarrow \infty$ for each $x \in S$ then $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$.
2. Suppose that f_n is a probability density function for a continuous distribution P_n on \mathbb{R} for each $n \in \mathbb{N}_+^*$. If $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$ for all $x \in \mathbb{R}$ (except perhaps on a set with Lebesgue measure 0) then $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$.

Proof

1. Fix $x \in \mathbb{R}$. Then $P_n(-\infty, x] = \sum_{y \in S, y \leq x} f(y)$ for $n \in \mathbb{N}_+$ and $P(-\infty, x] = \sum_{y \in S, y \leq x} f(y)$. It follows from [Scheffé's theorem](#) with the measure space $(S, \mathcal{P}(S), \#)$ that $P_n(-\infty, x] \rightarrow P(-\infty, x]$ as $n \rightarrow \infty$.
2. Fix $x \in \mathbb{R}$. Then $P_n(-\infty, x] = \int_{-\infty}^x f(y) dy$ for $n \in \mathbb{N}_+$ and $P(-\infty, x] = \int_{-\infty}^x f(y) dy$. It follows from [Scheffé's theorem](#) with the measure space $(\mathbb{R}, \mathcal{B}, \lambda)$ that $P_n(-\infty, x] \rightarrow P(-\infty, x]$ as $n \rightarrow \infty$.

Convergence in Probability

Naturally, we would like to compare convergence in distribution with other modes of convergence we have studied.

Suppose that X_n is a real-valued random variable for each $n \in \mathbb{N}_+^*$, all defined on the same probability space. If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in probability then $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.

Proof

Let F_n denote the distribution function of X_n for $n \in \mathbb{N}_+^*$. Fix $\epsilon > 0$. Note first that $\mathbb{P}(X_n \leq x) = \mathbb{P}(X_n \leq x, X_\infty \leq x + \epsilon) + \mathbb{P}(X_n \leq x, X_\infty > x + \epsilon)$. Hence $F_n(x) \leq F_\infty(x + \epsilon) + \mathbb{P}(|X_n - X_\infty| > \epsilon)$. Next, note that $\mathbb{P}(X_\infty \leq x - \epsilon) = \mathbb{P}(X_\infty \leq x - \epsilon, X_n \leq x) + \mathbb{P}(X_\infty \leq x - \epsilon, X_n > x)$. Hence $F_\infty(x - \epsilon) \leq F_n(x) + \mathbb{P}(|X_n - X_\infty| > \epsilon)$. From the last two results it follows that

$$F_\infty(x - \epsilon) - \mathbb{P}(|X_n - X_\infty| > \epsilon) \leq F_n(x) \leq F_\infty(x + \epsilon) + \mathbb{P}(|X_n - X_\infty| > \epsilon) \quad (3.8.1)$$

Letting $n \rightarrow \infty$ and using convergence in probability gives

$$F_\infty(x - \epsilon) \leq \liminf_{n \rightarrow \infty} F_n(x) \leq \limsup_{n \rightarrow \infty} F_n(x) \leq F_\infty(x + \epsilon) \quad (3.8.2)$$

Finally, letting $\epsilon \downarrow 0$ we see that if F_∞ is continuous at x then $F_n(x) \rightarrow F_\infty(x)$ as $n \rightarrow \infty$.

Our next example shows that even when the variables are defined on the same probability space, a sequence can converge in distribution, but not in any other way.

Let X be an indicator variable with $\mathbb{P}(X = 0) = \mathbb{P}(X = 1) = \frac{1}{2}$, so that X is the result of tossing a fair coin. Let $X_n = 1 - X$ for $n \in \mathbb{N}_+$. Then

1. $X_n \rightarrow X$ as $n \rightarrow \infty$ in distribution.
2. $\mathbb{P}(X_n \text{ does not converge to } X \text{ as } n \rightarrow \infty) = 1$.
3. X_n does not converge to X as $n \rightarrow \infty$ in probability.
4. X_n does not converge to X as $n \rightarrow \infty$ in mean.

Proof

1. This trivially holds since $1 - X$ has the same distribution as X .
2. This follows since $|X_n - X| = 1$ for every $n \in \mathbb{N}_+$.
3. This follows since $\mathbb{P}(|X_n - X| > \frac{1}{2}) = 1$ for each $n \in \mathbb{N}_+$.
4. This follows since $\mathbb{E}(|X_n - X|) = 1$ for each $n \in \mathbb{N}_+$.

The critical fact that makes this counterexample work is that $1 - X$ has the same distribution as X . Any random variable with this property would work just as well, so if you prefer a counterexample with continuous distributions, let X have probability density function f given by $f(x) = 6x(1 - x)$ for $0 \leq x \leq 1$. The distribution of X is an example of a beta distribution.

The following summary gives the implications for the various modes of convergence; no other implications hold in general.

Suppose that X_n is a real-valued random variable for each $n \in \mathbb{N}_+^*$, all defined on a common probability space.

1. If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ with probability 1 then $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in probability.
2. If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in mean then $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in probability.

3. If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in probability then $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.

It follows that convergence with probability 1, convergence in probability, and convergence in mean all imply convergence in distribution, so the latter mode of convergence is indeed the weakest. However, our next theorem gives an important converse to part (c) in (7), when the limiting variable is a constant. Of course, a constant can be viewed as a random variable defined on any probability space.

Suppose that X_n is a real-valued random variable for each $n \in \mathbb{N}_+$, defined on the same probability space, and that $c \in \mathbb{R}$. If $X_n \rightarrow c$ as $n \rightarrow \infty$ in distribution then $X_n \rightarrow c$ as $n \rightarrow \infty$ in probability.

Proof

Assume that the probability space is $(\Omega, \mathcal{F}, \mathbb{P})$. Note first that $\mathbb{P}(X_n \leq x) \rightarrow 0$ as $n \rightarrow \infty$ if $x < c$ and $\mathbb{P}(X_n \leq x) \rightarrow 1$ as $n \rightarrow \infty$ if $x > c$. It follows that $\mathbb{P}(|X_n - c| \leq \epsilon) \rightarrow 1$ as $n \rightarrow \infty$ for every $\epsilon > 0$.

The Skorohod Representation

As noted in the [summary above](#), convergence in distribution does not imply convergence with probability 1, even when the random variables are defined on the same probability space. However, the next theorem, known as the *Skorohod representation theorem*, gives an important partial result in this direction.

Suppose that P_n is a probability measure on $(\mathbb{R}, \mathcal{B})$ for each $n \in \mathbb{N}_+^*$ and that $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$. Then there exist real-valued random variables X_n for $n \in \mathbb{N}_+^*$, defined on the same probability space, such that

1. X_n has distribution P_n for $n \in \mathbb{N}_+^*$.
2. $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ with probability 1.

Proof

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and U a random variable defined on this space that is uniformly distributed on the interval $(0, 1)$. For a specific construction, we could take $\Omega = (0, 1)$, \mathcal{F} the σ -algebra of Borel measurable subsets of $(0, 1)$, and \mathbb{P} Lebesgue measure on (Ω, \mathcal{F}) (the uniform distribution on $(0, 1)$). Then let U be the identity function on Ω so that $U(\omega) = \omega$ for $\omega \in \Omega$, so that U has probability distribution \mathbb{P} . We have seen this construction many times before.

1. For $n \in \mathbb{N}_+^*$, let F_n denote the distribution function of P_n and define $X_n = F_n^{-1}(U)$ where F_n^{-1} is the quantile functions of F_n . Recall that X_n has distribution function F_n and therefore X_n has distribution P_n for $n \in \mathbb{N}_+^*$. Of course, these random variables are also defined on $(\Omega, \mathcal{F}, \mathbb{P})$.
2. Let $\epsilon > 0$ and let $u \in (0, 1)$. Pick a continuity point x of F_∞ such that $F_\infty^{-1}(u) - \epsilon < x < F_\infty^{-1}(u)$. Then $F_\infty(x) < u$ and hence $F_n(x) < u$ for all but finitely many $n \in \mathbb{N}_+$. It follows that $F_\infty^{-1}(u) - \epsilon < x < F_n^{-1}(u)$ for all but finitely many $n \in \mathbb{N}_+$. Let $n \rightarrow \infty$ and $u \downarrow 0$ to conclude that $F_\infty^{-1}(u) \leq \liminf_{n \rightarrow \infty} F_n^{-1}(u)$. Next, let v satisfy $0 < u < v < 1$ and let $\epsilon > 0$. Pick a continuity point x of F_∞ such that $F_\infty^{-1}(v) < x < F_\infty^{-1}(v) + \epsilon$. Then $u < v < F_\infty(x)$ and hence $u < F_n(x)$ for all but finitely many $n \in \mathbb{N}_+$. It follows that $F_n^{-1}(u) \leq x < F_\infty^{-1}(v) + \epsilon$ for all but finitely many $n \in \mathbb{N}_+$. Let $n \rightarrow \infty$ and $\epsilon \downarrow 0$ to conclude that $\limsup_{n \rightarrow \infty} F_n^{-1}(u) \leq F_\infty^{-1}(v)$. Letting $v \downarrow u$ it follows that $\limsup_{n \rightarrow \infty} F_n^{-1}(u) \leq F_\infty^{-1}(u)$ if u is a point of continuity of F_∞^{-1} . Therefore $F_n^{-1}(u) \rightarrow F_\infty^{-1}(u)$ as $n \rightarrow \infty$ if u is a point of continuity of F_∞^{-1} . Recall from analysis that since $F_\infty^{-1}(u)$ is increasing, the set $D \subseteq (0, 1)$ of discontinuities of F_∞^{-1} is countable. Since U has a continuous distribution, $\mathbb{P}(U \in D) = 0$. Finally, it follows that $\mathbb{P}(X_n \rightarrow X_\infty \text{ as } n \rightarrow \infty) = 1$.

The following theorem illustrates the value of the Skorohod representation and the usefulness of random variable notation for convergence in distribution. The theorem is also quite intuitive, since a basic idea is that *continuity* should preserve *convergence*.

Suppose that X_n is a real-valued random variable for each $n \in \mathbb{N}_+^*$ (not necessarily defined on the same probability space). Suppose also that $g: \mathbb{R} \rightarrow \mathbb{R}$ is measurable, and let D_g denote the set of discontinuities of g , and P_∞ the distribution of X_∞ . If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution and $P_\infty(D_g) = 0$, then $g(X_n) \rightarrow g(X_\infty)$ as $n \rightarrow \infty$ in distribution.

Proof

By Skorohod's theorem, there exists random variables Y_n for $n \in \mathbb{N}_+^*$, defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$, such that Y_n has the same distribution as X_n for $n \in \mathbb{N}_+^*$, and $Y_n \rightarrow Y_\infty$ as $n \rightarrow \infty$ with probability 1. Since

$\mathbb{P}(Y_\infty \in D_g) = P_\infty(D_g) = 0$ it follows that $g(Y_n) \rightarrow g(Y_\infty)$ as $n \rightarrow \infty$ with probability 1. Hence by the [theorem above](#), $g(Y_n) \rightarrow g(Y_\infty)$ as $n \rightarrow \infty$ in distribution. But $g(Y_n)$ has the same distribution as $g(X_n)$ for each $n \in \mathbb{N}_+^*$.

As a simple corollary, if X_n converges X_∞ as $n \rightarrow \infty$ in distribution, and if $a, b \in \mathbb{R}$ then $a + bX_n$ converges to $a + bX$ as $n \rightarrow \infty$ in distribution. But we can do a little better:

Suppose that X_n is a real-valued random variable and that $a_n, b_n \in \mathbb{R}$ for each $n \in \mathbb{N}_+^*$. If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution and if $a_n \rightarrow a_\infty$ and $b_n \rightarrow b_\infty$ as $n \rightarrow \infty$, then $a_n + b_n X_n \rightarrow a + bX_\infty$ as $n \rightarrow \infty$ in distribution.

Proof

Again by [Skorohod's theorem](#), there exist random variables Y_n for $n \in \mathbb{N}_+^*$, defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$ such that Y_n has the same distribution as X_n for $n \in \mathbb{N}_+^*$ and $Y_n \rightarrow Y_\infty$ as $n \rightarrow \infty$ with probability 1. Hence also $a_n + b_n Y_n \rightarrow a_\infty + b_\infty Y_\infty$ as $n \rightarrow \infty$ with probability 1. By the [result above](#), $a_n + b_n Y_n \rightarrow a_\infty + b_\infty Y_\infty$ as $n \rightarrow \infty$ in distribution. But $a_n + b_n Y_n$ has the same distribution as $a_n + b_n X_n$ for $n \in \mathbb{N}_+^*$.

The definition of convergence in distribution requires that the sequence of probability measures converge on sets of the form $(-\infty, x]$ for $x \in \mathbb{R}$ when the limiting distribution has probability 0 at x . It turns out that the probability measures will converge on lots of other sets as well, and this result points the way to extending convergence in distribution to more general spaces. To state the result, recall that if A is a subset of a topological space, then the *boundary* of A is $\partial A = \text{cl}(A) \setminus \text{int}(A)$ where $\text{cl}(A)$ is the *closure* of A (the smallest closed set that contains A) and $\text{int}(A)$ is the *interior* of A (the largest open set contained in A).

Suppose that P_n is a probability measure on $(\mathbb{R}, \mathcal{R})$ for $n \in \mathbb{N}_+^*$. Then $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$ if and only if $P_n(A) \rightarrow P_\infty(A)$ as $n \rightarrow \infty$ for every $A \in \mathcal{R}$ with $P(\partial A) = 0$.

Proof

Suppose that $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$. Let X_n be a random variable with distribution P_n for $n \in \mathbb{N}_+^*$. (We don't care about the underlying probability spaces.) If $A \in \mathcal{R}$ then the set of discontinuities of $\mathbf{1}_A$, the indicator function of A , is ∂A . So, suppose $\mathbb{P}_\infty(\partial A) = 0$. By the [continuity theorem above](#), $\mathbf{1}_A(X_n) \rightarrow \mathbf{1}_A(X_\infty)$ as $n \rightarrow \infty$ in distribution. Let G_n denote the CDF of $\mathbf{1}_A(X_n)$ for $n \in \mathbb{N}_+^*$. The only possible points of discontinuity of G_∞ are 0 and 1. Hence $G_n(\frac{1}{2}) \rightarrow G_\infty(\frac{1}{2})$ as $n \rightarrow \infty$. But $G_n(\frac{1}{2}) = P_n(A^c)$ for $n \in \mathbb{N}_+^*$. Hence $P_n(A^c) \rightarrow \mathbb{P}_\infty(A^c)$ and so also $P_n(A) \rightarrow P_\infty(A)$ as $n \rightarrow \infty$.

Conversely, suppose that the condition in the theorem holds. If $x \in \mathbb{R}$, then the boundary of $(-\infty, x]$ is $\{x\}$, so if $P_\infty\{x\} = 0$ then $P_n(-\infty, x] \rightarrow P_\infty(-\infty, x]$ as $n \rightarrow \infty$. So by definition, $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$.

In the context of this result, suppose that $a, b \in \mathbb{R}$ with $a < b$. If $P\{a\} = P\{b\} = 0$, then as $n \rightarrow \infty$ we have $P_n(a, b) \rightarrow P(a, b)$, $P_n[a, b) \rightarrow P[a, b)$, $P_n(a, b] \rightarrow P(a, b]$, and $P_n[a, b] \rightarrow P[a, b]$. Of course, the limiting values are all the same.

Examples and Applications

Next we will explore several interesting examples of the convergence of distributions on $(\mathbb{R}, \mathcal{R})$. There are several important cases where a special distribution converges to another special distribution as a parameter approaches a limiting value. Indeed, such convergence results are part of the reason why such distributions are *special* in the first place.

The Hypergeometric Distribution

Recall that the *hypergeometric distribution* with parameters m , r , and n is the distribution that governs the number of type 1 objects in a sample of size n , drawn without replacement from a population of m objects with r objects of type 1. It has discrete probability density function f given by

$$f(k) = \frac{\binom{r}{k} \binom{m-r}{n-k}}{\binom{m}{n}}, \quad k \in \{0, 1, \dots, n\} \quad (3.8.3)$$

The parameters m , r , and n are positive integers with $n \leq m$ and $r \leq m$. The hypergeometric distribution is studied in more detail in the chapter on Finite Sampling Models

Recall next that *Bernoulli trials* are independent trials, each with two possible outcomes, generically called *success* and *failure*. The probability of success $p \in [0, 1]$ is the same for each trial. The *binomial distribution* with parameters $n \in \mathbb{N}_+$ and p is the distribution of the number successes in n Bernoulli trials. This distribution has probability density function g given by

$$g(k) = \binom{n}{k} p^k (1-p)^{n-k}, \quad k \in \{0, 1, \dots, n\} \quad (3.8.4)$$

The binomial distribution is studied in more detail in the chapter on Bernoulli Trials. Note that the binomial distribution with parameters n and $p = r/m$ is the distribution that governs the number of type 1 objects in a sample of size n , drawn *with replacement* from a population of m objects with r objects of type 1. This fact is motivation for the following result:

Suppose that $r_m \in \{0, 1, \dots, m\}$ for each $m \in \mathbb{N}_+$ and that $r_m/m \rightarrow p$ as $m \rightarrow \infty$. For fixed $n \in \mathbb{N}_+$, the hypergeometric distribution with parameters m , r_m , and n converges to the binomial distribution with parameters n and p as $m \rightarrow \infty$.

Proof

Recall that for $a \in \mathbb{R}$ and $j \in \mathbb{N}$, we let $a^{(j)} = a(a-1) \cdots [a-(j-1)]$ denote the falling power of a of order j . The hypergeometric PDF can be written as

$$f_m(k) = \binom{n}{k} \frac{r_m^{(k)} (m-r_m)^{(n-k)}}{m^{(n)}}, \quad k \in \{0, 1, \dots, n\} \quad (3.8.5)$$

In the fraction above, the numerator and denominator both have n factors. Suppose that we group the k factors in $r_m^{(k)}$ with the first k factors of $m^{(n)}$ and the $n-k$ factors of $(m-r_m)^{(n-k)}$ with the last $n-k$ factors of $m^{(n)}$ to form a product of n fractions. The first k fractions have the form $(r_m - j)/(m - j)$ for some j that does not depend on m . Each of these converges to p as $m \rightarrow \infty$. The last $n-k$ fractions have the form $(m - r_m - j)/(m - k - j)$ for some j that does not depend on m . Each of these converges to $1 - p$ as $m \rightarrow \infty$. Hence

$$f_m(k) \rightarrow \binom{n}{k} p^k (1-p)^{n-k} \text{ as } m \rightarrow \infty \text{ for each } k \in \{0, 1, \dots, n\} \quad (3.8.6)$$

The result now follows from the [theorem above](#) on density functions.

From a practical point of view, the last result means that if the population size m is “large” compared to sample size n , then the hypergeometric distribution with parameters m , r , and n (which corresponds to sampling without replacement) is well approximated by the binomial distribution with parameters n and $p = r/m$ (which corresponds to sampling with replacement). This is often a useful result, not computationally, but rather because the binomial distribution has fewer parameters than the hypergeometric distribution (and often in real problems, the parameters may only be known approximately). Specifically, in the limiting binomial distribution, we do not need to know the population size m and the number of type 1 objects r *individually*, but only in the *ratio* r/m .

In the ball and urn experiment, set $m = 100$ and $r = 30$. For each of the following values of n (the sample size), switch between *sampling without replacement* (the hypergeometric distribution) and *sampling with replacement* (the binomial distribution). Note the difference in the probability density functions. Run the simulation 1000 times for each sampling mode and compare the relative frequency function to the probability density function.

1. 10
2. 20
3. 30
4. 40
5. 50

The Binomial Distribution

Recall again that the binomial distribution with parameters $n \in \mathbb{N}_+$ and $p \in [0, 1]$ is the distribution of the number successes in n Bernoulli trials, when p is the probability of success on a trial. This distribution has probability density function f given by

$$f(k) = \binom{n}{k} p^k (1-p)^{n-k}, \quad k \in \{0, 1, \dots, n\} \quad (3.8.7)$$

Recall also that the *Poisson distribution* with parameter $r \in (0, \infty)$ has probability density function g given by

$$g(k) = e^{-r} \frac{r^k}{k!}, \quad k \in \mathbb{N} \quad (3.8.8)$$

The distribution is named for Simeon Poisson and governs the number of “random points” in a region of time or space, under certain ideal conditions. The parameter r is proportional to the size of the region of time or space. The Poisson distribution is studied in more detail in the chapter on the Poisson Process.

Suppose that $p_n \in [0, 1]$ for $n \in \mathbb{N}_+$ and that $np_n \rightarrow r \in (0, \infty)$ as $n \rightarrow \infty$. Then the binomial distribution with parameters n and p_n converges to the Poisson distribution with parameter r as $n \rightarrow \infty$.

Proof

For $k, n \in \mathbb{N}$ with $k \leq n$, the binomial PDF can be written as

$$f_n(k) = \frac{n^{(k)}}{k!} p_n^k (1 - p_n)^{n-k} = \frac{1}{k!} (np_n) [(n-1)p_n] \cdots [(n-k+1)p_n] (1 - p_n)^{n-k} \quad (3.8.9)$$

First, $(n-j)p_n \rightarrow r$ as $n \rightarrow \infty$ for $j \in \{0, 1, \dots, n-1\}$. Next, by a famous limit from calculus, $(1 - p_n)^n = (1 - np_n/n)^n \rightarrow e^{-r}$ as $n \rightarrow \infty$. Hence also $(1 - p_n)^{n-k} \rightarrow e^{-r}$ as $n \rightarrow \infty$ for fixed $k \in \mathbb{N}_+$. Therefore $f_n(k) \rightarrow e^{-r} r^k / k!$ as $n \rightarrow \infty$ for each $k \in \mathbb{N}_+$. The result now follows from the [theorem above](#) on density functions.

From a practical point of view, the convergence of the binomial distribution to the Poisson means that if the number of trials n is “large” and the probability of success p “small”, so that np^2 is small, then the binomial distribution with parameters n and p is well approximated by the Poisson distribution with parameter $r = np$. This is often a useful result, again not computationally, but rather because the Poisson distribution has fewer parameters than the binomial distribution (and often in real problems, the parameters may only be known approximately). Specifically, in the approximating Poisson distribution, we do not need to know the number of trials n and the probability of success p *individually*, but only in the *product* np . As we will see in the next chapter, the condition that np^2 be small means that the variance of the binomial distribution, namely $np(1 - p) = np - np^2$ is approximately $r = np$, which is the variance of the approximating Poisson distribution.

In the binomial timeline experiment, set the parameter values as follows, and observe the graph of the probability density function. (Note that $np = 5$ in each case.) Run the experiment 1000 times in each case and compare the relative frequency function and the probability density function. Note also the successes represented as “random points” in discrete time.

1. $n = 10, p = 0.5$
2. $n = 20, p = 0.25$
3. $n = 100, p = 0.05$

In the Poisson experiment, set $r = 5$ and $t = 1$, to get the Poisson distribution with parameter 5. Note the shape of the probability density function. Run the experiment 1000 times and compare the relative frequency function to the probability density function. Note the similarity between this experiment and the one in the previous exercise.

The Geometric Distribution

Recall that the *geometric distribution* on \mathbb{N}_+ with success parameter $p \in (0, 1]$ has probability density function f given by

$$f(k) = p(1 - p)^{k-1}, \quad k \in \mathbb{N}_+ \quad (3.8.10)$$

The geometric distribution governs the trial number of the first success in a sequence of Bernoulli trials.

Suppose that U has the geometric distribution on \mathbb{N}_+ with success parameter $p \in (0, 1]$. For $n \in \mathbb{N}_+$, the conditional distribution of U given $U \leq n$ converges to the uniform distribution on $\{1, 2, \dots, n\}$ as $p \downarrow 0$.

Proof

The CDF F of U is given by $F(k) = 1 - (1 - p)^k$ for $k \in \mathbb{N}_+$. Hence for $n \in \mathbb{N}_+$, the conditional CDF of U given $U \leq n$ is

$$F_n(k) = \mathbb{P}(U \leq k | U \leq n) = \frac{\mathbb{P}(U \leq k)}{\mathbb{P}(U \leq n)} = \frac{1 - (1-p)^k}{1 - (1-p)^n}, \quad k \in \{1, 2, \dots, n\} \quad (3.8.11)$$

Using L'Hospital's rule, gives $F_n(k) \rightarrow k/n$ as $p \downarrow 0$ for $k \in \{1, 2, \dots, n\}$. As a function of k this is the CDF of the uniform distribution on $\{1, 2, \dots, n\}$.

Next, recall that the *exponential distribution* with rate parameter $r \in (0, \infty)$ has distribution function G given by

$$G(t) = 1 - e^{-rt}, \quad 0 \leq t < \infty \quad (3.8.12)$$

The exponential distribution governs the time between “arrivals” in the Poisson model of random points in time.

Suppose that U_n has the geometric distribution on \mathbb{N}_+ with success parameter $p_n \in (0, 1]$ for $n \in \mathbb{N}_+$, and that $np_n \rightarrow r \in (0, \infty)$ as $n \rightarrow \infty$. The distribution of U_n/n converges to the exponential distribution with parameter r as $n \rightarrow \infty$.

Proof

Let F_n denote the CDF of U_n/n . Then for $x \in [0, \infty)$

$$F_n(x) = \mathbb{P}\left(\frac{U_n}{n} \leq x\right) = \mathbb{P}(U_n \leq nx) = \mathbb{P}(U_n \leq \lfloor nx \rfloor) = 1 - (1-p_n)^{\lfloor nx \rfloor} \quad (3.8.13)$$

We showed in the proof of the [convergence of the binomial distribution](#) that $(1-p_n)^n \rightarrow e^{-r}$ as $n \rightarrow \infty$, and hence $(1-p_n)^{nx} \rightarrow e^{-rx}$ as $n \rightarrow \infty$. But by definition, $\lfloor nx \rfloor \leq nx < \lfloor nx \rfloor + 1$ or equivalently, $nx - 1 < \lfloor nx \rfloor \leq nx$ so it follows from the squeeze theorem that $(1-p_n)^{\lfloor nx \rfloor} \rightarrow e^{-rx}$ as $n \rightarrow \infty$. Hence $F_n(x) \rightarrow 1 - e^{-rx}$ as $n \rightarrow \infty$. As a function of $x \in [0, \infty)$, this is the CDF of the exponential distribution with parameter r .

Note that the limiting condition on n and p in the last result is precisely the same as the condition for the [convergence of the binomial distribution](#) to the Poisson distribution. For a deeper interpretation of both of these results, see the section on the Poisson distribution.

In the negative binomial experiment, set $k = 1$ to get the geometric distribution. Then decrease the value of p and note the shape of the probability density function. With $p = 0.5$ run the experiment 1000 times and compare the relative frequency function to the probability density function.

In the gamma experiment, set $k = 1$ to get the exponential distribution, and set $r = 5$. Note the shape of the probability density function. Run the experiment 1000 times and compare the empirical density function and the probability density function. Compare this experiment with the one in the previous exercise, and note the similarity, up to a change in scale.

The Matching Distribution

For $n \in \mathbb{N}_+$, consider a random permutation (X_1, X_2, \dots, X_n) of the elements in the set $\{1, 2, \dots, n\}$. We say that a *match* occurs at position i if $X_i = i$.

$$\mathbb{P}(X_i = i) = \frac{1}{n} \text{ for each } i \in \{1, 2, \dots, n\}.$$

Proof

The number of permutations of $\{1, 2, \dots, n\}$ is $n!$. For $i \in \{1, 2, \dots, n\}$, the number of such permutations with i in position i is $(n-1)!$. Hence $\mathbb{P}(X_i = i) = (n-1)!/n! = 1/n$. A more direct argument is that i is no more or less likely to end up in position i as any other number.

So the matching events all have the same probability, which varies inversely with the number of trials.

$$\mathbb{P}(X_i = i, X_j = j) = \frac{1}{n(n-1)} \text{ for } i, j \in \{1, 2, \dots, n\} \text{ with } i \neq j.$$

Proof

Again, the number of permutations of $\{1, 2, \dots, n\}$ is $n!$. For distinct $i, j \in \{1, 2, \dots, n\}$, the number of such permutations with i in position i and j in position j is $(n-2)!$. Hence $\mathbb{P}(X_i = i, X_j = j) = (n-2)!/n! = 1/n(n-1)$.

So the matching events are dependent, and in fact are positively correlated. In particular, the matching events do not form a sequence of Bernoulli trials. The matching problem is studied in detail in the chapter on Finite Sampling Models. In that section we show that the number of matches N_n has probability density function f_n given by:

$$f_n(k) = \frac{1}{k!} \sum_{j=0}^{n-k} \frac{(-1)^j}{j!}, \quad k \in \{0, 1, \dots, n\} \quad (3.8.14)$$

The distribution of N_n converges to the Poisson distribution with parameter 1 as $n \rightarrow \infty$.

Proof

For $k \in \mathbb{N}$,

$$f_n(k) = \frac{1}{k!} \sum_{j=0}^{n-k} \frac{(-1)^j}{j!} \rightarrow \frac{1}{k!} \sum_{j=0}^{\infty} \frac{(-1)^j}{j!} = \frac{1}{k!} e^{-1} \quad (3.8.15)$$

As a function of $k \in \mathbb{N}$, this is the PDF of the Poisson distribution with parameter 1. So the result follows from the [theorem above](#) on density functions.

In the matching experiment, increase n and note the apparent convergence of the probability density function for the number of matches. With selected values of n , run the experiment 1000 times and compare the relative frequency function and the probability density function.

The Extreme Value Distribution

Suppose that (X_1, X_2, \dots) is a sequence of independent random variables, each with the standard exponential distribution (parameter 1). Thus, recall that the common distribution function G is given by

$$G(x) = 1 - e^{-x}, \quad 0 \leq x < \infty \quad (3.8.16)$$

As $n \rightarrow \infty$, the distribution of $Y_n = \max\{X_1, X_2, \dots, X_n\} - \ln n$ converges to the distribution with distribution function F given by

$$F(x) = e^{-e^{-x}}, \quad x \in \mathbb{R} \quad (3.8.17)$$

Proof

Let $X_{(n)} = \max\{X_1, X_2, \dots, X_n\}$ and recall that $X_{(n)}$ has CDF G^n . Let F_n denote the CDF of Y_n . For $x \in \mathbb{R}$

$$F_n(x) = \mathbb{P}(Y_n \leq x) = \mathbb{P}(X_{(n)} \leq x + \ln n) = G^n(x + \ln n) = \left[1 - e^{-(x + \ln n)}\right]^n = \left(1 - \frac{e^{-x}}{n}\right)^n \quad (3.8.18)$$

By our famous limit from calculus again, $F_n(x) \rightarrow e^{-e^{-x}}$ as $n \rightarrow \infty$.

The limiting distribution in Exercise (27) is the *standard extreme value distribution*, also known as the *standard Gumbel distribution* in honor of Emil Gumbel. Extreme value distributions are studied in detail in the chapter on Special Distributions.

The Pareto Distribution

Recall that the *Pareto distribution* with shape parameter $a \in (0, \infty)$ has distribution function F given by

$$F(x) = 1 - \frac{1}{x^a}, \quad 1 \leq x < \infty \quad (3.8.19)$$

The Pareto distribution, named for Vilfredo Pareto, is a heavy-tailed distribution sometimes used to model financial variables. It is studied in more detail in the chapter on Special Distributions.

Suppose that X_n has the Pareto distribution with parameter n for each $n \in \mathbb{N}_+$. Then

1. $X_n \rightarrow 1$ as $n \rightarrow \infty$ in distribution (and hence also in probability).
2. The distribution of $Y_n = nX_n - n$ converges to the standard exponential distribution as $n \rightarrow \infty$.

Proof

1. The CDF of X_n is $F_n(x) = 1 - 1/x^n$ for $x \geq 1$. Hence $F_n(x) = 0$ for $n \in \mathbb{N}_+$ and $x \leq 1$ while $F_n(x) \rightarrow 1$ as $n \rightarrow \infty$ for $x > 1$. Thus the limit of F_n agrees with the CDF of the constant 1, except at $x = 1$, the point of discontinuity.
2. Let G_n denote the CDF of Y_n . For $x \geq 0$,

$$G_n(x) = \mathbb{P}(Y_n \leq x) = \mathbb{P}(X_n \leq 1 + x/n) = 1 - \frac{1}{(1 + x/n)^n} \quad (3.8.20)$$

By our famous theorem from calculus again, it follows that $G_n(x) \rightarrow 1 - 1/e^x = 1 - e^{-x}$ as $n \rightarrow \infty$. As a function of $x \in [0, \infty)$, this is the CDF of the standard exponential distribution.

Fundamental Theorems

The two fundamental theorems of basic probability theory, the law of large numbers and the central limit theorem, are studied in detail in the chapter on Random Samples. For this reason we will simply state the results in this section. So suppose that (X_1, X_2, \dots) is a sequence of independent, identically distributed, real-valued random variables (defined on the same probability space) with mean $\mu \in (-\infty, \infty)$ and standard deviation $\sigma \in (0, \infty)$. For $n \in \mathbb{N}_+$, let $Y_n = \sum_{i=1}^n X_i$ denote the sum of the first n variables, $M_n = Y_n/n$ the average of the first n variables, and $Z_n = (Y_n - n\mu)/\sqrt{n}\sigma$ the standard score of Y_n .

The fundamental theorems of probability

1. $M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1 (and hence also in probability and in distribution). This is the *law of large numbers*.
2. The distribution of Z_n converges to the standard normal distribution as $n \rightarrow \infty$. This is the *central limit theorem*.

In part (a), convergence with probability 1 is the *strong law of large numbers* while convergence in probability and in distribution are the *weak laws of large numbers*.

General Spaces

Our next goal is to define convergence of probability distributions on more general measurable spaces. For this discussion, you may need to refer to other sections in this chapter: the integral with respect to a positive measure, properties of the integral, and density functions. In turn, these sections depend on measure theory developed in the chapters on Foundations and Probability Measures.

Definition and Basic Properties

First we need to define the type of measurable spaces that we will use in this subsection.

We assume that (S, d) is a complete, separable metric space and let \mathcal{S} denote the Borel σ -algebra of subsets of S , that is, the σ -algebra generated by the topology. The standard spaces that we often use are special cases of the measurable space (S, \mathcal{S}) :

1. *Discrete*: S is countable and is given the discrete metric so \mathcal{S} is the collection of all subsets of S .
2. *Euclidean*: \mathbb{R}^n is given the standard Euclidean metric so \mathcal{B}_n is the usual σ -algebra of Borel measurable subsets of \mathbb{R}^n .

Additional details

Recall that the metric space (S, d) is *complete* if every Cauchy sequence in S converges to a point in S . The space is *separable* if there exists a countable subset that is dense. A complete, separable metric space is sometimes called a *Polish space* because such spaces were extensively studied by a group of Polish mathematicians in the 1930s, including Kazimierz Kuratowski.

As suggested by our setup, the definition for convergence in distribution involves both measure theory and topology. The motivation is the [theorem above](#) for the one-dimensional Euclidean space $(\mathbb{R}, \mathcal{B})$.

Convergence in distribution:

1. Suppose that P_n is a probability measure on (S, \mathcal{S}) for each $n \in \mathbb{N}_+$. Then P_n *converges (weakly)* to P_∞ as $n \rightarrow \infty$ if $P_n(A) \rightarrow P_\infty(A)$ as $n \rightarrow \infty$ for every $A \in \mathcal{S}$ with $P_\infty(\partial A) = 0$. We write $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$.

2. Suppose that X_n is a random variable with distribution P_n on (S, \mathcal{S}) for each $n \in \mathbb{N}_+^*$. Then X_n converges in distribution to X_∞ as $n \rightarrow \infty$ if $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$. We write $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.

Notes

1. The definition makes sense since $A \in \mathcal{S}$ implies $\partial A \in \mathcal{S}$. Specifically, $\text{cl}(A) \in \mathcal{S}$ because $\text{cl}(A)$ is closed, and $\text{int}(A) \in \mathcal{S}$ because $\text{int}(A)$ is open.
2. The random variables need not be defined on the same probability space.

Let's consider our two special cases. In the discrete case, as usual, the measure theory and topology are not really necessary.

Suppose that P_n is a probability measures on a discrete space (S, \mathcal{S}) for each $n \in \mathbb{N}_+^*$. Then $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$ if and only if $P_n(A) \rightarrow P_\infty(A)$ as $n \rightarrow \infty$ for every $A \subseteq S$.

Proof

This follows from the definition. Every subset is both open and closed so $\partial A = \emptyset$ for every $A \subseteq S$.

In the Euclidean case, it suffices to consider distribution functions, as in the one-dimensional case. If P is a probability measure on $(\mathbb{R}^n, \mathcal{B}_n)$, recall that the *distribution function* F of P is given by

$$F(x_1, x_2, \dots, x_n) = P((-\infty, x_1] \times (-\infty, x_2] \times \dots \times (-\infty, x_n]), \quad (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (3.8.21)$$

Suppose that P_n is a probability measures on $(\mathbb{R}^n, \mathcal{B}_n)$ with distribution function F_n for each $n \in \mathbb{N}_+^*$. Then $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$ if and only if $F_n(\mathbf{x}) \rightarrow F_\infty(\mathbf{x})$ as $n \rightarrow \infty$ for every $\mathbf{x} \in \mathbb{R}^n$ where F_∞ is continuous.

Convergence in Probability

As in the case of $(\mathbb{R}, \mathcal{B})$, convergence in probability implies convergence in distribution.

Suppose that X_n is a random variable with values in S for each $n \in \mathbb{N}_+^*$, all defined on the same probability space. If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in probability then $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.

Notes

Assume that the common probability space is $(\Omega, \mathcal{F}, \mathbb{P})$. Recall that convergence in probability means that $\mathbb{P}[d(X_n, X_\infty) > \epsilon] \rightarrow 0$ as $n \rightarrow \infty$ for every $\epsilon > 0$,

So as before, convergence with probability 1 implies convergence in probability which in turn implies convergence in distribution.

Skorohod's Representation Theorem

As you might guess, [Skorohod's theorem](#) for the one-dimensional Euclidean space $(\mathbb{R}, \mathcal{B})$ can be extended to the more general spaces. However the proof is not nearly as straightforward, because we no longer have the quantile function for constructing random variables on a common probability space.

Suppose that P_n is a probability measures on (S, \mathcal{S}) for each $n \in \mathbb{N}_+^*$ and that $P_n \Rightarrow P_\infty$ as $n \rightarrow \infty$. Then there exists a random variable X_n with values in S for each $n \in \mathbb{N}_+^*$, defined on a common probability space, such that

1. X_n has distribution P_n for $n \in \mathbb{N}_+^*$
2. $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ with probability 1.

One of the main consequences of Skorohod's representation, the preservation of convergence in distribution under continuous functions, is still true and has essentially the same proof. For the general setup, suppose that (S, d, \mathcal{S}) and (T, e, \mathcal{T}) are spaces of the type [described above](#).

Suppose that X_n is a random variable with values in S for each $n \in \mathbb{N}_+^*$ (not necessarily defined on the same probability space). Suppose also that $g: S \rightarrow T$ is measurable, and let D_g denote the set of discontinuities of g , and P_∞ the distribution of X_∞ . If $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution and $P_\infty(D_g) = 0$, then $g(X_n) \rightarrow g(X_\infty)$ as $n \rightarrow \infty$ in distribution.

Proof

By Skorohod's theorem, there exists random variables Y_n with values in S for $n \in \mathbb{N}_+^*$, defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$, such that Y_n has the same distribution as X_n for $n \in \mathbb{N}_+^*$, and $Y_n \rightarrow Y_\infty$ as $n \rightarrow \infty$ with probability 1. Since $\mathbb{P}(Y_\infty \in D_g) = P_\infty(D_g) = 0$ it follows that $g(Y_n) \rightarrow g(Y_\infty)$ as $n \rightarrow \infty$ with probability 1. Hence $g(Y_n) \rightarrow g(Y_\infty)$ as $n \rightarrow \infty$ in distribution. But $g(Y_n)$ has the same distribution as $g(X_n)$ for each $n \in \mathbb{N}_+^*$.

A simple consequence of the continuity theorem is that if a sequence of random vectors in \mathbb{R}^n converge in distribution, then the sequence of each coordinate also converges in distribution. Let's just consider the two-dimensional case to keep the notation simple.

Suppose that (X_n, Y_n) is a random variable with values in \mathbb{R}^2 for $n \in \mathbb{N}_+^*$ and that $(X_n, Y_n) \rightarrow (X_\infty, Y_\infty)$ as $n \rightarrow \infty$ in distribution. Then

1. $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ in distribution.
2. $Y_n \rightarrow Y_\infty$ as $n \rightarrow \infty$ in distribution.

Scheffé's Theorem

Our next discussion concerns an important result known as *Scheffé's theorem*, named after Henry Scheffé. To state our theorem, suppose that (S, \mathcal{S}, μ) is a measure space, so that S is a set, \mathcal{S} is a σ -algebra of subsets of S , and μ is a positive measure on (S, \mathcal{S}) . Further, suppose that P_n is a probability measure on (S, \mathcal{S}) that has density function f_n with respect to μ for each $n \in \mathbb{N}_+$, and that P is a probability measure on (S, \mathcal{S}) that has density function f with respect to μ .

If $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$ for almost all $x \in S$ (with respect to μ) then $P_n(A) \rightarrow P(A)$ as $n \rightarrow \infty$ uniformly in $A \in \mathcal{S}$.

Proof

From basic properties of the integral it follows that for $A \in \mathcal{S}$,

$$|P(A) - P_n(A)| = \left| \int_A f d\mu - \int_A f_n d\mu \right| = \left| \int_A (f - f_n) d\mu \right| \leq \int_A |f - f_n| d\mu \leq \int_S |f - f_n| d\mu \quad (3.8.22)$$

Let $g_n = f - f_n$, and let g_n^+ denote the positive part of g_n and g_n^- the negative part of g_n . Note that $g_n^+ \leq f$ and $g_n^+ \rightarrow 0$ as $n \rightarrow \infty$ almost everywhere on S . Since f is a probability density function, it is trivially integrable, so by the dominated convergence theorem, $\int_S g_n^+ d\mu \rightarrow 0$ as $n \rightarrow \infty$. But $\int_{\mathbb{R}} g_n d\mu = 0$ so $\int_{\mathbb{R}} g_n^+ d\mu = \int_{\mathbb{R}} g_n^- d\mu$. Therefore $\int_S |g_n| d\mu = 2 \int_S g_n^+ d\mu \rightarrow 0$ as $n \rightarrow \infty$. Hence $P_n(A) \rightarrow P(A)$ as $n \rightarrow \infty$ uniformly in $A \in \mathcal{S}$.

Of course, the most important special cases of Scheffé's theorem are to discrete distributions and to continuous distributions on a subset of \mathbb{R}^n , as in the [theorem above](#) on density functions.

Expected Value

Generating functions are studied in the chapter on Expected Value. In part, the importance of generating functions stems from the fact that ordinary (pointwise) convergence of a sequence of generating functions corresponds to the convergence of the distributions in the sense of this section. Often it is easier to show convergence in distribution using generating functions than directly from the definition.

In addition, convergence in distribution has elegant characterizations in terms of the convergence of the expected values of certain types of functions of the underlying random variables.

This page titled [3.8: Convergence in Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.9: General Distribution Functions

Our goal in this section is to define and study functions that play the same role for positive measures on \mathbb{R} that (cumulative) distribution functions do for probability measures on \mathbb{R} . Of course probability measures on \mathbb{R} are usually associated with real-valued random variables. These general *distribution functions* are useful for constructing measures on \mathbb{R} and will appear in our study of integrals with respect to a measure in the next section, as well as non-homogeneous Poisson processes and general renewal processes.

Basic Theory

Throughout this section, our basic measurable space is $(\mathbb{R}, \mathcal{R})$, where \mathcal{R} is the σ -algebra of Borel measurable subsets of \mathbb{R} , and as usual, we will let λ denote Lebesgue measure on $(\mathbb{R}, \mathcal{R})$. As with cumulative distribution functions, it's convenient to have compact notation for the limits of a function $F : \mathbb{R} \rightarrow \mathbb{R}$ from the left and right at $x \in \mathbb{R}$, and at ∞ and $-\infty$ (assuming of course that these limits exist):

$$F(x^+) = \lim_{t \downarrow x} F(t), \quad F(x^-) = \lim_{t \uparrow x} F(t), \quad F(\infty) = \lim_{t \rightarrow \infty} F(t), \quad F(-\infty) = \lim_{t \rightarrow -\infty} F(t) \quad (3.9.1)$$

Distribution Functions and Their Measures

A function $F : \mathbb{R} \rightarrow \mathbb{R}$ that satisfies the following properties is a *distribution function* on \mathbb{R}

1. F is increasing: if $x \leq y$ then $F(x) \leq F(y)$.
2. F is continuous from the right: $F(x^+) = F(x)$ for all $x \in \mathbb{R}$.

Since F is increasing, $F(x^-)$ exists in \mathbb{R} . Similarly $F(\infty)$ exists, as a real number or ∞ , and $F(-\infty)$ exists, as a real number or $-\infty$.

If F is a distribution function on \mathbb{R} , then there exists a unique positive measure μ on \mathcal{R} that satisfies

$$\mu(a, b] = F(b) - F(a), \quad a, b \in \mathbb{R}, \quad a \leq b \quad (3.9.2)$$

Proof

Let \mathcal{J} denote the collection of subsets of \mathbb{R} consisting of intervals of the form $(a, b]$ where $a, b \in \mathbb{R}$ with $a \leq b$, and intervals of the form $(-\infty, a]$ and (a, ∞) where $a \in \mathbb{R}$. Then \mathcal{J} is a semi-algebra. That is, if $A, B \in \mathcal{J}$ then $A \cap B \in \mathcal{J}$, and if $A \in \mathcal{J}$ then A^c is the union of a finite number (actually one or two) sets in \mathcal{J} . We define μ on \mathcal{J} by $\mu(a, b] = F(b) - F(a)$, $\mu(-\infty, a] = F(a) - F(-\infty)$ and $\mu(a, \infty) = F(\infty) - F(a)$. Note that \mathcal{J} contains the empty set via intervals of the form $(a, a]$ where $a \in \mathbb{R}$, but the definition gives $\mu(\emptyset) = 0$. Next, μ is *finitely additive* on \mathcal{J} . That is, if $\{A_i : i \in I\}$ is a finite, disjoint collection of sets in \mathcal{J} and $\bigcup_{i \in I} A_i \in \mathcal{J}$, then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i) \quad (3.9.3)$$

Next, μ is *countably subadditive* on \mathcal{J} . That is, if $A \in \mathcal{J}$ and $A \subseteq \bigcup_{i \in I} A_i$ where $\{A_i : i \in I\}$ is a countable collection of sets in \mathcal{J} then

$$\mu(A) \leq \sum_{i \in I} \mu(A_i) \quad (3.9.4)$$

Finally, μ is clearly σ -finite on \mathcal{J} since $\mu(a, b] < \infty$ for $a, b \in \mathbb{R}$ with $a < b$, and \mathbb{R} is a countable, disjoint union of intervals of this form. Hence it follows from the basic extension and uniqueness theorems that μ can be extended uniquely to a measure on the $\mathcal{R} = \sigma(\mathcal{J})$.

For the final uniqueness part, suppose that μ is a measure on \mathcal{R} satisfying $\mu(a, b] = F(b) - F(a)$ for $a, b \in \mathbb{R}$ with $a < b$. Then by the continuity theorem for increasing sets, $\mu(-\infty, a] = F(a) - F(-\infty)$ and $\mu(a, \infty) = F(\infty) - F(a)$ for $a \in \mathbb{R}$. Hence μ is the unique measure constructed above.

The measure μ is called the *Lebesgue-Stieltjes measure* associated with F , named for Henri Lebesgue and Thomas Joannes Stieltjes. A very rich variety of measures on \mathbb{R} can be constructed in this way. In particular, when the function F takes values in $[0, 1]$, the associated measure \mathbb{P} is a probability measure. Another special case of interest is the distribution function defined by $F(x) = x$ for $x \in \mathbb{R}$, in which case $\mu(a, b]$ is the length of the interval $(a, b]$ and therefore $\mu = \lambda$, Lebesgue measure on \mathcal{R} . But although the measure associated with a distribution function is unique, the distribution function itself is not. Note that if $c \in \mathbb{R}$ then the distribution function defined by $F(x) = x + c$ for $x \in \mathbb{R}$ also generates Lebesgue measure. This example captures the general situation.

Suppose that F and G are distribution functions that generate the same measure μ on \mathbb{R} . Then there exists $c \in \mathbb{R}$ such that $G = F + c$.

Proof

For $x \in \mathbb{R}$, note that $F(x) - F(0) = G(x) - G(0)$. The common value is $\mu(0, x]$ if $x \geq 0$ and $-\mu(x, 0]$ if $x < 0$. Thus $G(x) = F(x) - F(0) + G(0)$ for $x \in \mathbb{R}$.

Returning to the case of a probability measure \mathbb{P} on \mathbb{R} , the cumulative distribution function F that we studied in this chapter is the unique distribution function satisfying $F(-\infty) = 0$. More generally, having constructed a measure from a distribution function, let's now consider the complementary problem of finding a distribution function for a given measure. The proof of the last theorem points the way.

Suppose that μ is a positive measure on $(\mathbb{R}, \mathcal{R})$ with the property that $\mu(A) < \infty$ if A is bounded. Then there exists a distribution function that generates μ .

Proof

Define F on \mathbb{R} by

$$F(x) = \begin{cases} \mu(0, x], & x \geq 0 \\ -\mu(x, 0], & x < 0 \end{cases} \quad (3.9.5)$$

Then $F : \mathbb{R} \rightarrow \mathbb{R}$ by the assumption on μ . Also F is increasing: if $0 \leq x \leq y$ then $\mu(0, x] \leq \mu(0, y]$ by the increasing property of a positive measure. Similarly, if $x \leq y \leq 0$, the $\mu(x, 0] \geq \mu(y, 0]$, so $-\mu(x, 0] \leq -\mu(y, 0]$. Finally, if $x \leq 0 \leq y$, then $-\mu(x, 0] \leq 0$ and $\mu(0, y] \geq 0$. Next, F is continuous from the right: Suppose that $x_n \in \mathbb{R}$ for $n \in \mathbb{N}_+$ and $x_n \downarrow x$ as $n \rightarrow \infty$. If $x \geq 0$ then $\mu(0, x_n] \downarrow \mu(0, x]$ by the continuity theorem for decreasing sets, which applies since the measures are finite. If $x < 0$ then $\mu(x_n, 0] \uparrow \mu(x, 0]$ by the continuity theorem for increasing sets. So in both cases, $F(x_n) \downarrow F(x)$ as $n \rightarrow \infty$. Hence F is a distribution function, and it remains to show that it generates μ . Let $a, b \in \mathbb{R}$ with $a \leq b$. If $a \geq 0$ then $\mu(a, b] = \mu(0, b] - \mu(0, a] = F(b) - F(a)$ by the difference property of a positive measure. Similarly, if $b \leq 0$ then $\mu(a, b] = \mu(a, 0] - \mu(b, 0] = -F(a) + F(b)$. Finally, if $a \leq 0$ and $b \geq 0$, then $\mu(a, b] = \mu(a, 0] + \mu(0, b] = -F(a) + F(b)$.

In the proof of the last theorem, the use of 0 as a “reference point” is arbitrary, of course. Any other point in \mathbb{R} would do as well, and would produce a distribution function that differs from the one in the proof by a constant. If μ has the property that $\mu(-\infty, x] < \infty$ for $x \in \mathbb{R}$, then it's easy to see that F defined by $F(x) = \mu(-\infty, x]$ for $x \in \mathbb{R}$ is a distribution function that generates μ , and is the unique distribution function with $F(-\infty) = 0$. Of course, in the case of a probability measure, this is the cumulative distribution function, as noted above.

Properties

General distribution functions enjoy many of the same properties as the cumulative distribution function (but not all because of the lack of uniqueness). In particular, we can easily compute the measure of any interval from the distribution function.

Suppose that F is a distribution function and μ is the positive measure on $(\mathbb{R}, \mathcal{R})$ associated with F . For $a, b \in \mathbb{R}$ with $a < b$,

1. $\mu[a, b] = F(b) - F(a^-)$
2. $\mu\{a\} = F(a) - F(a^-)$
3. $\mu(a, b) = F(b^-) - F(a)$
4. $\mu[a, b) = F(b^-) - F(a^-)$

Proof

All of these results follow from the continuity theorems for a positive measure. Suppose that (x_1, x_2, \dots) is a sequence of distinct points in \mathbb{R} .

1. If $x_n \uparrow a$ as $n \rightarrow \infty$ then $(x_n, b] \uparrow [a, b]$ so $\mu(x_n, b] \uparrow \mu[a, b]$ as $n \rightarrow \infty$. But also $\mu(x_n, b] = F(b) - F(x_n) \rightarrow F(b) - F(a^-)$ as $n \rightarrow \infty$.
2. This follows from (a) by taking $a = b$
3. If $x_n \uparrow b$ as $n \rightarrow \infty$ then $(a, x_n] \uparrow (a, b)$ so $\mu(a, x_n] \uparrow \mu(a, b)$ as $n \rightarrow \infty$. But also $\mu(a, x_n] = F(x_n) - F(a) \rightarrow F(b^-) - F(a)$ as $n \rightarrow \infty$.
4. From (a) and (b) and the difference rule,

$$\mu[a, b] = \mu[a, b] - \mu\{b\} = F(b) - F(a^-) - [F(b) - F(b^-)] = F(b^-) - F(a^-) \quad (3.9.6)$$

Note that F is continuous at $x \in \mathbb{R}$ if and only if $\mu\{x\} = 0$. In particular, μ is a *continuous measure* (recall that this means that $\mu\{x\} = 0$ for all $x \in \mathbb{R}$) if and only if F is continuous on \mathbb{R} . On the other hand, F is discontinuous at $x \in \mathbb{R}$ if and only if $\mu\{x\} > 0$, so that μ has an *atom* at x . So μ is a *discrete measure* (recall that this means that μ has countable support) if and only if F is a step function.

Suppose again that F is a distribution function and μ is the positive measure on $(\mathbb{R}, \mathcal{R})$ associated with F . If $a \in \mathbb{R}$ then

1. $\mu(a, \infty) = F(\infty) - F(a)$
2. $\mu[a, \infty) = F(\infty) - F(a^-)$
3. $\mu(-\infty, a] = F(a) - F(-\infty)$
4. $\mu(-\infty, a) = F(a^-) - F(-\infty)$
5. $\mu(\mathbb{R}) = F(\infty) - F(-\infty)$

Proof

The proofs, as before, just use the continuity theorems. Suppose that (x_1, x_2, \dots) is a sequence of distinct points in \mathbb{R}

1. If $x_n \uparrow \infty$ as $n \rightarrow \infty$ then $(a, x_n] \uparrow (a, \infty)$ so $\mu(a, x_n] \uparrow \mu(a, \infty)$ as $n \rightarrow \infty$. But also $\mu(a, x_n] = F(x_n) - F(a) \rightarrow F(\infty) - F(a)$ as $n \rightarrow \infty$
2. Similarly, if $x_n \uparrow \infty$ as $n \rightarrow \infty$ then $[a, x_n] \uparrow [a, \infty)$ so $\mu[a, x_n] \uparrow \mu[a, \infty)$ as $n \rightarrow \infty$. But also $\mu[a, x_n] = F(x_n) - F(a^-) \rightarrow F(\infty) - F(a^-)$ as $n \rightarrow \infty$
3. If $x_n \downarrow -\infty$ as $n \rightarrow \infty$ then $(x_n, a] \uparrow (-\infty, a]$ so $\mu(x_n, a] \uparrow \mu(-\infty, a]$ as $n \rightarrow \infty$. But also $\mu(x_n, a] = F(a) - F(x_n) \rightarrow F(a) - F(-\infty)$ as $n \rightarrow \infty$
4. Similarly, if $x_n \downarrow -\infty$ as $n \rightarrow \infty$ then $(x_n, a) \uparrow (-\infty, a)$ so $\mu(x_n, a) \uparrow \mu(-\infty, a)$ as $n \rightarrow \infty$. But also $\mu(x_n, a) = F(a^-) - F(x_n) \rightarrow F(a^-) - F(-\infty)$ as $n \rightarrow \infty$
5. $\mu(\mathbb{R}) = \mu(-\infty, 0] + \mu(0, \infty) = [F(0) - F(-\infty)] + [F(\infty) - F(0)] = F(\infty) - F(-\infty)$

Distribution Functions on $[0, \infty)$

Positive measures and distribution functions on $[0, \infty)$ are particularly important in renewal theory and Poisson processes, because they model random times.

The discrete case. Suppose that G is discrete, so that there exists a countable set $C \subset [0, \infty)$ with $G(C^c) = 0$. Let $g(t) = G\{t\}$ for $t \in C$ so that g is the density function of G with respect to counting measure on C . If $u : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded then

$$\int_0^t u(s) dG(s) = \sum_{s \in C \cap [0, t]} u(s)g(s) \quad (3.9.7)$$



Figure 3.9.1: A discrete measure

In the discrete case, the distribution is often *arithmetic*. Recall that this means that the countable set C is of the form $\{nd : n \in \mathbb{N}\}$ for some $d \in (0, \infty)$. In the following results,

The continuous case. Suppose that G is absolutely continuous with respect to Lebesgue measure on $[0, \infty)$ with density function $g : [0, \infty) \rightarrow [0, \infty)$. If $u : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded then

$$\int_0^t u(s) dG(s) = \int_0^t u(s)g(s) ds \quad (3.9.8)$$



Figure 3.9.2: A continuous measure

The mixed case. Suppose that there exists a countable set $C \subset [0, \infty)$ with $G(C) > 0$ and $G(C^c) > 0$, and that G restricted to subsets of C^c is absolutely continuous with respect to Lebesgue measure. Let $g(t) = G'\{t\}$ for $t \in C$ and let h be a density with respect to Lebesgue measure of G restricted to subsets of C^c . If $u : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded then,

$$\int_0^t u(s) dG(s) = \sum_{s \in C \cap [0, t]} u(s)g(s) + \int_0^t u(s)h(s) ds \quad (3.9.9)$$



Figure 3.9.3: A mixed measure

The three special cases do not exhaust the possibilities, but are by far the most common cases in applied problems.

This page titled [3.9: General Distribution Functions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.10: The Integral With Respect to a Measure

Probability density functions have very different interpretations for discrete distributions as opposed to continuous distributions. For a discrete distribution, the probability of an event is computed by summing the density function over the outcomes in the event, while for a continuous distribution, the probability is computed by integrating the density function over the outcomes. For a mixed distributions, we have partial discrete and continuous density functions and the probability of an event is computed by summing and integrating. The various types of density functions can unified under a general theory of integration, which is the subject of this section. This theory has enormous importance in probability, far beyond just density functions. Expected value, which we consider in the next chapter, can be interpreted as an integral with respect to a probability measure. Beyond probability, the general theory of integration is of fundamental importance in many areas of mathematics.

Basic Theory

Definitions

Our starting point is a measure space (S, \mathcal{S}, μ) . That is, S is a set, \mathcal{S} is a σ -algebra of subsets of S , and μ is a positive measure on \mathcal{S} . As usual, the most important special cases are

- *Euclidean space*: $S = \mathbb{R}^n$ for some $n \in \mathbb{N}_+$, $\mathcal{S} = \mathcal{R}_n$, the σ -algebra of Lebesgue measurable subsets of \mathbb{R}^n , and $\mu = \lambda_n$, standard n -dimensional Lebesgue measure.
- *Discrete space*: S is a countable set, \mathcal{S} is the collection of all subsets of S , and $\mu = \#$, counting measure.
- *Probability space*: S is the set of outcomes of a random experiment, \mathcal{S} is the σ -algebra of events, and $\mu = \mathbb{P}$, a probability measure.

The following definition reflects the fact that in measure theory, sets of measure 0 are often considered unimportant.

Consider a statement with $x \in S$ as a free variable. Technically such a statement is a *predicate* on S . Suppose that $A \in \mathcal{S}$.

1. The statement holds *on* A if it is true for every $x \in A$.
2. The statement holds *almost everywhere on* A (with respect to μ) if there exists $B \in \mathcal{S}$ with $B \subseteq A$ such that the statement holds on B and $\mu(A \setminus B) = 0$.

A typical statement that we have in mind is an equation or an inequality with $x \in S$ as a free variable. Our goal is to define the *integral* of certain measurable functions $f : S \rightarrow \mathbb{R}$, with respect to the measure μ . The integral may exist as a number in \mathbb{R} (in which case we say that f is *integrable*), or may exist as ∞ or $-\infty$, or may not exist at all. When it exists, the integral is denoted variously by

$$\int_S f d\mu, \int_S f(x) d\mu(x), \int_S f(x) \mu(dx) \quad (3.10.1)$$

We will use the first two.

Since the set of extended real numbers $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, \infty\}$ plays an important role in the theory, we need to recall the arithmetic of ∞ and $-\infty$. Here are the conventions that are appropriate for integration:

Arithmetic on \mathbb{R}^*

1. If $a \in (0, \infty]$ then $a \cdot \infty = \infty$ and $a \cdot (-\infty) = -\infty$
2. If $a \in [-\infty, 0)$ then $a \cdot \infty = -\infty$ and $a \cdot (-\infty) = \infty$
3. $0 \cdot \infty = 0$ and $0 \cdot (-\infty) = 0$
4. If $a \in \mathbb{R}$ then $a + \infty = \infty$ and $a + (-\infty) = -\infty$
5. $\infty + \infty = \infty$
6. $-\infty + (-\infty) = -\infty$

However, $\infty - \infty$ is not defined (because it does not make consistent sense) and we must be careful never to produce this indeterminate form. You might recall from calculus that $0 \cdot \infty$ is also an indeterminate form. However, for the theory of integration, the convention that $0 \cdot \infty = 0$ is convenient and consistent. In terms of order of course, $-\infty < a < \infty$ for $a \in \mathbb{R}$.

We also need to extend topology and measure to \mathbb{R}^* . In terms of the first, $(a, \infty]$ is an open neighborhood of ∞ and $[-\infty, a)$ is an open neighborhood of $-\infty$ for every $a \in \mathbb{R}$. This ensures that if $x_n \in \mathbb{R}$ for $n \in \mathbb{N}_+$ then $x_n \rightarrow \infty$ or $x_n \rightarrow -\infty$ as $n \rightarrow \infty$ has its usual calculus meaning. Technically this topology results in the *two-point compactification* of \mathbb{R} . Now we can give \mathbb{R}^* the Borel σ -algebra \mathcal{R}^* , that is, the σ -algebra generated by the topology. Basically, this simply means that if $A \in \mathcal{R}$ then $A \cup \{\infty\}$, $A \cup \{-\infty\}$, and $A \cup \{-\infty, \infty\}$ are all in \mathcal{R}^* .

Desired Properties

As motivation for the definition, every version of integration should satisfy some basic properties. First, the integral of the indicator function of a measurable set should simply be the size of the set, as measured by μ . This gives our first definition:

If $A \in \mathcal{S}$ then $\int_S \mathbf{1}_A d\mu = \mu(A)$.

This definition hints at the intimate relationship between measure and integration. We will construct the integral from the measure μ in this section, but this first property shows that if we started with the integral, we could recover the measure. This property also shows why we need ∞ as a possible value of the integral, and coupled with some of the properties below, why $-\infty$ is also needed. Here is a simple corollary of our first definition.

$$\int_S 0 d\mu = 0$$

Proof

Note that $\int_S 0 d\mu = \int_S \mathbf{1}_\emptyset d\mu = \mu(\emptyset) = 0$.

We give three more essential properties that we want. First are the *linearity properties* in two parts—part (a) is the *additive property* and part (b) is the *scaling property*.

If $f, g : S \rightarrow \mathbb{R}$ are measurable functions whose integrals exist, and $c \in \mathbb{R}$, then

1. $\int_S (f + g) d\mu = \int_S f d\mu + \int_S g d\mu$ as long as the right side is not of the form $\infty - \infty$
2. $\int_S cf d\mu = c \int_S f d\mu$.

The additive property almost implies the scaling property

The steps below do not constitute a proof because questions of the existence of the integrals are ignored and because the limit interchange in the last step is not justified. Still, the argument shows the close relationship between the additive property and the scaling property.

1. If $n \in \mathbb{N}_+$, then by (a) and induction, $\int_S nf d\mu = n \int_S f d\mu$.
2. From step (1), if $n \in \mathbb{N}_+$ then $\int_S f d\mu = \int_S n \frac{1}{n} f d\mu = n \int_S \frac{1}{n} f d\mu$ so $\int_S \frac{1}{n} f d\mu = \frac{1}{n} \int_S f d\mu$.
3. If $m, n \in \mathbb{N}_+$ then from steps (1) and (2) $\int_S \frac{m}{n} f d\mu = m \int_S \frac{1}{n} f d\mu = \frac{m}{n} \int_S f d\mu$.
4. $0 = \int_S 0 d\mu = \int_S (f - f) d\mu = \int_S f d\mu + \int_S -f d\mu$ so $\int_S -f d\mu = - \int_S f d\mu$.
5. By steps (3) and (4), $\int_S cf d\mu = c \int_S f d\mu$ for every $c \in \mathbb{Q}$ (the set of rational real numbers).
6. If $c \in \mathbb{R}$ there exists $c_n \in \mathbb{Q}$ for $n \in \mathbb{N}_+$ with $c_n \rightarrow c$ as $n \rightarrow \infty$. By step (5), $\int_S c_n f d\mu = c_n \int_S f d\mu$.
7. Taking limits in step (6) suggests $\int_S cf d\mu = c \int_S f d\mu$.

To be more explicit, we want the additivity property (a) to hold if at least one of the integrals on the right is finite, or if both are ∞ or if both are $-\infty$. What is ruled out are the two cases where one integral is ∞ and the other is $-\infty$, and this is what is meant by the indeterminate form $\infty - \infty$. Our next essential properties are the *order properties*, again in two parts—part (a) is the *positive property* and part (b) is the *increasing property*.

Suppose that $f, g : S \rightarrow \mathbb{R}$ are measurable.

1. If $f \geq 0$ on S then $\int_S f d\mu \geq 0$.
2. If the integrals of f and g exist and $f \leq g$ on S then $\int_S f d\mu \leq \int_S g d\mu$

The positive property and the additive property imply the increasing property

Implicit in part (a) is that the integral of a nonnegative, measurable function always exists in $[0, \infty]$. Suppose that the integrals of f and g exist and $f \leq g$ on S . Then $g - f \geq 0$ on S and $g = f + (g - f)$. If $\int_S f d\mu = -\infty$, then trivially, $\int_S f d\mu \leq \int_S g d\mu$. Otherwise, by the additivity property,

$$\int_S g d\mu = \int_S f d\mu + \int_S (g - f) d\mu \quad (3.10.2)$$

But $\int_S (g - f) d\mu \geq 0$ (so in particular the right side is not $-\infty + \infty$), and hence $\int_S g d\mu \geq \int_S f d\mu$

Our last essential property is perhaps the least intuitive, but is a type of *continuity property* of integration, and is closely related to the continuity property of positive measure. The official name is the *monotone convergence theorem*.

Suppose that $f_n : S \rightarrow [0, \infty)$ is measurable for $n \in \mathbb{N}_+$ and that f_n is increasing in n . Then

$$\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.10.3)$$

Note that since f_n is increasing in n , $\lim_{n \rightarrow \infty} f_n(x)$ exists in $\mathbb{R} \cup \{\infty\}$ for each $x \in \mathbb{R}$ (and the limit defines a measurable function). This property shows that it is sometimes convenient to allow nonnegative functions to take the value ∞ . Note also that by the [increasing property](#), $\int_S f_n d\mu$ is increasing in n and hence also has a limit in $\mathbb{R} \cup \{\infty\}$.

To see the connection with measure, suppose that (A_1, A_2, \dots) is an increasing sequence of sets in \mathcal{S} , and let $A = \bigcup_{i=1}^{\infty} A_i$. Note that $\mathbf{1}_{A_n}$ is increasing in $n \in \mathbb{N}_+$ and $\mathbf{1}_{A_n} \rightarrow \mathbf{1}_A$ as $n \rightarrow \infty$. For this reason, the union A is sometimes called the *limit* of A_n as $n \rightarrow \infty$. The continuity theorem of positive measure states that $\mu(A_n) \rightarrow \mu(A)$ as $n \rightarrow \infty$. Equivalently, $\int_S \mathbf{1}_{A_n} d\mu \rightarrow \int_S \mathbf{1}_A d\mu$ as $n \rightarrow \infty$, so the continuity theorem of positive measure is a special case of the monotone convergence theorem.

Armed with the properties that we want, the definition of the integral is fairly straightforward, and proceeds in stages. We give the definition successively for

1. Nonnegative simple functions
2. Nonnegative measurable functions
3. Measurable real-valued functions

Of course, each definition should agree with the previous one on the functions that are in both collections.

Simple Functions

A *simple function* on S is simply a measurable, real-valued function with finite range. Simple functions are usually expressed as linear combinations of indicator functions.

Representations of simple functions

1. Suppose that I is a finite index set, $a_i \in \mathbb{R}$ for each $i \in I$, and $\{A_i : i \in I\}$ is a collection of sets in \mathcal{S} that partition S . Then $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ is a simple function. Expressing a simple function in this form is a *representation* of f .
2. A simple function f has a unique representation as $f = \sum_{j \in J} b_j \mathbf{1}_{B_j}$ where J is a finite index set, $\{b_j : j \in J\}$ is a set of distinct real numbers, and $\{B_j : j \in J\}$ is a collection of nonempty sets in \mathcal{S} that partition S . This representation is known as the *canonical representation*.

Proof

1. Note that f is measurable since $A_i \in \mathcal{S}$ for each $i \in I$. Also f has finite range since I is finite. Specifically, the range of f consists of the distinct a_i for $i \in I$ with $A_i \neq \emptyset$.
2. Suppose that f is simple. Let $\{b_j : j \in J\}$ denote the (distinct) values in the range of f and let $B_j = f^{-1}\{b_j\}$ for $j \in J$. Then J is finite, $\{B_j : j \in J\}$ is a collection of nonempty sets in \mathcal{S} that partition S , and $f = \sum_{j \in J} b_j \mathbf{1}_{B_j}$. Conversely, suppose that f has a representation of this form. Then $\{b_j : j \in J\}$ is the range of f and $B_j = f^{-1}\{b_j\}$ so the representation is unique.

You might wonder why we don't just always use the canonical representation for simple functions. The problem is that even if we start with canonical representations, when we combine simple functions in various ways, the resulting representations may not be canonical. The collection of simple functions is closed under the basic arithmetic operations, and in particular, forms a vector space.

Suppose that f and g are simple functions with representations $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ and $g = \sum_{j \in J} b_j \mathbf{1}_{B_j}$, and that $c \in \mathbb{R}$. Then

1. $f + g$ is simple, with representation $f + g = \sum_{(i,j) \in I \times J} (a_i + b_j) \mathbf{1}_{A_i \cap B_j}$.
2. fg is simple, with representation $fg = \sum_{(i,j) \in I \times J} (a_i b_j) \mathbf{1}_{A_i \cap B_j}$.
3. cf is simple, with representation $cf = \sum_{i \in I} ca_i \mathbf{1}_{A_i}$.

Proof

Since f and g are measurable, so are $f + g$, fg , and cf . Moreover, since f and g have finite range, so do $f + g$, fg , and cf . For the representations in parts (a) and (b), note that $I \times J$ is finite, $\{A_i \cap B_j : (i, j) \in I \times J\}$ is a collection of sets in \mathcal{S} that partition S , and on $A_i \cap B_j$, $f + g = a_i + b_j$ and $fg = a_i b_j$.

As we alluded to earlier, note that even if the representations of f and g are canonical, the representations for $f + g$ and fg may not be. The next result treats composition, and will be important for the change of variables theorem in the next section.

Suppose that (T, \mathcal{T}) is another measurable space, and that $f : S \rightarrow T$ is measurable. If g is a simple function on T with representation $g = \sum_{i \in I} b_i \mathbf{1}_{B_i}$, then $g \circ f$ is a simple function on S with representation $g \circ f = \sum_{i \in I} b_i \mathbf{1}_{f^{-1}(B_i)}$.

Proof

Recall that $g \circ f : S \rightarrow \mathbb{R}$ and $\text{range}(g \circ f) \subseteq \text{range}(g)$ so $g \circ f$ has finite range. f is measurable, and inverse images preserve all set operations, so $\{f^{-1}(B_i) : i \in I\}$ is a measurable partition of S . Finally, if $x \in f^{-1}(B_i)$ then $f(x) \in B_i$ so $g[f(x)] = b_i$.

Given the definition of the integral of an indicator function in (3) and that we want the linearity property (5) to hold, there is no question as to how we should define the integral of a nonnegative simple function.

Suppose that f is a nonnegative simple function, with the representation $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ where $a_i \geq 0$ for $i \in I$. We define

$$\int_S f d\mu = \sum_{i \in I} a_i \mu(A_i) \quad (3.10.4)$$

The definition is consistent

Consistency refers to the fact that a simple function can have more than one representation as a linear combination of indicator functions, and hence we must show that all such representations lead to the same value for the integral. Let $\{b_j : j \in J\}$ denote the set of distinct elements among the numbers a_i where $i \in I$ and $A_i \neq \emptyset$. For $j \in J$, let $I_j = \{i \in I : a_i = b_j\}$ and let $B_j = \bigcup_{i \in I_j} A_i$. Thus, $f = \sum_{j \in J} b_j \mathbf{1}_{B_j}$, and this is the canonical representation. Note that

$$\sum_{i \in I} a_i \mu(A_i) = \sum_{j \in J} \sum_{i \in I_j} a_i \mu(A_i) = \sum_{j \in J} b_j \sum_{i \in I_j} \mu(A_i) = \sum_{j \in J} b_j \mu(B_j) \quad (3.10.5)$$

The first sum is the integral defined in terms of the general representation $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ while the last sum is the integral defined in terms of the unique canonical representation $f = \sum_{j \in J} b_j \mathbf{1}_{B_j}$. Thus, any representation of a simple function f leads to the same value for the integral.

Note that if f is a nonnegative simple function, then $\int_S f d\mu$ exists in $[0, \infty]$, so the [order properties](#) holds. We next show that the [linearity properties](#) are satisfied for nonnegative simple functions.

Suppose that f and g are nonnegative simple functions, and that $c \in [0, \infty)$. Then

1. $\int_S (f+g) d\mu = \int_S f d\mu + \int_S g d\mu$
2. $\int_S cf d\mu = c \int_S f d\mu$

Proof

Suppose that f and g are nonnegative simple functions with the representations $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ and $g = \sum_{j \in J} b_j \mathbf{1}_{B_j}$. Thus $a_i \geq 0$ for $i \in I$, $b_j \geq 0$ for $j \in J$, and $\int_S f d\mu = \sum_{i \in I} a_i \mu(A_i)$ and $\int_S g d\mu = \sum_{j \in J} b_j \mu(B_j)$.

1. As noted above, $f+g$ has the representation

$$f+g = \sum_{(i,j) \in I \times J} (a_i + b_j) \mathbf{1}_{A_i \cap B_j} \quad (3.10.6)$$

Note that $\{A_i \cap B_j : j \in J\}$ is a partition of A_i for each $i \in I$, and similarly $\{A_i \cap B_j : i \in I\}$ is a partition of B_j for each $j \in J$. Hence

$$\int_S (f+g) d\mu = \sum_{(i,j) \in I \times J} (a_i + b_j) \mu(A_i \cap B_j) \quad (3.10.7)$$

$$= \sum_{i \in I} \sum_{j \in J} a_i \mu(A_i \cap B_j) + \sum_{j \in J} \sum_{i \in I} b_j \mu(A_i \cap B_j) \quad (3.10.8)$$

$$= \sum_{i \in I} a_i \mu(A_i \cap B) + \sum_{j \in J} b_j \mu(B_j \cap A) \quad (3.10.9)$$

$$= \sum_{i \in I} a_i \mu(A_i) + \sum_{j \in J} b_j \mu(B_j) = \int_S f d\mu + \int_S g d\mu \quad (3.10.10)$$

Note that all the terms are nonnegative (although some may be ∞), so there are no problems with rearranging the order of the terms.

2. This part is easier. For $c \in [0, \infty)$, recall that cf has the representation $cf = \sum_{i \in I} ca_i \mathbf{1}_{A_i}$ so

$$\int_S cf d\mu = \sum_{i \in I} ca_i \mu(A_i) = c \sum_{i \in I} a_i \mu(A_i) = c \int_S f d\mu \quad (3.10.11)$$

The [increasing property](#) holds for nonnegative simple functions.

Suppose that f and g are nonnegative simple functions and $f \leq g$ on S . Then $\int_S f d\mu \leq \int_S g d\mu$

Proof

The proof from the [additive property](#) above works. Note that $g-f$ is a nonnegative simple function, and $g = f + (g-f)$. By the additivity property, $\int_S g d\mu = \int_S f d\mu + \int_S (g-f) d\mu \geq \int_S f d\mu$.

Next we give a version of the continuity theorem in (7) for simple functions. It's not completely general, but will be needed for the next subsection where we do prove the general version.

Suppose that f is a nonnegative simple function and that (A_1, A_2, \dots) is an increasing sequence of sets in \mathcal{S} with $A = \bigcup_{n=1}^{\infty} A_n$. then

$$\int_S \mathbf{1}_{A_n} f d\mu \rightarrow \int_S \mathbf{1}_A f d\mu \text{ as } n \rightarrow \infty \quad (3.10.12)$$

Proof

Suppose that f has the representation $f = \sum_{i \in I} b_i \mathbf{1}_{B_i}$. Then $\mathbf{1}_{A_n} f = \sum_{i \in I} b_i \mathbf{1}_{A_n} \mathbf{1}_{B_i} = \sum_{i \in I} b_i \mathbf{1}_{A_n \cap B_i}$ and similarly, $\mathbf{1}_A f = \sum_{i \in I} b_i \mathbf{1}_{A \cap B_i}$. But for each $i \in I$, $B_i \cap A_n$ is increasing in $n \in \mathbb{N}_+$ and $\bigcup_{n=1}^{\infty} (B_i \cap A_n) = B_i \cap A$. By the continuity theorem for positive measures, $\mu(B_i \cap A_n) \rightarrow \mu(B_i \cap A)$ as $n \rightarrow \infty$ for each $i \in I$. Since I is finite,

$$\int_{A_n} f d\mu = \sum_{i \in I} b_i \mu(A_n \cap B_i) \rightarrow \sum_{i \in I} b_i \mu(A \cap B_i) = \int_A f d\mu \text{ as } n \rightarrow \infty \quad (3.10.13)$$

Note that $\mathbf{1}_{A_n} f$ is increasing in $n \in \mathbb{N}_+$ and $\mathbf{1}_{A_n} f \rightarrow \mathbf{1}_A f$ as $n \rightarrow \infty$, so this really is a special case of the monotone convergence theorem.

Nonnegative Functions

Next we will consider nonnegative measurable functions on S . First we note that a function of this type is the limit of nonnegative simple functions.

Suppose that $f : S \rightarrow [0, \infty)$ is measurable. Then there exists an increasing sequence (f_1, f_2, \dots) of nonnegative simple functions with $f_n \rightarrow f$ on S as $n \rightarrow \infty$.

Proof

For $n \in \mathbb{N}_+$ and $k \in \{1, 2, \dots, n2^n\}$ Let $I_{n,k} = [(k-1)/2^n, k/2^n)$ and $I_n = [n, \infty)$. Note that

1. $\{I_{n,k} : k = 1, \dots, n2^n\} \cup \{I_n\}$ is a partition of $[0, \infty)$ for each $n \in \mathbb{N}_+$.
2. $I_{n,k} = I_{n+1, 2k-1} \cup I_{n+1, 2k}$ for $k \in \{1, 2, \dots, n2^n\}$.
3. $I_n = \left(\bigcup_{k=n2^{n+1}+1}^{(n+1)2^{n+1}} I_{n+1, k} \right) \cup I_{n+1}$ for $n \in \mathbb{N}_+$.

Note that the n th partition divides the interval $[0, n)$ into $n2^n$ subintervals of length $1/2^n$. Thus, (b) follows because the $(n+1)$ st partition divides each of the first 2^n intervals of the n th partition in half, and (c) follows because the $(n+1)$ st partition divides the interval $[n, n+1)$ into subintervals of length $1/2^{n+1}$. Now let $A_{n,k} = f^{-1}(I_{n,k})$ and $A_n = f^{-1}(I_n)$ for $n \in \mathbb{N}_+$ and $k \in \{1, 2, \dots, n2^n\}$. Since inverse images preserve all set operations, (a), (b), and (c) hold with A replacing I everywhere, and S replacing $[0, \infty)$ in (a). Moreover, since f is measurable, $A_n \in \mathcal{S}$ and $A_{n,k} \in \mathcal{S}$ for each n and k . Now, define

$$f_n = \sum_{k=1}^{n2^n} \frac{k-1}{2^n} \mathbf{1}_{A_{n,k}} + n \mathbf{1}_{A_n} \quad (3.10.14)$$

Then f_n is a simple function and $0 \leq f_n \leq f$ for each $n \in \mathbb{N}_+$. To show convergence, fix $x \in S$. If $n > f(x)$ then $|f(x) - f_n(x)| \leq 2^{-n}$ and hence $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$. All that remains is to show that f_n is increasing in n . Let $x \in S$ and $n \in \mathbb{N}_+$. If $x \in A_{n,k}$ for some $k \in \{1, 2, \dots, n2^n\}$, then $f_n(x) = (k-1)/2^n$. But either $f_{n+1}(x) = (2k-2)/2^{n+1}$ or $f_{n+1}(x) = (2k-1)/2^{n+1}$. If $x \in A_n$ then $f_n(x) = n$. But either $f_{n+1}(x) = (k-1)/2^{n+1}$ for some $k \in \{n2^{n+1}+1, \dots, (n+1)2^{n+1}\}$ or $f_{n+1}(x) = n+1$. In all cases, $f_{n+1}(x) \geq f_n(x)$.

The last result points the way towards the definition of the integral of a measurable function $f : S \rightarrow [0, \infty)$ in terms of the integrals of simple functions. If g is a nonnegative simple function with $g \leq f$, then by the [order property](#), we need $\int_S g d\mu \leq \int_S f d\mu$. On the other hand, there exists a sequence of nonnegative simple function converging to f . Thus the [continuity property](#) suggests the following definition:

If $f : S \rightarrow [0, \infty)$ is measurable, define

$$\int_S f d\mu = \sup \left\{ \int_S g d\mu : g \text{ is simple and } 0 \leq g \leq f \right\} \quad (3.10.15)$$

Note that $\int_S f d\mu$ exists in $[0, \infty]$ so the [positive property](#) holds. Note also that if f is simple, the new definition agrees with the old one. As always, we need to establish the essential properties. First, the [increasing property](#) holds.

If $f, g : S \rightarrow [0, \infty)$ are measurable and $f \leq g$ on S then $\int_S f d\mu \leq \int_S g d\mu$.

Proof

Note that $\{h : h \text{ is simple and } 0 \leq h \leq f\} \subseteq \{h : h \text{ is simple and } 0 \leq h \leq g\}$. therefore

$$\int_S f d\mu = \sup \left\{ \int_S h d\mu : h \text{ is simple and } 0 \leq h \leq f \right\} \leq \sup \left\{ \int_S h d\mu : h \text{ is simple and } 0 \leq h \leq g \right\} = \int_S g d\mu \quad (3.10.16)$$

We can now prove the [continuity property](#) known as the *monotone convergence theorem* in full generality.

Suppose that $f_n : S \rightarrow [0, \infty)$ is measurable for $n \in \mathbb{N}_+$ and that f_n is increasing in n . Then

$$\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.10.17)$$

Proof

Let $f = \lim_{n \rightarrow \infty} f_n$. By the order property, note that $\int_S f_n d\mu$ is increasing in $n \in \mathbb{N}_+$ and hence has a limit in \mathbb{R}^* , which we will denote by c . Note that $f_n \leq f$ on S for $n \in \mathbb{N}_+$, so by the order property again, $\int_S f_n d\mu \leq \int_S f d\mu$ for $n \in \mathbb{N}_+$. Letting $n \rightarrow \infty$ gives $c \leq \int_S f d\mu$. To show that $c \geq \int_S f d\mu$ we need to show that $c \geq \int_S g d\mu$ for every simple function g with $0 \leq g \leq f$. Fix $a \in (0, 1)$ and let $A_n = \{x \in S : f_n(x) \geq ag(x)\}$. Since f_n is increasing in n , $A_n \subseteq A_{n+1}$. Moreover, since $f_n \rightarrow f$ as $n \rightarrow \infty$ on S and $g \leq f$ on S , $\bigcup_{n=1}^{\infty} A_n = S$. But by definition, $ag \leq f_n$ on A_n so

$$\alpha \int_S \mathbf{1}_{A_n} g d\mu = \int_S \alpha \mathbf{1}_{A_n} g d\mu \leq \int_S \mathbf{1}_{A_n} f_n d\mu \leq \int_S f_n d\mu \quad (3.10.18)$$

Letting $n \rightarrow \infty$ in the extreme parts of the displayed inequality and using the version of the monotone convergence theorem for [simple functions](#), we have $a \int_S g d\mu \leq c$ for every $a \in (0, 1)$. Finally, letting $a \uparrow 1$ gives $\int_S g d\mu \leq c$

If $f : S \rightarrow [0, \infty)$ is measurable, then by the [theorem above](#), there exists an increasing sequence (f_1, f_2, \dots) of simple functions with $f_n \rightarrow f$ as $n \rightarrow \infty$. By the monotone convergence theorem in (18), $\int_S f_n d\mu \rightarrow \int_S f d\mu$ as $n \rightarrow \infty$. These two facts can be used to establish other properties of the integral of a nonnegative function based on our knowledge that the properties hold for simple functions. This type of argument is known as *bootstrapping*. We use bootstrapping to show that the [linearity properties](#) hold:

If $f, g : S \rightarrow [0, \infty)$ are measurable and $c \in [0, \infty)$, then

1. $\int_S (f + g) d\mu = \int_S f d\mu + \int_S g d\mu$
2. $\int_S cf d\mu = c \int_S f d\mu$

Proof

1. Let (f_1, f_2, \dots) and (g_1, g_2, \dots) be increasing sequences of nonnegative simple functions with $f_n \rightarrow f$ and $g_n \rightarrow g$ as $n \rightarrow \infty$. Then $(f_1 + g_1, f_2 + g_2, \dots)$ is also an increasing sequence of simple functions, and $f_n + g_n \rightarrow f + g$ as $n \rightarrow \infty$. By the monotone convergence theorem, $\int_S f_n d\mu \rightarrow \int_S f d\mu$, $\int_S g_n d\mu \rightarrow \int_S g d\mu$, and $\int_S (f_n + g_n) d\mu \rightarrow \int_S (f + g) d\mu$ as $n \rightarrow \infty$. But $\int_S (f_n + g_n) d\mu = \int_S f_n d\mu + \int_S g_n d\mu$ for each $n \in \mathbb{N}_+$ so taking limits gives $\int_S (f + g) d\mu = \int_S f d\mu + \int_S g d\mu$.
2. Similarly, (cf_1, cf_2, \dots) is an increasing sequence of nonnegative simple functions with $cf_n \rightarrow cf$ as $n \rightarrow \infty$. Again, by the MCT, $\int_S cf_n d\mu \rightarrow \int_S cf d\mu$ and $\int_S cf_n d\mu = c \int_S f_n d\mu$ so taking limits gives $\int_S cf d\mu = c \int_S f d\mu$.

General Functions

Our final step is to define the integral of a measurable function $f : S \rightarrow \mathbb{R}$. First, recall the positive and negative parts of $x \in \mathbb{R}$:

$$x^+ = \max\{x, 0\}, \quad x^- = \max\{-x, 0\} \quad (3.10.19)$$

Note that $x^+ \geq 0$, $x^- \geq 0$, $x = x^+ - x^-$, and $|x| = x^+ + x^-$. Given that we want the integral to have the linearity properties in (5), there is no question as to how we should define the integral of f in terms of the integrals of f^+ and f^- , which being nonnegative, are defined by the previous subsection.

If $f : S \rightarrow \mathbb{R}$ is measurable, we define

$$\int_S f d\mu = \int_S f^+ d\mu - \int_S f^- d\mu \quad (3.10.20)$$

assuming that at least one of the integrals on the right is finite. If both are finite, then f is said to be *integrable*.

Assuming that either the integral of the positive part or the integral of the negative part is finite ensures that we do not get the dreaded indeterminate form $\infty - \infty$.

Suppose that $f : S \rightarrow \mathbb{R}$ is measurable. Then f is integrable if and only if $\int_S |f| d\mu < \infty$.

Proof

Suppose that f is integrable. Recall that $|f| = f^+ + f^-$. By the additive property for nonnegative functions, $\int_S |f| d\mu = \int_S f^+ d\mu + \int_S f^- d\mu < \infty$. Conversely, suppose that $\int_S |f| d\mu < \infty$. Then $f^+ \leq |f|$ and $f^- \leq |f|$ so by the increasing property for nonnegative functions, $\int_S f^+ d\mu \leq \int_S |f| d\mu < \infty$ and $\int_S f^- d\mu \leq \int_S |f| d\mu < \infty$.

Note that if f is nonnegative, then our new definition agrees with our old one, since $f^+ = f$ and $f^- = 0$. For simple functions the integral has the same basic form as for nonnegative simple functions:

Suppose that f is a simple function with the representation $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$. Then

$$\int_S f d\mu = \sum_{i \in I} a_i \mu(A_i) \quad (3.10.21)$$

assuming that the sum does not have both ∞ and $-\infty$ terms.

Proof

Note that f^+ and f^- are also simple, with the representations $f^+ = \sum_{i \in I} a_i^+ \mathbf{1}_{A_i}$ and $f^- = \sum_{i \in I} a_i^- \mathbf{1}_{A_i}$. Hence

$$\int_S f d\mu = \sum_{i \in I} a_i^+ \mu(A_i) - \sum_{i \in I} a_i^- \mu(A_i) \quad (3.10.22)$$

as long as one of the sums is finite. Given that this is the case, we can recombine the sums to get

$$\int_S f d\mu = \sum_{i \in I} a_i \mu(A_i) \quad (3.10.23)$$

Once again, we need to establish the essential properties. Our first result is an intermediate step towards linearity.

If $f, g: S \rightarrow [0, \infty)$ are measurable then $\int_S (f - g) d\mu = \int_S f d\mu - \int_S g d\mu$ as long as at least one of the integrals on the right is finite.

Proof

We take cases. Suppose first that $\int_S f d\mu < \infty$ and $\int_S g d\mu < \infty$. Note that $(f - g)^+ \leq f$ and $(f - g)^- \leq g$. By the increasing property for nonnegative functions, $\int_S (f - g)^+ d\mu \leq \int_S f d\mu < \infty$ and $\int_S (f - g)^- d\mu \leq \int_S g d\mu < \infty$. Thus $f - g$ is integrable. Next we have $f - g = (f - g)^+ - (f - g)^-$ and therefore $f + (f - g)^- = g + (f - g)^+$. All four of the functions in the last equation are nonnegative, and therefore by additivity property for nonnegative functions, we have

$$\int_S f d\mu + \int_S (f - g)^- d\mu = \int_S g d\mu + \int_S (f - g)^+ d\mu \quad (3.10.24)$$

All of these integrals are finite, and hence

$$\int_S (f - g) d\mu = \int_S (f - g)^+ d\mu - \int_S (f - g)^- d\mu = \int_S f d\mu - \int_S g d\mu \quad (3.10.25)$$

Next suppose that $\int_S f d\mu = \infty$ and $\int_S g d\mu < \infty$. Then $f - g \leq (f - g)^+$ and hence $f \leq (f - g)^+ + g$. Using the additivity and increasing properties for nonnegative functions, we have $\infty = \int_S f d\mu \leq \int_S (f - g)^+ d\mu + \int_S g d\mu$. Since $\int_S g d\mu < \infty$ we must have $\int_S (f - g)^+ d\mu = \infty$. On the other hand, $(f - g)^- \leq g$ so $\int_S (f - g)^- d\mu \leq \int_S g d\mu < \infty$. Hence $\int_S (f - g) d\mu = \infty = \int_S f d\mu - \int_S g d\mu$.

Finally, suppose that $\int_S f d\mu < \infty$ and $\int_S g d\mu = \infty$. By the argument in the last paragraph, we have $\int_S (g - f)^+ d\mu = \infty$ and $\int_S (g - f)^- d\mu < \infty$. Equivalently, $\int_S (f - g)^+ d\mu < \infty$ and $\int_S (f - g)^- d\mu = \infty$. Hence $\int_S (f - g) d\mu = -\infty = \int_S f d\mu - \int_S g d\mu$.

We finally have the [linearity properties](#) in full generality.

If $f, g: S \rightarrow \mathbb{R}$ are measurable functions whose integrals exist, and $c \in \mathbb{R}$, then

1. $\int_S (f + g) d\mu = \int_S f d\mu + \int_S g d\mu$ as long as the right side is not of the form $\infty - \infty$.
2. $\int_S cf d\mu = c \int_S f d\mu$

Proof

1. Note that $f + g = (f^+ - f^-) + (g^+ - g^-) = (f^+ + g^+) - (f^- + g^-)$ and the two functions in parentheses in the last expression are nonnegative. By the previous lemma and the additivity property for nonnegative functions, we have

$$\int_S (f + g) d\mu = \int_S (f^+ + g^+) d\mu - \int_S (f^- + g^-) d\mu \quad (3.10.26)$$

$$= \left(\int_S f^+ d\mu + \int_S g^+ d\mu \right) - \left(\int_S f^- d\mu + \int_S g^- d\mu \right) \quad (3.10.27)$$

assuming that either both integrals in the first parentheses are finite or both integrals in the second parentheses are finite. In either case, we can group the terms (without worrying about the dreaded $\infty - \infty$) to get

$$\int_S (f+g) d\mu = \left(\int_S f^+ d\mu - \int_S f^- d\mu \right) + \left(\int_S g^+ d\mu - \int_S g^- d\mu \right) = \int_S f d\mu + \int_S g d\mu \quad (3.10.28)$$

2. Note that if $c \geq 0$ then $(cf)^+ = cf^+$ and $(cf)^- = cf^-$. Hence using the scaling property for nonnegative functions,

$$\int_S cf d\mu = \int_S (cf)^+ d\mu - \int_S (cf)^- d\mu = \int_S cf^+ d\mu - \int_S cf^- d\mu = c \int_S f^+ d\mu - c \int_S f^- d\mu = c \int_S f d\mu \quad (3.10.29)$$

On the other hand, if $c < 0$, $(cf)^+ = -cf^-$ and $(cf)^- = -cf^+$. Again using the scaling property for nonnegative functions,

$$\int_S cf d\mu = \int_S (cf)^+ d\mu - \int_S (cf)^- d\mu = \int_S -cf^- d\mu - \int_S -cf^+ d\mu = -c \int_S f^- d\mu + c \int_S f^+ d\mu = c \int_S f d\mu \quad (3.10.30)$$

In particular, note that if f and g are integrable, then so are $f+g$ and cf for $c \in \mathbb{R}$. Thus, the set of integrable functions on (S, \mathcal{S}, μ) forms a vector space, which is denoted $\mathcal{L}(S, \mathcal{S}, \mu)$. The \mathcal{L} is in honor of Henri Lebesgue, who first developed the theory. This vector space, and other related ones, will be studied in more detail in the section on function spaces.

We also have the [increasing property](#) in full generality.

If $f, g : S \rightarrow \mathbb{R}$ are measurable functions whose integrals exist, and if $f \leq g$ on S then $\int_S f d\mu \leq \int_S g d\mu$

Proof

We can use the proof based on the additive property from (6). First $g = f + (g - f)$ and $g - f \geq 0$ on S . If $\int_S f d\mu = -\infty$ then trivially, $\int_S f d\mu \leq \int_S g d\mu$. Otherwise $\int_S (g - f) d\mu \geq 0$ and therefore $\int_S g d\mu = \int_S f d\mu + \int_S (g - f) d\mu \geq \int_S f d\mu$.

The Integral Over a Set

Now that we have defined the integral of a measurable function f over all of S , there is a natural extension to the integral of f over a measurable subset

If $f : S \rightarrow \mathbb{R}$ is measurable and $A \in \mathcal{S}$, we define

$$\int_A f d\mu = \int_S \mathbf{1}_A f d\mu \quad (3.10.31)$$

assuming that the integral on the right exists.

If $f : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists and $A \in \mathcal{S}$, then the integral of f over A exists.

Proof

Note that $(\mathbf{1}_A f)^+ = \mathbf{1}_A f^+$ and $(\mathbf{1}_A f)^- = \mathbf{1}_A f^-$. Also $\mathbf{1}_A f^+ \leq f^+$ and $\mathbf{1}_A f^- \leq f^-$. If $\int_S f d\mu$ exists, then either $\int_S f^+ d\mu < \infty$ or $\int_S f^- d\mu < \infty$. By the increasing property, it follows that either $\int_S \mathbf{1}_A f^+ d\mu < \infty$ or $\int_S \mathbf{1}_A f^- d\mu < \infty$, so $\int_A f d\mu$ exists.

On the other hand, it's clearly possible for $\int_A f d\mu$ to exist for some $A \in \mathcal{S}$, but not $\int_S f d\mu$.

We could also simply think of $\int_A f d\mu$ as the integral of a measurable function $f : A \rightarrow \mathbb{R}$ over the measure space $(A, \mathcal{S}_A, \mu_A)$, where $\mathcal{S}_A = \{B \in \mathcal{S} : B \subseteq A\} = \{C \cap A : C \in \mathcal{S}\}$ is the σ -algebra of measurable subsets of A , and where μ_A is the restriction of μ to \mathcal{S}_A . It follows that all of the essential properties hold for integrals over A : the [linearity properties](#), the [order properties](#), and the [monotone convergence theorem](#). The following property is a simple consequence of the general additive property, and is known as *additive property for disjoint domains*.

Suppose that $f : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists, and that $A, B \in \mathcal{S}$ are disjoint. then

$$\int_{A \cup B} f d\mu = \int_A f d\mu + \int_B f d\mu \quad (3.10.32)$$

Proof

Recall that $\mathbf{1}_{A \cup B} = \mathbf{1}_A + \mathbf{1}_B$. Hence by the additive property and the previous result,

$$\int_{A \cup B} f d\mu = \int_S \mathbf{1}_{A \cup B} f d\mu = \int_S (\mathbf{1}_A f + \mathbf{1}_B f) d\mu = \int_S \mathbf{1}_A f d\mu + \int_S \mathbf{1}_B f d\mu = \int_A f d\mu + \int_B f d\mu \quad (3.10.33)$$

By induction, the additive property holds for a finite collection of disjoint domains. The extension to a countably infinite collection of disjoint domains will be considered in the next section on properties of the integral.

Special Cases

Discrete Spaces

Recall again that the measure space $(S, \mathcal{S}, \#)$ is *discrete* if S is countable, \mathcal{S} is the collection of all subsets of S , and $\#$ is counting measure on \mathcal{S} . Thus all functions $f : S \rightarrow \mathbb{R}$ are measurable, and as we will see, integrals with respect to $\#$ are simply sums.

If $f : S \rightarrow \mathbb{R}$ then

$$\int_S f d\# = \sum_{x \in S} f(x) \quad (3.10.34)$$

as long as either the sum of the positive terms or the sum of the negative terms is finite.

Proof

The proof is a bootstrapping argument.

1. Suppose first that S is finite. In this case, every function $f : S \rightarrow \mathbb{R}$ is simple and has the representation $f = \sum_{x \in S} f(x) \mathbf{1}_x$ where $\mathbf{1}_x$ is an abbreviation of $\mathbf{1}_{\{x\}}$. Thus the result follows from the definition of the integral.
2. Next suppose that S is countable infinite and $f : S \rightarrow [0, \infty)$. Let (A_1, A_2, \dots) be an increasing sequence of finite subsets of S with $\bigcup_{i=1}^{\infty} A_i = S$. Define $f_n = \sum_{x \in A_n} f(x) \mathbf{1}_x$. Then (f_1, f_2, \dots) is an increasing sequence of simple functions with $f_n \rightarrow f$ as $n \rightarrow \infty$. Thus

$$\int_S f d\# = \lim_{n \rightarrow \infty} \int_S f_n d\# = \lim_{n \rightarrow \infty} \sum_{x \in A_n} f(x) \quad (3.10.35)$$

But by definition, the last limit on the right is just $\sum_{x \in S} f(x)$.

3. Finally consider the general case where S is countable and $f : S \rightarrow \mathbb{R}$. In this case the result follows from the definition of the integral as $\int_S f d\# = \int_S f^+ d\# - \int_S f^- d\#$ as long as one of the integrals on the right is finite. By (b), $\int_S f^+ d\#$ is the sum of the positive terms and $-\int_S f^- d\#$ is the sum of the negative terms.

If the sum of the positive terms and the sum of the negative terms are both finite, then f is integrable with respect to $\#$, but the usual term from calculus is that the series $\sum_{x \in S} f(x)$ is *absolutely convergent*. The result will look more familiar in the special case $S = \mathbb{N}_+$. Functions on S are simply sequences, so we can use the more familiar notation a_i rather than $a(i)$ for a function $a : S \rightarrow \mathbb{R}$. Part (b) of the proof (with $A_n = \{1, 2, \dots, n\}$) is just the definition of an infinite series of nonnegative terms as the limit of the partial sums:

$$\sum_{i=1}^{\infty} a_i = \lim_{n \rightarrow \infty} \sum_{i=1}^n a_i \quad (3.10.36)$$

Part (c) of the proof is just the definition of a general infinite series

$$\sum_{i=1}^{\infty} a_i = \sum_{i=1}^{\infty} a_i^+ - \sum_{i=1}^{\infty} a_i^- \quad (3.10.37)$$

as long as one of the series on the right is finite. Again, when both are finite, the series is absolutely convergent. In calculus we also consider *conditionally convergent* series. This means that $\sum_{i=1}^{\infty} a_i^+ = \infty$, $\sum_{i=1}^{\infty} a_i^- = \infty$, but $\lim_{n \rightarrow \infty} \sum_{i=1}^n a_i$ exists in \mathbb{R} . Such series have no place in general integration theory. Also, you may recall that such series are pathological in the sense that, given any number in \mathbb{R}^* , there exists a rearrangement of the terms so that the rearranged series converges to the given number.

The Lebesgue and Riemann Integrals on \mathbb{R}

Consider the one-dimensional Euclidean space $(\mathbb{R}, \mathcal{R}, \lambda)$ where \mathcal{R} is the usual σ -algebra of Lebesgue measurable sets and λ is Lebesgue measure. The theory developed above applies, of course, for the integral $\int_A f d\mu$ of a measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$ over a set $A \in \mathcal{R}$. It's not surprising that in this special case, the theory of integration is referred to as *Lebesgue integration* in honor of our good friend Henri Lebesgue, who first developed the theory.

On the other hand, we already have a theory of integration on \mathbb{R} , namely the Riemann integral of calculus, named for our other good friend Georg Riemann. For a suitable function f and domain A this integral is denoted $\int_A f(x) dx$, as we all remember from calculus. How are the two integrals related? As we will see, the Lebesgue integral generalizes the Riemann integral.

To understand the connection we need to review the definition of the Riemann integral. Consider first the standard case where the domain of integration is a closed, bounded interval. Here are the preliminary definitions that we will need.

Suppose that $f : [a, b] \rightarrow \mathbb{R}$, where $a, b \in \mathbb{R}$ and $a < b$.

1. A *partition* $\mathcal{A} = \{A_i : i \in I\}$ of $[a, b]$ is a finite collection of disjoint *subintervals* whose union is $[a, b]$.

2. The *norm* of a partition \mathcal{A} is $\|\mathcal{A}\| = \max\{\lambda(A_i) : i \in I\}$, the length of the largest subinterval of \mathcal{A} .
3. A set of points $B = \{x_i : i \in I\}$ where $x_i \in A_i$ for each $i \in I$ is said to be *associated* with the partition \mathcal{A} .
4. The *Riemann sum* of f corresponding to a partition \mathcal{A} and a set B associated with \mathcal{A} is

$$R(f, \mathcal{A}, B) = \sum_{i \in I} f(x_i) \lambda(A_i) \quad (3.10.38)$$

Note that the Riemann sum is simply the integral of the simple function $g = \sum_{i \in I} f(x_i) \mathbf{1}_{A_i}$. Moreover, since A_i is an interval for each $i \in I$, g is a *step function*, since it is constant on a finite collection of disjoint intervals. Moreover, again since A_i is an interval for each $i \in I$, $\lambda(A_i)$ is simply the length of the subinterval A_i , so of course measure theory per se is not needed for Riemann integration. Now for the definition from calculus:

f is *Riemann integrable* on $[a, b]$ if there exists $r \in \mathbb{R}$ with the property that for every $\epsilon > 0$ there exists $\delta > 0$ such that if \mathcal{A} is a partition of $[a, b]$ with $\|\mathcal{A}\| < \delta$ then $|r - R(f, \mathcal{A}, B)| < \epsilon$ for every set of points B associated with \mathcal{A} . Then of course we define the integral by

$$\int_a^b f(x) dx = r \quad (3.10.39)$$

Here is our main theorem of this subsection.

If $f : [a, b] \rightarrow \mathbb{R}$ is Riemann integrable on $[a, b]$ then f is Lebesgue integrable on $[a, b]$ and

$$\int_{[a,b]} f d\lambda = \int_a^b f(x) dx \quad (3.10.40)$$

On the other hand, there are lots of functions that are Lebesgue integrable but not Riemann integrable. In fact there are indicator functions of this type, the simplest of functions from the point of view of Lebesgue integration.

Consider the function $\mathbf{1}_{\mathbb{Q}}$ where as usual, \mathbb{Q} is the set of rational number in \mathbb{R} . Then

1. $\int_{\mathbb{R}} \mathbf{1}_{\mathbb{Q}} d\lambda = 0$.
2. $\mathbf{1}_{\mathbb{Q}}$ is not Riemann integrable on any interval $[a, b]$ with $a < b$.

Proof

Part (a) follows from the definition of the Lebesgue integral:

$$\int_{\mathbb{R}} \mathbf{1}_{\mathbb{Q}} d\lambda = \lambda(\mathbb{Q}) = 0 \quad (3.10.41)$$

For part (b), note that there are rational and irrational numbers in every interval of \mathbb{R} of positive length (the rational numbers and the irrational numbers are *dense* in \mathbb{R}). Thus, given any partition $\mathcal{A} = \{A_i : i \in I\}$ of $[a, b]$, no matter how small the norm, there are Riemann sums that are 0 (take $x_i \in A_i$ irrational for each $i \in I$), and Riemann sums that are $b - a$ (take $x_i \in A_i$ rational for each $i \in I$).

The following fundamental theorem completes the picture.

$f : [a, b] \rightarrow \mathbb{R}$ is Riemann integrable on $[a, b]$ if and only if f is bounded on $[a, b]$ and f is continuous almost everywhere on $[a, b]$.

Now that the Riemann integral is defined for a closed bounded interval, it can be extended to other domains.

Extensions of the Riemann integral.

1. If f is defined on $[a, b)$ and Riemann integrable on $[a, t]$ for $a < t < b$, we define $\int_a^b f(x) dx = \lim_{t \uparrow b} \int_a^t f(x) dx$ if the limit exists in \mathbb{R}^* .
2. If f is defined on $(a, b]$ and Riemann integrable on $[t, b]$ for $a < t < b$, we define $\int_a^b f(x) dx = \lim_{t \downarrow a} \int_t^b f(x) dx$ if the limit exists in \mathbb{R}^* .
3. If f is defined on (a, b) , we select $c \in (a, b)$ and define $\int_a^b f(x) dx = \int_a^c f(x) dx + \int_c^b f(x) dx$ if the integrals on the right exist in \mathbb{R}^* by (a) and (b), and are not of the form $\infty - \infty$.
4. If f is defined on $[a, \infty)$ and Riemann integrable on $[a, t]$ for $a < t < \infty$ we define $\int_a^\infty f(x) dx = \lim_{t \rightarrow \infty} \int_a^t f(x) dx$.
5. If f is defined on $(-\infty, b]$ and Riemann integrable on $[t, b]$ for $-\infty < t < b$ we define $\int_{-\infty}^b f(x) dx = \lim_{t \rightarrow -\infty} \int_t^b f(x) dx$ if the limit exists in \mathbb{R}^* .
6. If f is defined on \mathbb{R} we select $c \in \mathbb{R}$ and define $\int_{-\infty}^\infty f(x) dx = \int_{-\infty}^c f(x) dx + \int_c^\infty f(x) dx$ if both integrals on the right exist by (d) and (e), and are not of the form $\infty - \infty$.

7. The integral is defined for a domain that is the union of a finite collection of disjoint intervals by the requirement that the integral be additive over disjoint domains

As another indication of its superiority, note that *none* of these convolutions is necessary for the Lebesgue integral. Once and for all, we have defined $\int_A f(x) dx$ for a general measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$ and a general domain $A \in \mathcal{R}$

The Lebesgue-Stieltjes Integral

Consider again the measurable space $(\mathbb{R}, \mathcal{R})$ where \mathcal{R} is the usual σ -algebra of Lebesgue measurable subsets of \mathbb{R} . Suppose that $F : \mathbb{R} \rightarrow \mathbb{R}$ is a general distribution function, so that by definition, F is increasing and continuous from the right. Recall that the Lebesgue-Stieltjes measure μ associated with F is the unique measure on \mathcal{R} that satisfies

$$\mu(a, b] = F(b) - F(a); \quad a, b \in \mathbb{R}, \quad a < b \quad (3.10.42)$$

Recall that F satisfies some, but not necessarily all of the properties of a probability distribution function. The properties not necessarily satisfied are the normalizing properties

- $F(x) \rightarrow 0$ as $x \rightarrow -\infty$
- $F(x) \rightarrow 1$ as $x \rightarrow \infty$

If F does satisfy these two additional properties, then μ is a probability measure and F its probability distribution function.

The integral with respect to the measure μ is, appropriately enough, referred to as the *Lebesgue-Stieltjes* integral with respect to F , and like the measure, is named for the ubiquitous Henri Lebesgue and for Thomas Stieltjes. In addition to our usual notation $\int_S f d\mu$, the Lebesgue-Stieltjes integral is also denoted $\int_S f dF$ and $\int_S f(x) dF(x)$.

Probability Spaces

Suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space, so that S is the set of outcomes of a random experiment, \mathcal{S} is the σ -algebra of events, and \mathbb{P} the probability measure on the sample space (S, \mathcal{S}) . A measurable, real-valued function X on S is, of course, a real-valued random variable. The integral with respect to \mathbb{P} , if it exists, is the *expected value* of X and is denoted

$$\mathbb{E}(X) = \int_S X d\mathbb{P} \quad (3.10.43)$$

This concept is of fundamental importance in probability theory and is studied in detail in a separate chapter on Expected Value, mostly from an elementary point of view that does not involve abstract integration. However an advanced section treats expected value as an integral over the underlying probability measure, as above.

Suppose next that $(T, \mathcal{T}, \#)$ is a discrete space and that X is a random variable for the experiment, taking values in T . In this case X has a discrete distribution and the probability density function f of X is given by $f(x) = \mathbb{P}(X = x)$ for $x \in T$. More generally,

$$\mathbb{P}(X \in A) = \sum_{x \in A} f(x) = \int_A f d\#, \quad A \subseteq T \quad (3.10.44)$$

On the other hand, suppose that X is a random variable with values in \mathbb{R}^n , where as usual, $(\mathbb{R}^n, \mathcal{R}_n, \lambda_n)$ is n -dimensional Euclidean space. If X has a continuous distribution, then $f : T \rightarrow [0, \infty)$ is a probability density function of X if

$$\mathbb{P}(X \in A) = \int_A f d\lambda_n, \quad A \in \mathcal{R}^n \quad (3.10.45)$$

Technically, f is the density function of X with respect to counting measure $\#$ in the discrete case, and f is the density function of X with respect to Lebesgue measure λ_n in the continuous case. In both cases, the probability of an event A is computed by integrating the density function, with respect to the appropriate measure, over A . There are still differences, however. In the discrete case, the existence of the density function with respect to counting measure is guaranteed, and indeed we have an explicit formula for it. In the continuous case, the existence of a density function with respect to Lebesgue measure is not guaranteed, and indeed there might not be one. More generally, suppose that we have a measure space (T, \mathcal{T}, μ) and a random variable X with values in T . A measurable function $f : T \rightarrow [0, \infty)$ is a *probability density function* of X (or more precisely, the *distribution* of X) with respect to μ if

$$\mathbb{P}(X \in A) = \int_A f d\mu, \quad A \in \mathcal{T} \quad (3.10.46)$$

This fundamental question of the existence of a density function will be clarified in the section on absolute continuity and density functions.

Suppose again that X is a real-valued random variable with distribution function F . Then, by definition, the distribution of X is the Lebesgue-Stieltjes measure associated with F :

$$\mathbb{P}(a < X \leq b) = F(b) - F(a), \quad a, b \in \mathbb{R}, \quad a < b \quad (3.10.47)$$

regardless of whether the distribution is discrete, continuous, or mixed. Trivially, $\mathbb{P}(X \in A) = \int_S \mathbf{1}_A dF$ for $A \in \mathcal{R}$ and the expected value of X defined above can also be written as $\mathbb{E}(X) = \int_{\mathbb{R}} x dF(x)$. Again, all of this will be explained in much more detail in the next chapter on Expected Value.

Computational Exercises

Let $g(x) = \frac{1}{1+x^2}$ for $x \in \mathbb{R}$.

1. Find $\int_{-\infty}^{\infty} g(x) dx$.
2. Show that $\int_{-\infty}^{\infty} xg(x) dx$ does not exist.

Answer

1. $\int_{-\infty}^{\infty} g(x) dx = \pi$
2. $\int_0^{\infty} xg(x) dx = \infty$, $\int_{-\infty}^0 xg(x) dx = -\infty$

You may recall that the function g in the last exercise is important in the study of the *Cauchy distribution*, named for Augustin Cauchy. You may also remember that the graph of g is known as the *witch of Agnesi*, named for Maria Agnesi.

Let $g(x) = \frac{1}{x^b}$ for $x \in [1, \infty)$ where $b > 0$ is a parameter. Find $\int_1^{\infty} g(x) dx$

Answer

$$\int_1^{\infty} g(x) dx = \begin{cases} \infty, & 0 < b \leq 1 \\ \frac{1}{b-1}, & b > 1 \end{cases}$$

You may recall that the function g in the last exercise is important in the study of the *Pareto distribution*, named for Vilfredo Pareto.

Suppose that $f(x) = 0$ if $x \in \mathbb{Q}$ and $f(x) = \sin(x)$ if $x \in \mathbb{R} - \mathbb{Q}$.

1. Find $\int_{[0, \pi]} f(x) d\lambda(x)$
2. Does $\int_0^{\pi} f(x) dx$ exist?

Answer

1. 2
2. No

This page titled [3.10: The Integral With Respect to a Measure](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.11: Properties of the Integral

Basic Theory

Again our starting point is a measure space (S, \mathcal{S}, μ) . That is, S is a set, \mathcal{S} is a σ -algebra of subsets of S , and μ is a positive measure on \mathcal{S} .

Definition

In the last section we defined the integral of certain measurable functions $f : S \rightarrow \mathbb{R}$ with respect to the measure μ . Recall that the integral, denoted $\int_S f d\mu$, may exist as a number in \mathbb{R} (in which case f is *integrable*), or may exist as ∞ or $-\infty$, or may fail to exist. Here is a review of how the definition is built up in stages:

Definition of the integral

1. If f is a nonnegative simple function, so that $f = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ where I is a finite index set, $a_i \in [0, \infty)$ for $i \in I$, and $\{A_i : i \in I\}$ is measurable partition of S , then

$$\int_S f d\mu = \sum_{i \in I} a_i \mu(A_i) \quad (3.11.1)$$

2. If $f : S \rightarrow [0, \infty)$ is measurable, then

$$\int_S f d\mu = \sup \left\{ \int_S g d\mu : g \text{ is simple and } 0 \leq g \leq f \right\} \quad (3.11.2)$$

3. If $f : S \rightarrow \mathbb{R}$ is measurable, then

$$\int_S f d\mu = \int_S f^+ d\mu - \int_S f^- d\mu \quad (3.11.3)$$

as long as the right side is not of the form $\infty - \infty$, and where f^+ and f^- denote the positive and negative parts of f .

4. If $f : S \rightarrow \mathbb{R}$ is measurable and $A \in \mathcal{S}$, then the integral of f over A is defined by

$$\int_A f d\mu = \int_S \mathbf{1}_A f d\mu \quad (3.11.4)$$

assuming that the integral on the right exists.

Consider a statement on the elements of S , for example an equation or an inequality with $x \in S$ as a free variable. (Technically such a statement is a *predicate* on S .) For $A \in \mathcal{S}$, we say that the statement holds *on* A if it is true for every $x \in A$. We say that the statement holds *almost everywhere on* A (with respect to μ) if there exists $B \in \mathcal{S}$ with $B \subseteq A$ such that the statement holds on B and $\mu(A \setminus B) = 0$.

Basic Properties

A few properties of the integral that were essential to the motivation of the definition were given in the last section. In this section, we extend some of those properties and we study a number of new ones. As a review, here is what we know so far.

Properties of the integral

1. If $f, g : S \rightarrow \mathbb{R}$ are measurable functions whose integrals exist, then $\int_S (f + g) d\mu = \int_S f d\mu + \int_S g d\mu$ as long as the right side is not of the form $\infty - \infty$.
2. If $f : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists and $c \in \mathbb{R}$, then $\int_S cf d\mu = c \int_S f d\mu$.
3. If $f : S \rightarrow \mathbb{R}$ is measurable and $f \geq 0$ on S then $\int_S f d\mu \geq 0$.
4. If $f, g : S \rightarrow \mathbb{R}$ are measurable functions whose integrals exist and $f \leq g$ on S then $\int_S f d\mu \leq \int_S g d\mu$.
5. If $f_n : S \rightarrow [0, \infty)$ is measurable for $n \in \mathbb{N}_+$ and f_n is increasing in n on S then $\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu$.
6. $f : S \rightarrow \mathbb{R}$ is measurable and the the integral of f on $A \cup B$ exists, where $A, B \in \mathcal{S}$ are disjoint, then $\int_{A \cup B} f d\mu = \int_A f d\mu + \int_B f d\mu$.

Parts (a) and (b) are the *linearity properties*; part (a) is the *additivity property* and part (b) is the *scaling property*. Parts (c) and (d) are the *order properties*; part (c) is the *positive property* and part (d) is the *increasing property*. Part (e) is a continuity property known as the *monotone convergence theorem*. Part (f) is the *additive property for disjoint domains*. Properties (a)–(e) hold with S replaced by $A \in \mathcal{S}$.

Equality and Order

Our first new results are extensions dealing with equality and order. The integral of a function over a null set is 0:

Suppose that $f : S \rightarrow \mathbb{R}$ is measurable and $A \in \mathcal{S}$ with $\mu(A) = 0$. Then $\int_A f d\mu = 0$.

Proof

The proof proceeds in stages via the definition of the integral.

1. Suppose that g is a nonnegative simple function with $g = 0$ on A^c . Then g has the representation $g = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ where $a_i \in (0, \infty)$ and $A_i \subseteq A$ for $i \in I$. But $\mu(A_i) = 0$ for each $i \in I$ and so $\int_S g d\mu = \sum_{i \in I} a_i \mu(A_i) = 0$.
2. Suppose that $f : S \rightarrow [0, \infty)$ is measurable. If g is a nonnegative simple function with $g \leq \mathbf{1}_A f$, then $g = 0$ on A^c so by (a), $\int_S g d\mu = 0$. Hence by part (b) of (1), $\int_A f d\mu = \int_S \mathbf{1}_A f d\mu = 0$.
3. Finally, suppose that $f : S \rightarrow \mathbb{R}$ is measurable. Then $\int_A f d\mu = \int_A f^+ d\mu - \int_A f^- d\mu$. But both integrals on the right are 0 by part (b).

Two functions that are indistinguishable from the point of view of μ must have the same integral.

Suppose that $f : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists. If $g : S \rightarrow \mathbb{R}$ is measurable and $g = f$ almost everywhere on S , then $\int_S g d\mu = \int_S f d\mu$.

Proof

Note that $g = f$ if and only if $g^+ = f^+$ and $g^- = f^-$. Let $A = \{x \in S : g^+(x) = f^+(x)\}$. Then $A \in \mathcal{S}$ and $\mu(A^c) = 0$. Hence by the additivity property and (3),

$$\int_S g^+ d\mu = \int_A g^+ d\mu + \int_{A^c} g^+ d\mu = \int_A f^+ d\mu + 0 = \int_A f^+ d\mu + \int_{A^c} f^+ d\mu = \int_S f^+ d\mu \quad (3.11.5)$$

Similarly $\int_S g^- d\mu = \int_S f^- d\mu$. Hence the integral of g exists and $\int_S g d\mu = \int_S f d\mu$.

Next we have a simple extension of the positive property.

Suppose that $f : S \rightarrow \mathbb{R}$ is measurable and $f \geq 0$ almost everywhere on S . Then

1. $\int_S f d\mu \geq 0$
2. $\int_S f = 0$ if and only if $f = 0$ almost everywhere on S .

Proof

1. Let $A = \{x \in S : f(x) \geq 0\}$. Then $A \in \mathcal{S}$ and $\mu(A^c) = 0$. By the additivity of the integral over disjoint sets we have

$$\int_S f d\mu = \int_A f d\mu + \int_{A^c} f d\mu \quad (3.11.6)$$

But $\int_A f d\mu \geq 0$ by the [positive property](#) and $\int_{A^c} f d\mu = 0$ by the [null property](#), so $\int_S f d\mu \geq 0$.

2. Note first that if $\mu(A) = 0$ then both integrals in the displayed equation are 0 so $\int_S f d\mu = 0$. For the converse, let $B_n = \{x \in S : f(x) \geq \frac{1}{n}\}$ for $n \in \mathbb{N}_+$ and $B = \{x \in S : f(x) > 0\}$. Then B_n is increasing in n and $\bigcup_{n=1}^{\infty} B_n = B$. If $\mu(B) > 0$ then $\mu(B_n) > 0$ for some $n \in \mathbb{N}_+$. But $f \geq \frac{1}{n} \mathbf{1}_{B_n}$ on A , so by the increasing property, $\int_S f d\mu = \int_A f d\mu \geq \int_A \frac{1}{n} \mathbf{1}_{B_n} d\mu = \frac{1}{n} \mu(B_n) > 0$.

So, if $f \geq 0$ almost everywhere on S then $\int_S f d\mu > 0$ if and only if $\mu\{x \in S : f(x) > 0\} > 0$. The simple extension of the positive property in turn leads to a simple extension of the increasing property.

Suppose that $f, g : S \rightarrow \mathbb{R}$ are measurable functions whose integrals exist, and that $f \leq g$ almost everywhere on S . Then

1. $\int_S f \leq \int_S g$

2. Except in the case that both integrals are ∞ or both $-\infty$, $\int_S f d\mu = \int_S g d\mu$ if and only if $f = g$ almost everywhere on S .

Proof

1. Note that $g = f + (g - f)$ and $g - f \geq 0$ almost everywhere on S . If $\int_S f d\mu = -\infty$ then trivially $\int_S f d\mu \leq \int_S g d\mu$. Otherwise, by the additive property,

$$\int_S g d\mu = \int_S f d\mu + \int_S (g - f) d\mu \quad (3.11.7)$$

By the [positive property](#), $\int_S (g - f) d\mu \geq 0$ so $\int_S g d\mu \geq \int_S f d\mu$.

2. Except in the case that both integrals are ∞ or both are $-\infty$ we have

$$\int_S g d\mu - \int_S f d\mu = \int_S (g - f) d\mu \quad (3.11.8)$$

By assumption $g - f \geq 0$ almost everywhere on S , and hence by the [positive property](#), the integral on the right is 0 if and only if $g - f = 0$ almost everywhere on S .

So if $f \leq g$ almost everywhere on S then, except in the two cases mentioned, $\int_S f d\mu < \int_S g d\mu$ if and only if $\mu\{x \in S : f(x) < g(x)\} > 0$. The exclusion when both integrals are ∞ or $-\infty$ is important. A counterexample when this condition does not hold is given [below](#). The next result is the absolute value inequality.

Suppose that $f : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists. Then

$$\left| \int_S f d\mu \right| \leq \int_S |f| d\mu \quad (3.11.9)$$

If f is integrable, then equality holds if and only if $f \geq 0$ almost everywhere on S or $f \leq 0$ almost everywhere on S .

Proof

First note that $-|f| \leq f \leq |f|$ on S . The integrals of all three functions exist, so the increasing property and scaling properties give

$$-\int_S |f| d\mu \leq \int_S f d\mu \leq \int_S |f| d\mu \quad (3.11.10)$$

which is equivalent to the inequality above. If f is integrable, then by the [increasing property](#), equality holds if and only if $f = -|f|$ almost everywhere on S or $f = |f|$ almost everywhere on S . In the first case, $f \leq 0$ almost everywhere on S and in the second case, $f \geq 0$ almost everywhere on S .

Change of Variables

Suppose that (T, \mathcal{T}) is another measurable space and that $u : S \rightarrow T$ is measurable. As we saw in our first study of positive measures, ν defined by

$$\nu(B) = \mu[u^{-1}(B)], \quad B \in \mathcal{T} \quad (3.11.11)$$

is a positive measure on (T, \mathcal{T}) . The following result is known as the *change of variables* theorem.

If $f : T \rightarrow \mathbb{R}$ is measurable then, assuming that the integrals exist,

$$\int_T f d\nu = \int_S (f \circ u) d\mu \quad (3.11.12)$$

Proof

We will show that if either of the integrals exist then they both do, and are equal. The proof is a classical bootstrapping argument that parallels the definition of the integral.

1. Suppose first that f is a nonnegative simple function on T with the representation $f = \sum_{i \in I} b_i \mathbf{1}_{B_i}$ where I is a finite index set, $\{B_i : i \in I\}$ is a measurable partition of T , and $b_i \in [0, \infty)$ for $i \in I$. Recall that $f \circ u$ is a nonnegative simple function on S , with representation $f \circ u = \sum_{i \in I} b_i \mathbf{1}_{u^{-1}(B_i)}$. Hence

$$\int_T f \, d\nu = \sum_{i \in I} b_i \nu(B_i) = \sum_{i \in I} b_i \mu[u^{-1}(B_i)] = \int_S (f \circ u) \, d\mu \quad (3.11.13)$$

2. Next suppose that $f : T \rightarrow [0, \infty)$ is measurable, so that $f \circ u : S \rightarrow [0, \infty)$ is also measurable. There exists an increasing sequence (f_1, f_2, \dots) of nonnegative simple functions on T with $f_n \rightarrow f$ as $n \rightarrow \infty$. Then $(f_1 \circ u, f_2 \circ u, \dots)$ is an increasing sequence of simple functions on S with $f_n \circ u \rightarrow f \circ u$ as $n \rightarrow \infty$. By step (a), $\int_T f_n \, d\nu = \int_S (f_n \circ u) \, d\mu$ for each $n \in \mathbb{N}_+$. But by the monotone convergence theorem, $\int_T f_n \, d\nu \rightarrow \int_T f \, d\nu$ as $n \rightarrow \infty$ and $\int_S (f_n \circ u) \, d\mu \rightarrow \int_S (f \circ u) \, d\mu$ so we conclude that $\int_T f \, d\nu = \int_S (f \circ u) \, d\mu$.
3. Finally, suppose that $f : T \rightarrow \mathbb{R}$ is measurable, so that $f \circ u : S \rightarrow \mathbb{R}$ is also measurable. Note that $(f \circ u)^+ = f^+ \circ u$ and $(f \circ u)^- = f^- \circ u$. By part (b),

$$\int_T f^+ \, d\nu = \int_S (f^+ \circ u) \, d\mu = \int_S (f \circ u)^+ \, d\mu \quad (3.11.14)$$

$$\int_T f^- \, d\nu = \int_S (f^- \circ u) \, d\mu = \int_S (f \circ u)^- \, d\mu \quad (3.11.15)$$

Assuming that at least one of the integrals in the displayed equations is finite, we have

$$\int_T f \, d\nu = \int_T f^+ \, d\nu - \int_T f^- \, d\nu = \int_S (f \circ u)^+ \, d\mu - \int_S (f \circ u)^- \, d\mu = \int_S (f \circ u) \, d\mu \quad (3.11.16)$$

The change of variables theorem will look more familiar if we give the variables explicitly. Thus, suppose that we want to evaluate

$$\int_S f[u(x)] \, d\mu(x) \quad (3.11.17)$$

where again, $u : S \rightarrow T$ and $f : T \rightarrow \mathbb{R}$. One way is to use the substitution $u = u(x)$, find the new measure ν , and then evaluate

$$\int_T g(u) \, d\nu(u) \quad (3.11.18)$$

Convergence Properties

We start with a simple but important corollary of the monotone convergence theorem that extends the additivity property to a countably infinite sum of nonnegative functions.

Suppose that $f_n : S \rightarrow [0, \infty)$ is measurable for $n \in \mathbb{N}_+$. Then

$$\int_S \sum_{n=1}^{\infty} f_n \, d\mu = \sum_{n=1}^{\infty} \int_S f_n \, d\mu \quad (3.11.19)$$

Proof

Let $g_n = \sum_{i=1}^n f_i$ for $n \in \mathbb{N}_+$. Then $g_n : S \rightarrow [0, \infty)$ is measurable and g_n is increasing in n . Moreover, by definition, $g_n \rightarrow \sum_{i=1}^{\infty} f_i$ as $n \rightarrow \infty$. Hence by the MCT, $\int_S g_n \, d\mu \rightarrow \int_S \sum_{i=1}^{\infty} f_i \, d\mu$ as $n \rightarrow \infty$. But we know the additivity property holds for finite sums, so $\int_S g_n \, d\mu = \sum_{i=1}^n \int_S f_i \, d\mu$ and again, by definition, this sum converges to $\sum_{i=1}^{\infty} \int_S f_i \, d\mu$ as $n \rightarrow \infty$.

A [theorem below](#) gives a related result that relaxes the assumption that f be nonnegative, but imposes a stricter integrability requirement. Our next result is the additivity of the integral over a countably infinite collection of disjoint domains.

Suppose that $f : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists, and that $\{A_n : n \in \mathbb{N}_+\}$ is a disjoint collection of sets in \mathcal{S} . Let $A = \bigcup_{n=1}^{\infty} A_n$. Then

$$\int_A f \, d\mu = \sum_{n=1}^{\infty} \int_{A_n} f \, d\mu \quad (3.11.20)$$

Proof

Suppose first that f is nonnegative. Note that $\mathbf{1}_A = \sum_{n=1}^{\infty} \mathbf{1}_{A_n}$ and hence $\mathbf{1}_A f = \sum_{n=1}^{\infty} \mathbf{1}_{A_n} f$. Thus from the [theorem above](#),

$$\int_A f d\mu = \int_S \mathbf{1}_A f d\mu = \int_S \sum_{n=1}^{\infty} \mathbf{1}_{A_n} f d\mu = \sum_{n=1}^{\infty} \int_S \mathbf{1}_{A_n} f d\mu = \sum_{n=1}^{\infty} \int_{A_n} f d\mu \quad (3.11.21)$$

Suppose now that $f : S \rightarrow \mathbb{R}$ is measurable and $\int_S f d\mu$ exists. Note that for $B \in \mathcal{S}$, $(\mathbf{1}_B f)^+ = \mathbf{1}_B f^+$ and $(\mathbf{1}_B f)^- = \mathbf{1}_B f^-$. Hence from the previous argument,

$$\int_A f^+ d\mu = \sum_{n=1}^{\infty} \int_{A_n} f^+ d\mu, \quad \int_A f^- d\mu = \sum_{n=1}^{\infty} \int_{A_n} f^- d\mu \quad (3.11.22)$$

Both of these are sums of nonnegative terms, and one of the sums, at least, is finite. Hence we can group the terms to get

$$\int_A f d\mu = \int_A f^+ d\mu - \int_A f^- d\mu = \sum_{n=1}^{\infty} \int_{A_n} (f^+ - f^-) d\mu = \sum_{n=1}^{\infty} \int_{A_n} f d\mu \quad (3.11.23)$$

Of course, the previous theorem applies if f is nonnegative or if f is integrable. Next we give a minor extension of the monotone convergence theorem that relaxes the assumption that the functions be nonnegative.

Monotone Convergence Theorem. Suppose that $f_n : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists for each $n \in \mathbb{N}_+$ and that f_n is increasing in n on S . If $\int_S f_1 d\mu > -\infty$ then

$$\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.11.24)$$

Proof

Let $f(x) = \lim_{n \rightarrow \infty} f_n(x)$ for $x \in S$ which exists in $\mathbb{R} \cup \{\infty\}$ since $f_n(x)$ is increasing in $n \in \mathbb{N}_+$. If $\int_S f_1 d\mu = \infty$, then by the increasing property, $\int_S f_n d\mu = \infty$ for all $n \in \mathbb{N}_+$ and $\int_S f d\mu = \infty$, so the conclusion of the MCT trivially holds. Thus suppose that f_1 is integrable. Let $g_n = f_n - f_1$ for $n \in \mathbb{N}$ and let $g = f - f_1$. Then g_n is nonnegative and increasing in n on S , and $g_n \rightarrow g$ as $n \rightarrow \infty$ on S . By the ordinary MCT, $\int_S g_n d\mu \rightarrow \int_S g d\mu$ as $n \rightarrow \infty$. But since $\int_S f_1 d\mu$ is finite, $\int_S g_n d\mu = \int_S f_n d\mu - \int_S f_1 d\mu$ and $\int_S g d\mu = \int_S f d\mu - \int_S f_1 d\mu$. Again since $\int_S f_1 d\mu$ is finite, it follows that $\int_S f_n d\mu \rightarrow \int_S f d\mu$ as $n \rightarrow \infty$.

Here is the complementary result for decreasing functions.

Suppose that $f_n : S \rightarrow \mathbb{R}$ is a measurable function whose integral exists for each $n \in \mathbb{N}_+$ and that f_n is decreasing in n on S . If $\int_S f_1 d\mu < \infty$ then

$$\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.11.25)$$

Proof

The functions $-f_n$ for $n \in \mathbb{N}_+$ satisfy the hypotheses of the [MCT for increasing functions](#) and hence $\int_S \lim_{n \rightarrow \infty} -f_n d\mu = \lim_{n \rightarrow \infty} \int_S -f_n d\mu$. By the scaling property, $\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu$.

The additional assumptions on the integral of f_1 in the last two extensions of the monotone convergence theorem are necessary. An example is given in [below](#).

Our next result is also a consequence of the monotone convergence theorem, and is called *Fatou's lemma* in honor of Pierre Fatou. Its usefulness stems from the fact that no assumptions are placed on the integrand functions, except that they be nonnegative and measurable.

Fatou's Lemma. Suppose that $f_n : S \rightarrow [0, \infty)$ is measurable for $n \in \mathbb{N}_+$. Then

$$\int_S \liminf_{n \rightarrow \infty} f_n d\mu \leq \liminf_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.11.26)$$

Proof

Let $g_n = \inf \{f_k : k \in \{n, n+1, \dots\}\}$ for $n \in \mathbb{N}_+$. Then $g_n : S \rightarrow [0, \infty)$ is measurable for $n \in \mathbb{N}_+$, g_n is increasing in n , and by definition, $\lim_{n \rightarrow \infty} g_n = \liminf_{n \rightarrow \infty} f_n$. By the MCT,

$$\int_S \liminf_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S g_n d\mu \quad (3.11.27)$$

But $g_n \leq f_k$ on S for $n \in \mathbb{N}_+$ and $k \in \{n, n+1, \dots\}$ so by the increasing property, $\int_S g_n d\mu \leq \int_S f_k d\mu$ for $n \in \mathbb{N}_+$ and $k \in \{n, n+1, \dots\}$. Hence $\int_S g_n d\mu \leq \inf \{\int_S f_k d\mu : k \in \{n, n+1, \dots\}\}$ for $n \in \mathbb{N}_+$ and therefore

$$\lim_{n \rightarrow \infty} \int_S g_n d\mu \leq \liminf_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.11.28)$$

Given the weakness of the hypotheses, it's hardly surprising that strict inequality can easily occur in Fatou's lemma. An example is given [below](#).

Our next convergence result is one of the most important and is known as the *dominated convergence theorem*. It's sometimes also known as *Lebesgue's dominated convergence theorem* in honor of Henri Lebesgue, who first developed all of this stuff in the context of \mathbb{R}^n . The dominated convergence theorem gives a basic condition under which we may interchange the limit and integration operators.

Dominated Convergence Theorem. Suppose that $f_n : S \rightarrow \mathbb{R}$ is measurable for $n \in \mathbb{N}_+$ and that $\lim_{n \rightarrow \infty} f_n$ exists on S . Suppose also that $|f_n| \leq g$ for $n \in \mathbb{N}$ where $g : S \rightarrow [0, \infty)$ is integrable. Then

$$\int_S \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu \quad (3.11.29)$$

Proof

First note that by the increasing property, $\int_S |f_n| d\mu \leq \int_S g d\mu < \infty$ and hence f_n is integrable for $n \in \mathbb{N}_+$. Let $f = \lim_{n \rightarrow \infty} f_n$. Then f is measurable, and by the increasing property again, $\int_S |f| d\mu \leq \int_S g d\mu < \infty$, so f is integrable.

Now for $n \in \mathbb{N}_+$, let $u_n = \inf \{f_k : k \in \{n, n+1, \dots\}\}$ and let $v_n = \sup \{f_k : k \in \{n, n+1, \dots\}\}$. Then $u_n \leq f_n \leq v_n$ for $n \in \mathbb{N}_+$, u_n is increasing in n , v_n is decreasing in n , and $u_n \rightarrow f$ and $v_n \rightarrow f$ as $n \rightarrow \infty$. Moreover, $\int_S u_1 d\mu \geq -\int_S g d\mu > -\infty$ so by the version of the MCT above, $\int_S u_n d\mu \rightarrow \int_S f d\mu$ as $n \rightarrow \infty$. Similarly, $\int_S v_1 d\mu < \int_S g d\mu < \infty$, so by the MCT in (11), $\int_S v_n d\mu \rightarrow \int_S f d\mu$ as $n \rightarrow \infty$. But by the increasing property, $\int_S u_n d\mu \leq \int_S f_n d\mu \leq \int_S v_n d\mu$ for $n \in \mathbb{N}_+$ so by the squeeze theorem for limits, $\int_S f_n d\mu \rightarrow \int_S f d\mu$ as $n \rightarrow \infty$.

As you might guess, the assumption that $|f_n|$ is uniformly bounded in n by an integrable function is critical. A counterexample when this assumption is missing is given [below](#) when this assumption is missing. The dominated convergence theorem remains true if $\lim_{n \rightarrow \infty} f_n$ exists almost everywhere on S . The following corollary of the dominated convergence theorem gives a condition for the interchange of infinite sum and integral.

Suppose that $f_i : S \rightarrow \mathbb{R}$ is measurable for $i \in \mathbb{N}_+$ and that $\sum_{i=1}^{\infty} |f_i|$ is integrable. then

$$\int_S \sum_{i=1}^{\infty} f_i d\mu = \sum_{i=1}^{\infty} \int_S f_i d\mu \quad (3.11.30)$$

Proof

The assumption that $g = \sum_{i=1}^{\infty} |f_i|$ is integrable implies that $g < \infty$ almost everywhere on S . In turn, this means that $\sum_{i=1}^{\infty} f_i$ is absolutely convergent almost everywhere on S . Let $f(x) = \sum_{i=1}^{\infty} f_i(x)$ if $g(x) < \infty$, and for completeness, let $f(x) = 0$ if $g(x) = \infty$. Since only the integral of f appears in the theorem, it doesn't matter how we define f on the null set where $g = \infty$. Now let $g_n = \sum_{i=1}^n |f_i|$. Then $g_n \rightarrow g$ as $n \rightarrow \infty$ almost everywhere on S and $|g_n| \leq g$ on S . Hence by the [dominated convergence theorem](#), $\int_S g_n d\mu \rightarrow \int_S g d\mu$ as $n \rightarrow \infty$. But we know the additivity property holds for finite sums, so $\int_S g_n d\mu = \sum_{i=1}^n \int_S |f_i| d\mu$, and in turn this converges to $\sum_{i=1}^{\infty} \int_S |f_i| d\mu$ as $n \rightarrow \infty$. Thus we have $\sum_{i=1}^{\infty} \int_S |f_i| d\mu = \int_S g d\mu$.

The following corollary of the dominated convergence theorem is known as the *bounded convergence theorem*.

Bounded Convergence Theorem. Suppose that $f_n : S \rightarrow \mathbb{R}$ is measurable for $n \in \mathbb{N}_+$ and there exists $A \in \mathcal{S}$ such that $\mu(A) < \infty$, $\lim_{n \rightarrow \infty} f_n$ exists on A , and $|f_n|$ is bounded in $n \in \mathbb{N}_+$ on A . Then

$$\int_A \lim_{n \rightarrow \infty} f_n d\mu = \lim_{n \rightarrow \infty} \int_A f_n d\mu \quad (3.11.31)$$

Proof

Suppose that $|f_n|$ is bounded in n on A by $c \in (0, \infty)$. The constant c is integrable on A since $\int_A c d\mu = c\mu(A) < \infty$, and $|f_n| \leq c$ on A for $n \in \mathbb{N}_+$. Thus the result follows from the [dominated convergence theorem](#).

Again, the bounded convergence remains true if $\lim_{n \rightarrow \infty} f_n$ exists almost everywhere on A . For a finite measure space (and in particular for a probability space), the condition that $\mu(A) < \infty$ automatically holds.

Product Spaces

Suppose now that (S, \mathcal{S}, μ) and (T, \mathcal{T}, ν) are σ -finite measure spaces. Please recall the basic facts about the product σ -algebra $\mathcal{S} \otimes \mathcal{T}$ of subsets of $S \times T$, and the product measure $\mu \otimes \nu$ on $\mathcal{S} \otimes \mathcal{T}$. The product measure space $(S \times T, \mathcal{S} \otimes \mathcal{T}, \mu \otimes \nu)$ is the standard one that we use for product spaces. If $f : S \times T \rightarrow \mathbb{R}$ is measurable, there are three integrals we might consider. First, of course, is the integral of f with respect to the product measure $\mu \otimes \nu$

$$\int_{S \times T} f(x, y) d(\mu \otimes \nu)(x, y) \quad (3.11.32)$$

sometimes called a *double integral* in this context. But also we have the nested or *iterated integrals* where we integrate with respect to one variable at a time:

$$\int_S \left(\int_T f(x, y) d\nu(y) \right) d\mu(x), \quad \int_T \left(\int_S f(x, y) d\mu(x) \right) d\nu(y) \quad (3.11.33)$$

How are these integrals related? Well, just as in calculus with ordinary Riemann integrals, under mild conditions the three integrals are the same. The resulting important theorem is known as *Fubini's Theorem* in honor of the Italian mathematician Guido Fubini.

Fubini's Theorem. Suppose that $f : S \times T \rightarrow \mathbb{R}$ is measurable. If the double integral on the left exists, then

$$\int_{S \times T} f(x, y) d(\mu \otimes \nu)(x, y) = \int_S \int_T f(x, y) d\nu(y) d\mu(x) = \int_T \int_S f(x, y) d\mu(x) d\nu(y) \quad (3.11.34)$$

Proof

We will show that

$$\int_{S \times T} f(x, y) d(\mu \otimes \nu)(x, y) = \int_S \int_T f(x, y) d\nu(y) d\mu(x) \quad (3.11.35)$$

The proof with the other iterated integral is symmetric. The proof proceeds in stages, paralleling the definition of the integral.

1. Suppose that $f = \mathbf{1}_{A \times B}$ where $A \in \mathcal{S}$ and $B \in \mathcal{T}$. The equation holds by definition of the product measure, since the double integral is $(\mu \otimes \nu)(A \times B)$ and the iterated integral is

$$\int_S \int_T \mathbf{1}_{A \times B}(x, y) d\nu(y) d\mu(x) = \int_S \int_T \mathbf{1}_A(x) \mathbf{1}_B(y) d\nu(y) d\mu(x) = \int_S \mathbf{1}_A(x) \nu(B) d\mu(x) = \mu(A) \nu(B) \quad (3.11.36)$$

2. Consider $f = \mathbf{1}_C$ where $C \in \mathcal{S} \otimes \mathcal{T}$. The double integral is $(\mu \otimes \nu)(C)$, and so as a function of $C \in \mathcal{S} \otimes \mathcal{T}$ defines the measure $\mu \otimes \nu$. On the other hand, the iterated integral is

$$\int_S \int_T \mathbf{1}_C(x, y) d\nu(y) d\mu(x) = \int_S \int_T \mathbf{1}_{C_x}(y) d\nu(y) d\mu(x) = \int_S \nu(C_x) d\mu(x) \quad (3.11.37)$$

where $C_x = \{y \in T : (x, y) \in C\}$ is the cross-section of C at $x \in S$. Recall that $x \mapsto \nu(C_x)$ is a nonnegative, measurable function of x , so $C \mapsto \int_S \nu(C_x) d\mu(x)$ makes sense. Moreover, as a function of $C \in \mathcal{S} \otimes \mathcal{T}$, this integral also forms a measure: If $\{C^i : i \in I\}$ is a countable, disjoint collection sets in $\mathcal{S} \otimes \mathcal{T}$, then $\{C_x^i : i \in I\}$ is a countable,

disjoint collection of sets in \mathcal{T} . Cross-sections preserve set operations, so if $C = \bigcup_{i \in I} C^i$ then $C_x = \bigcup_{i \in I} C_x^i$. By the additivity of the measure ν and the integral we have

$$\int_S \nu(C_x) d\mu(x) = \int_S \nu\left(\bigcup_{i \in I} C_x^i\right) d\mu(x) = \int_S \sum_{i \in I} \nu(C_x^i) d\mu(x) = \sum_{i \in I} \int_S \nu(C_x^i) d\mu(x) \quad (3.11.38)$$

To summarize, the double integral and the iterated integral define positive measures on $\mathcal{S} \otimes \mathcal{T}$. By (a), these measures agree on the measurable rectangles. By the uniqueness theorem, they must be the same measure. Thus the double integral and the iterated integral agree with integrand $f = \mathbf{1}_C$ for every $C \in \mathcal{S} \otimes \mathcal{T}$.

3. Suppose $f = \sum_{i \in I} c_i \mathbf{1}_{C_i}$ is a nonnegative simple function on $S \times T$. Thus, I is a finite index set, $c_i \in [0, \infty)$ for $i \in I$, and $\{C_i : i \in I\}$ is a disjoint collection of sets in $\mathcal{S} \otimes \mathcal{T}$. The double integral and the iterated integral satisfy the linearity properties, and hence by (b), agree with integrand f .
4. Suppose that $f : S \times T \rightarrow [0, \infty)$ is measurable. Then there exists a sequence of nonnegative simple functions g_n , $n \in \mathbb{N}_+$ such that g_n is increasing in $n \in \mathbb{N}_+$ on $S \times T$, and $g_n \rightarrow f$ as $n \rightarrow \infty$ on $S \times T$. By the monotone convergence theorem, $\int_{S \times T} g_n d(\mu \otimes \nu) \rightarrow \int_{S \times T} f d(\mu \otimes \nu)$. But for fixed $x \in S$, $y \mapsto g_n(x, y)$ is increasing in n on T and has limit $f(x, y)$ as $n \rightarrow \infty$. By another application of the monotone convergence theorem, $\int_T g_n(x, y) d\nu(y) \rightarrow \int_T f(x, y) d\nu(y)$ as $n \rightarrow \infty$. But $x \mapsto \int_T g_n(x, y) d\nu(y)$ is measurable and is increasing in $n \in \mathbb{N}_+$ on S , so by yet another application of the monotone convergence theorem, $\int_S \int_T g_n(x, y) d\nu(y) d\mu(x) \rightarrow \int_S \int_T f(x, y) d\nu(y) d\mu(x)$ as $n \rightarrow \infty$. But the double integral and the iterated integral agree with integrand g_n by (c) for each $n \in \mathbb{N}_+$, so it follows that the double integral and the iterated integral agree with integrand f .
5. Suppose that $f : S \times T \rightarrow \mathbb{R}$ is measurable. By (d), the double integral and the iterated integral agree with integrand functions f^+ and f^- . Assuming that at least one of these is finite, then by the additivity property, they agree with integrand function $f = f^+ - f^-$.

Of course, the double integral exists, and so Fubini's theorem applies, if either f is nonnegative or integrable with respect to $\mu \otimes \nu$. When f is nonnegative, the result is sometimes called *Tonelli's theorem* in honor of another Italian mathematician, Leonida Tonelli. On the other hand, the iterated integrals may exist, and may be different, when the double integral does not exist. A [counterexample](#) and a [second counterexample](#) are given below.

A special case of Fubini's theorem (and indeed part of the proof) is that we can compute the measure of a set in the product space by integrating the cross-sectional measures.

If $C \in \mathcal{S} \otimes \mathcal{T}$ then

$$(\mu \otimes \nu)(C) = \int_S \nu(C_x) d\mu(x) = \int_T \mu(C^y) d\nu(y) \quad (3.11.39)$$

where $C_x = \{y \in T : (x, y) \in C\}$ for $x \in S$, and $C^y = \{x \in S : (x, y) \in C\}$ for $y \in T$.

In particular, if $C, D \in \mathcal{S} \otimes \mathcal{T}$ have the property that $\nu(C_x) = \nu(D_x)$ for all $x \in S$, or $\mu(C^y) = \mu(D^y)$ for all $y \in T$ (that is, C and D have the same cross-sectional measures with respect to one of the variables), then $(\mu \otimes \nu)(C) = (\mu \otimes \nu)(D)$. In \mathbb{R}^2 with area, and in \mathbb{R}^3 with volume (Lebesgue measure in both cases), this is known as *Cavalieri's principle*, named for Bonaventura Cavalieri, yet a third Italian mathematician. Clearly, Italian mathematicians cornered the market on theorems of this sort.

A simple corollary of Fubini's theorem is that the double integral of a product function over a product set is the product of the integrals. This result has important applications to independent random variables.

Suppose that $g : S \rightarrow \mathbb{R}$ and $h : T \rightarrow \mathbb{R}$ are measurable, and are either nonnegative or integrable with respect to μ and ν , respectively. Then

$$\int_{S \times T} g(x)h(y) d(\mu \otimes \nu)(x, y) = \left(\int_S g(x) d\mu(x) \right) \left(\int_T h(y) d\nu(y) \right) \quad (3.11.40)$$

Recall that a *discrete measure space* consists of a countable set with the σ -algebra of all subsets and with counting measure. In such a space, integrals are simply sums and so Fubini's theorem allows us to rearrange the order of summation in a double sum.

Suppose that I and J are countable and that $a_{ij} \in \mathbb{R}$ for $i \in I$ and $j \in J$. If the sum of the positive terms or the sum of the negative terms is finite, then

$$\sum_{(i,j) \in I \times J} a_{ij} = \sum_{i \in I} \sum_{j \in J} a_{ij} = \sum_{j \in J} \sum_{i \in I} a_{ij} \quad (3.11.41)$$

Often $I = J = \mathbb{N}_+$, and in this case, a_{ij} can be viewed as an infinite array, with $i \in \mathbb{N}_+$ the row number and $j \in \mathbb{N}_+$ the column number:

a_{11}	a_{12}	a_{13}	\dots
a_{21}	a_{22}	a_{23}	\dots
a_{31}	a_{32}	a_{33}	\dots
\vdots	\vdots	\vdots	\vdots

The significant point is that \mathbb{N}_+ is totally ordered. While there is no implied order of summation in the double sum $\sum_{(i,j) \in \mathbb{N}_+^2} a_{ij}$, the iterated sum $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}$ is obtained by summing over the rows in order and then summing the results by column in order, while the iterated sum $\sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij}$ is obtained by summing over the columns in order and then summing the results by row in order.

Of course, only one of the product spaces might be discrete. Theorems (9) and (15) which give conditions for the interchange of sum and integral can be viewed as applications of Fubini's theorem, where one of the measure spaces is (S, \mathcal{S}, μ) and the other is \mathbb{N}_+ with counting measure.

Examples and Applications

Probability Spaces

Suppose that $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, so that Ω is the set of outcomes of a random experiment, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is a probability measure on the sample space (Ω, \mathcal{F}) . Suppose also that (S, \mathcal{S}) is another measurable space, and that X is a random variable for the experiment, taking values in S . Of course, this simply means that X is a measurable function from Ω to S . Recall that the *probability distribution* of X is the probability measure P_X on (S, \mathcal{S}) defined by

$$P_X(A) = \mathbb{P}(X \in A), \quad A \in \mathcal{S} \quad (3.11.42)$$

Since $\{X \in A\}$ is just probability notation for the inverse image of A under X , P_X is simply a special case of constructing a new positive measure from a given positive measure via a change of variables. Suppose now that $r : S \rightarrow \mathbb{R}$ is measurable, so that $r(X)$ is a real-valued random variable. The integral of $r(X)$ (assuming that it exists) is known as the *expected value* of $r(X)$ and is of fundamental importance. We will study expected values in detail in the next chapter. Here, we simply note different ways to write the integral. By the change of variables formula (8) we have

$$\int_{\Omega} r[X(\omega)] d\mathbb{P}(\omega) = \int_S r(x) dP_X(x) \quad (3.11.43)$$

Now let F_Y denote the distribution function of $Y = r(X)$. By another change of variables, Y has a probability distribution P_Y on \mathbb{R} , which is also a *Lebesgue-Stieltjes* measure, named for Henri Lebesgue and Thomas Stieltjes. Recall that this probability measure is characterized by

$$P_Y(a, b] = \mathbb{P}(a < Y \leq b) = F_Y(b) - F_Y(a); \quad a, b \in \mathbb{R}, \quad a < b \quad (3.11.44)$$

With another application of our change of variables theorem, we can add to our chain of integrals:

$$\int_{\Omega} r[X(\omega)] d\mathbb{P}(\omega) = \int_S r(x) dP_X(x) = \int_{\mathbb{R}} y dP_Y(y) = \int_{\mathbb{R}} y dF_Y(y) \quad (3.11.45)$$

Of course, the last two integrals are simply different notations for exactly the same thing. In the section on absolute continuity and density functions, we will see other ways to write the integral.

Counterexamples

In the first three exercises below, $(\mathbb{R}, \mathcal{R}, \lambda)$ is the standard one-dimensional Euclidean space, so \mathcal{R} is σ -algebra of Lebesgue measurable sets and λ is Lebesgue measure.

Let $f = \mathbf{1}_{[1, \infty)}$ and $g = \mathbf{1}_{[0, \infty)}$. Show that

1. $f \leq g$ on \mathbb{R}
2. $\lambda\{x \in \mathbb{R} : f(x) < g(x)\} = 1$
3. $\int_{\mathbb{R}} f d\lambda = \int_{\mathbb{R}} g d\lambda = \infty$

This example shows that the [strict increasing property](#) can fail when the integrals are infinite.

Let $f_n = \mathbf{1}_{[n, \infty)}$ for $n \in \mathbb{N}_+$. Show that

1. f_n is decreasing in $n \in \mathbb{N}_+$ on \mathbb{R} .
2. $f_n \rightarrow 0$ as $n \rightarrow \infty$ on \mathbb{R} .
3. $\int_{\mathbb{R}} f_n d\lambda = \infty$ for each $n \in \mathbb{N}_+$.

This example shows that the [monotone convergence theorem](#) can fail if the first integral is infinite. It also illustrates strict inequality in [Fatou's lemma](#).

Let $f_n = \mathbf{1}_{[n, n+1]}$ for $n \in \mathbb{N}_+$. Show that

1. $\lim_{n \rightarrow \infty} f_n = 0$ on \mathbb{R} so $\int_{\mathbb{R}} \lim_{n \rightarrow \infty} f_n d\mu = 0$
2. $\int_{\mathbb{R}} f_n d\lambda = 1$ for $n \in \mathbb{N}_+$ so $\lim_{n \rightarrow \infty} \int_{\mathbb{R}} f_n d\lambda = 1$
3. $\sup\{f_n : n \in \mathbb{N}_+\} = \mathbf{1}_{[1, \infty)}$ on \mathbb{R}

This example shows that the [dominated convergence theorem](#) can fail if $|f_n|$ is not bounded by an integrable function. It also shows that strict inequality can hold in [Fatou's lemma](#).

Consider the product space $[0, 1]^2$ with the usual Lebesgue measurable subsets and Lebesgue measure. Let $f : [0, 1]^2 \rightarrow \mathbb{R}$ be defined by

$$f(x, y) = \frac{x^2 - y^2}{(x^2 + y^2)^2} \quad (3.11.46)$$

Show that

1. $\int_{[0, 1]^2} f(x, y) d(x, y)$ does not exist.
2. $\int_0^1 \int_0^1 f(x, y) dx dy = -\frac{\pi}{4}$
3. $\int_0^1 \int_0^1 f(x, y) dy dx = \frac{\pi}{4}$

This example shows that [Fubini's theorem](#) can fail if the double integral does not exist.

For $i, j \in \mathbb{N}_+$ define the sequence a_{ij} as follows: $a_{ii} = 1$ and $a_{i+1, i} = -1$ for $i \in \mathbb{N}_+$, $a_{ij} = 0$ otherwise.

1. Give a_{ij} in array form with $i \in \mathbb{N}_+$ as the row number and $j \in \mathbb{N}_+$ as the column number
2. Show that $\sum_{(i, j) \in \mathbb{N}_+^2} a_{ij}$ does not exist
3. Show that $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij} = 1$
4. Show that $\sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij} = 0$

This example shows that the iterated sums can exist and be different when the double sum does not exist, a counterexample to the [corollary to Fubini's theorem](#) for sums when the hypotheses are not satisfied.

Computational Exercises

Compute $\int_D f(x, y) d(x, y)$ in each case below for the given $D \subseteq \mathbb{R}^2$ and $f : D \rightarrow \mathbb{R}$.

1. $f(x, y) = e^{-2x} e^{-3y}$, $D = [0, \infty) \times [0, \infty)$

$$2. f(x, y) = e^{-2x} e^{-3y}, D = \{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq y < \infty\}$$

Integrals of the type in the last exercise are useful in the study of exponential distributions.

This page titled [3.11: Properties of the Integral](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.12: General Measures

Basic Theory

Our starting point in this section is a measurable space (S, \mathcal{S}) . That is, S is a set and \mathcal{S} is a σ -algebra of subsets of S . So far, we have only considered positive measures on such spaces. Positive measures have applications, as we know, to length, area, volume, mass, probability, counting, and similar concepts of the nonnegative “size” of a set. Moreover, we have defined the integral of a measurable function $f : S \rightarrow \mathbb{R}$ with respect to a positive measure, and we have studied properties of the integral.

Definition

But now we will consider measures that can take negative values as well as positive values. These measures have applications to electric charge, monetary value, and other similar concepts of the “content” of a set that might be positive or negative. Also, this generalization will help in our study of density functions in the next section. The definition is exactly the same as for a positive measure, except that values in $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, \infty\}$ are allowed.

A *measure* on (S, \mathcal{S}) is a function $\mu : \mathcal{S} \rightarrow \mathbb{R}^*$ that satisfies the following properties:

1. $\mu(\emptyset) = 0$
2. If $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} then $\mu(\bigcup_{i \in I} A_i) = \sum_{i \in I} \mu(A_i)$

As before, (b) is known as *countable additivity* and is the critical assumption: the measure of a set that consists of a countable number of disjoint pieces is the sum of the measures of the pieces. Implicit in the statement of this assumption is that the sum in (b) exists for every countable disjoint collection $\{A_i : i \in I\}$. That is, either the sum of the positive terms is finite or the sum of the negative terms is finite. In turn, this means that the order of the terms in the sum does not matter (a good thing, since there is no implied order). The term *signed measure* is used by many, but we will just use the simple term *measure*, and add appropriate adjectives for the special cases. Note that if $\mu(A) \geq 0$ for all $A \in \mathcal{S}$, then μ is a *positive measure*, the kind we have already studied (and so the new definition really is a generalization). In this case, the sum in (b) always exists in $[0, \infty]$. If $\mu(A) \in \mathbb{R}$ for all $A \in \mathcal{S}$ then μ is a *finite measure*. Note that in this case, the sum in (b) is absolutely convergent for every countable disjoint collection $\{A_i : i \in I\}$. If μ is a positive measure and $\mu(S) = 1$ then μ is a *probability measure*, our favorite kind. Finally, as with positive measures, μ is σ -finite if there exists a countable collection $\{A_i : i \in I\}$ of sets in \mathcal{S} such that $S = \bigcup_{i \in I} A_i$ and $\mu(A_i) \in \mathbb{R}$ for $i \in I$.

Basic Properties

We give a few simple properties of general measures; hopefully many of these will look familiar. Throughout, we assume that μ is a measure on (S, \mathcal{S}) . Our first result is that although μ can take the value ∞ or $-\infty$, it turns out that it cannot take both of these values.

Either $\mu(A) > -\infty$ for all $A \in \mathcal{S}$ or $\mu(A) < \infty$ for all $A \in \mathcal{S}$.

Proof

Suppose that there exist $A, B \in \mathcal{S}$ with $\mu(A) = \infty$ and $\mu(B) = -\infty$. Then $A = (A \cap B) \cup (A \setminus B)$ and the sets in the union are disjoint. By the additivity assumption, $\mu(A) = \mu(A \cap B) + \mu(A \setminus B)$. Similarly, $\mu(B) = \mu(A \cap B) + \mu(B \setminus A)$. The only way that both of these equations can make sense is for $\mu(A \setminus B) = \infty$, $\mu(B \setminus A) = -\infty$, and $\mu(A \cap B) \in \mathbb{R}$. But then $\mu(A \triangle B) = \mu(A \setminus B) + \mu(B \setminus A)$ is undefined, and so we have a contradiction.

We will say that two measures are *of the same type* if neither takes the value ∞ or if neither takes the value $-\infty$. *Being of the same type* is trivially an equivalence relation on the collection of measures on (S, \mathcal{S}) .

The difference rule holds, as long as the sets have finite measure:

Suppose that $A, B \in \mathcal{S}$. If $\mu(B) \in \mathbb{R}$ then $\mu(B \setminus A) = \mu(B) - \mu(A \cap B)$.

Proof

Note that $B = (A \cap B) \cup (B \setminus A)$ and the sets in the union are disjoint. Thus $\mu(B) = \mu(A \cap B) + \mu(B \setminus A)$. Since $\mu(B) \in \mathbb{R}$, we must have $\mu(A \cap B) \in \mathbb{R}$ and $\mu(B \setminus A) \in \mathbb{R}$ also, and then the difference rule holds by subtraction.

The following corollary is the difference rule for subsets, and will be needed below.

Suppose that $A, B \in \mathcal{S}$ and $A \subseteq B$. If $\mu(B) \in \mathbb{R}$ then $\mu(A) \in \mathbb{R}$ and $\mu(B \setminus A) = \mu(B) - \mu(A)$.

Proof

Note that $B = A \cup (B \setminus A)$ and the sets in the union are disjoint. Thus $\mu(B) = \mu(A) + \mu(B \setminus A)$. Since $\mu(B) \in \mathbb{R}$, we must have $\mu(A) \in \mathbb{R}$ and $\mu(B \setminus A) \in \mathbb{R}$ also, and then the difference rule holds by subtraction.

As a consequence, suppose that $A, B \in \mathcal{S}$ and $A \subseteq B$. If $\mu(A) = \infty$, then by the [infinity rule](#) we cannot have $\mu(B) = -\infty$ and by the [difference rule](#) we cannot have $\mu(B) \in \mathbb{R}$, so we must have $\mu(B) = \infty$. Similarly, if $\mu(A) = -\infty$ then $\mu(B) = -\infty$. The *inclusion-exclusion* rules hold for general measures, as long as the sets have finite measure.

Suppose that $A_i \in \mathcal{S}$ for each $i \in I$ where $\#(I) = n$, and that $\mu(A_i) \in \mathbb{R}$ for $i \in I$. Then

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{k=1}^n (-1)^{k-1} \sum_{J \subseteq I, \#(J)=k} \mu\left(\bigcap_{j \in J} A_j\right) \quad (3.12.1)$$

Proof

For $n = 2$, note that $A_1 \cup A_2 = A_1 \cup (A_2 \setminus A_1)$ and the sets in the last union are disjoint. By the additivity axiom and the difference rule (3),

$$\mu(A_1 \cup A_2) = \mu(A_1) + \mu(A_2 \setminus A_1) = \mu(A_1) + \mu(A_2) - \mu(A_1 \cap A_2) \quad (3.12.2)$$

The general result then follows by induction, just like the proof for probability measures.

The *continuity properties* hold for general measures. Part (a) is the continuity property for *increasing sets*, and part (b) is the continuity property for *decreasing sets*.

Suppose that $A_n \in \mathcal{S}$ for $n \in \mathbb{N}_+$.

1. If $A_n \subseteq A_{n+1}$ for $n \in \mathbb{N}_+$ then $\lim_{n \rightarrow \infty} \mu(A_n) = \mu(\bigcup_{i=1}^{\infty} A_i)$.
2. If $A_{n+1} \subseteq A_n$ for $n \in \mathbb{N}_+$ and $\mu(A_1) \in \mathbb{R}$, then $\lim_{n \rightarrow \infty} \mu(A_n) = \mu(\bigcap_{i=1}^{\infty} A_i)$.

Proof

The proofs are almost the same as for positive measures, except for technicalities involving ∞ and $-\infty$.

1. Let $A = \bigcup_{i=1}^{\infty} A_i$. From the [infinity rule](#) and the [difference rule](#), if $\mu(A_m) = \infty$ (respectively $-\infty$) for some $m \in \mathbb{N}_+$, then $\mu(A_n) = \infty$ ($-\infty$) for $n \geq m$ and $\mu(A) = \infty$ ($-\infty$), so the result trivially holds. Thus, assume that $\mu(A_n) \in \mathbb{R}$ for all $n \in \mathbb{N}_+$. Let $B_1 = A_1$ and let $B_i = A_i \setminus A_{i-1}$ for $i \in \{2, 3, \dots\}$. Then $\{B_i : i \in \mathbb{N}_+\}$ is a disjoint collection of sets and also has union A . Moreover, from the [difference rule](#), $\mu(B_i) = \mu(A_{i+1}) - \mu(A_i)$ for $i \in \{2, 3, \dots\}$. Thus

$$\mu(A) = \sum_{i=1}^{\infty} \mu(B_i) = \lim_{n \rightarrow \infty} \sum_{i=1}^n \mu(B_i) = \lim_{n \rightarrow \infty} \left(\mu(A_1) + \sum_{i=2}^n [\mu(A_i) - \mu(A_{i-1})] \right) = \lim_{n \rightarrow \infty} \mu(A_n) \quad (3.12.3)$$

2. Let $C_n = A_1 \setminus A_n$ for $n \in \mathbb{N}_+$. Then $C_n \subseteq C_{n+1}$ for $n \in \mathbb{N}_+$ and $\bigcup_{i=1}^{\infty} C_i = A_1 \setminus \bigcap_{i=1}^{\infty} A_i$. Part (a) applies, so $\lim_{n \rightarrow \infty} \mu(C_n) = \mu(\bigcup_{i=1}^{\infty} C_i)$. But by the [difference rule](#), $\mu(C_n) = \mu(A_1) - \mu(A_n)$ for $n \in \mathbb{N}_+$ and $\mu(\bigcup_{i=1}^{\infty} C_i) = \mu(A_1) - \mu(\bigcap_{i=1}^{\infty} A_i)$. All of these are real numbers, so subtracting $\mu(A_1)$ gives the result.

Recall that a positive measure is an increasing function, relative to the subset partial order on \mathcal{S} and the ordinary order on $[0, \infty]$, and this property follows from the [difference rule](#). But for general measures, the increasing property fails, and so do other properties that flow from it, including the subadditive property (Boole's inequality in probability) and the Bonferroni inequalities.

Constructions

It's easy to construct general measures as differences of positive measures.

Suppose that μ and ν are positive measures on (S, \mathcal{S}) and that at least one of them is finite. Then $\delta = \mu - \nu$ is a measure.

Proof

Suppose that ν is a finite measure; the proof when μ is finite is similar. First, $\delta(\emptyset) = \mu(\emptyset) - \nu(\emptyset) = 0$. Suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} and let $A = \bigcup_{i \in I} A_i$. Then

$$\delta(A) = \mu(A) - \nu(A) = \sum_{i \in I} \mu(A_i) - \sum_{i \in I} \nu(A_i) \quad (3.12.4)$$

Since $\nu(A_i) < \infty$ for $i \in I$, we can combine terms to get

$$\delta(A) = \sum_{i \in I} [\mu(A_i) - \nu(A_i)] = \sum_{i \in I} \delta(A_i) \quad (3.12.5)$$

The collection of measures on our space is closed under scalar multiplication.

If μ is a measure on (S, \mathcal{S}) and $c \in \mathbb{R}$, then $c\mu$ is a measure on (S, \mathcal{S})

Proof

First, $(c\mu)(\emptyset) = c\mu(\emptyset) = c0 = 0$. Next suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} . Then

$$(c\mu)\left(\bigcup_{i \in I} A_i\right) = c\mu\left(\bigcup_{i \in I} A_i\right) = c \sum_{i \in I} \mu(A_i) = \sum_{i \in I} c\mu(A_i) = \sum_{i \in I} (c\mu)(A_i) \quad (3.12.6)$$

The last step is the important one, and holds since the sum exists.

If μ is a finite measure, then so is $c\mu$ for $c \in \mathbb{R}$. If μ is not finite then μ and $c\mu$ are of the same type if $c > 0$ and are of opposite types if $c < 0$. We can add two measures to get another measure, as long as they are of the same type. In particular, the collection of *finite* measures is closed under addition as well as scalar multiplication, and hence forms a vector space.

If μ and ν are measures on (S, \mathcal{S}) of the same type then $\mu + \nu$ is a measure on (S, \mathcal{S}) .

Proof

First, $(\mu + \nu)(\emptyset) = \mu(\emptyset) + \nu(\emptyset) = 0 + 0 = 0$. Next suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} . Then

$$\begin{aligned} (\mu + \nu)\left(\bigcup_{i \in I} A_i\right) &= \mu\left(\bigcup_{i \in I} A_i\right) + \nu\left(\bigcup_{i \in I} A_i\right) \\ &= \sum_{i \in I} \mu(A_i) + \sum_{i \in I} \nu(A_i) = \sum_{i \in I} [\mu(A_i) + \nu(A_i)] = \sum_{i \in I} (\mu + \nu)(A_i) \end{aligned}$$

The sums can be combined because the measures are of the same type. That is, either the sum of all of the positive terms is finite or the sum of all the negative terms is finite. In short, we don't have to worry about the dreaded indeterminate form $\infty - \infty$.

Finally, it is easy to explicitly construct measures on a σ -algebra generated by a countable partition. Such σ -algebras are important for counterexamples and to gain insight, and also because many σ -algebras that occur in applications can be constructed from them.

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of S into nonempty sets, and that $\mathcal{S} = \sigma(\mathcal{A})$. For $i \in I$, define $\mu(A_i) \in \mathbb{R}^*$ arbitrarily, subject only to the condition that the sum of the positive terms is finite, or the sum of the negative terms is finite. For $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$, define

$$\mu(A) = \sum_{j \in J} \mu(A_j) \quad (3.12.7)$$

Then μ is a measure on (S, \mathcal{S}) .

Proof

Recall that every $A \in \mathcal{S}$ has a unique representation of the form $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$.

1. $J = \emptyset$ in the representation gives $A = \emptyset$. The sum over an empty index set is 0, so $\mu(\emptyset) = 0$.

2. Suppose that $\{B_k : k \in K\}$ is a countable, disjoint collection of events in \mathcal{S} . Then for each $k \in K$ there exists $J_k \subseteq I$ and $\{A_j^k : j \in J_k\} \subseteq \mathcal{A}$ such that $B_k = \bigcup_{j \in J_k} A_j^k$. Hence

$$\mu\left(\bigcup_{k \in K} B_k\right) = \mu\left(\bigcup_{k \in K} \bigcup_{j \in J_k} A_j^k\right) = \sum_{k \in K} \sum_{j \in J_k} \mu(A_j^k) = \sum_{k \in K} \mu(B_k) \quad (3.12.8)$$

The fact that either the sum of all positive terms is finite or the sum of all the negative terms is finite means that we do not have to worry about the order of summation.

Positive, Negative, and Null Sets

To understand the structure of general measures, we need some basic definitions and properties. As before, we assume that μ is a measure on (S, \mathcal{S}) .

Definitions

1. $A \in \mathcal{S}$ is a *positive set* for μ if $\mu(B) \geq 0$ for every $B \in \mathcal{S}$ with $B \subseteq A$.
2. $A \in \mathcal{S}$ is a *negative set* for μ if $\mu(B) \leq 0$ for every $B \in \mathcal{S}$ with $B \subseteq A$.
3. $A \in \mathcal{S}$ is a *null set* for μ if $\mu(B) = 0$ for every $B \in \mathcal{S}$ with $B \subseteq A$.

Note that *positive* and *negative* are used in the weak sense (just as we use the terms *increasing* and *decreasing* in this text). Of course, if μ is a positive measure, then every $A \in \mathcal{S}$ is positive for μ , and $A \in \mathcal{S}$ is negative for μ if and only if A is null for μ if and only if $\mu(A) = 0$. For a general measure, $A \in \mathcal{S}$ is both positive and negative for μ if and only if A is null for μ . In particular, \emptyset is null for μ . A set $A \in \mathcal{S}$ is a *support set* for μ if and only if A^c is a null set for μ . A support set is a set where the measure “lives” in a sense. Positive, negative, and null sets for μ have a basic inheritance property that is essentially equivalent to the definition.

Suppose $A \in \mathcal{S}$.

1. If A is positive for μ then B is positive for μ for every $B \in \mathcal{S}$ with $B \subseteq A$.
2. If A is negative for μ then B is negative for μ for every $B \in \mathcal{S}$ with $B \subseteq A$.
3. If A is null for μ then B is null for μ for every $B \in \mathcal{S}$ with $B \subseteq A$.

The collections of positive sets, negative sets, and null sets for μ are closed under countable unions.

Suppose that $\{A_i : i \in I\}$ is a countable collection of sets in \mathcal{S} .

1. If A_i is positive for μ for $i \in I$ then $\bigcup_{i \in I} A_i$ is positive for μ .
2. If A_i is negative for μ for $i \in I$ then $\bigcup_{i \in I} A_i$ is negative for μ .
3. If A_i is null for μ for $i \in I$ then $\bigcup_{i \in I} A_i$ is null for μ .

Proof

We will prove (a); the proofs for (b) and (c) are analogous. Without loss of generality, we can suppose that $I = \mathbb{N}_+$. Let $A = \bigcup_{n=1}^{\infty} A_n$. Now let $B_1 = A_1$ and $B_n = A_n \setminus \left(\bigcup_{i=1}^{n-1} A_i\right)$ for $n \in \{2, 3, \dots\}$. Then $\{B_n : n \in \mathbb{N}_+\}$ is a countable, disjoint collection in \mathcal{S} , and $\bigcup_{n=1}^{\infty} B_n = A$. If $C \subseteq A$ then $C = \bigcup_{n=1}^{\infty} (C \cap B_n)$ and the sets in this union are disjoint. Hence by additivity, $\mu(C) = \sum_{n=1}^{\infty} \mu(C \cap B_n)$. But $C \cap B_n \subseteq B_n \subseteq A_n$ so $\mu(C \cap B_n) \geq 0$. Hence $\mu(C) \geq 0$.

It's easy to see what happens to the positive, negative, and null sets when a measure is multiplied by a non-zero constant.

Suppose that μ is a measure on (S, \mathcal{S}) , $c \in \mathbb{R}$, and $A \in \mathcal{S}$.

1. If $c > 0$ then A is positive (negative) for μ if and only if A is positive (negative) for $c\mu$.
2. If $c < 0$ then A is positive (negative) for μ if and only if A is negative (positive) for $c\mu$.
3. If $c \neq 0$ then A is null for μ if and only if A is null for $c\mu$.

Positive, negative, and null sets are also preserved under countable sums, assuming that the measures make sense.

Suppose that μ_i is a measure on (S, \mathcal{S}) for each i in a countable index set I , and that $\mu = \sum_{i \in I} \mu_i$ is a well-defined measure on (S, \mathcal{S}) . Let $A \in \mathcal{S}$.

1. If A is positive for μ_i for every $i \in I$ then A is positive for μ .
2. If A is negative for μ_i for every $i \in I$ then A is negative for μ .
3. If A is null for μ_i for every $i \in I$ then A is null for μ .

In particular, note that $\mu = \sum_{i \in I} \mu_i$ is a well-defined measure if μ_i is a positive measure for each $i \in I$, or if I is finite and μ_i is a finite measure for each $i \in I$. It's easy to understand the positive, negative, and null sets for a σ -algebra generated by a [countable partition](#).

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of S into nonempty sets, and that $\mathcal{S} = \sigma(\mathcal{A})$. Suppose that μ is a measure on (S, \mathcal{S}) . Define

$$I_+ = \{i \in I : \mu(A_i) > 0\}, \quad I_- = \{i \in I : \mu(A_i) < 0\}, \quad I_0 = \{i \in I : \mu(A_i) = 0\} \quad (3.12.9)$$

Let $A \in \mathcal{S}$, so that $A = \bigcup_{j \in J} A_j$ for some $J \subseteq I$ (and this representation is unique). Then

1. A is positive for μ if and only if $J \subseteq I_+ \cup I_0$.
2. A is negative for μ if and only if $J \subseteq I_- \cup I_0$.
3. A is null for μ if and only if $J \subseteq I_0$.

The Hahn Decomposition

The fundamental results in this section and the next are two decomposition theorems that show precisely the relationship between general measures and positive measures. First we show that if a set has finite, positive measure, then it has a positive subset with at least that measure.

If $A \in \mathcal{S}$ and $0 \leq \mu(A) < \infty$ then there exists $P \in \mathcal{S}$ with $P \subseteq A$ such that P is positive for μ and $\mu(P) \geq \mu(A)$.

Proof

The proof is recursive, and works by successively removing sets of negative measure from A . For the initialization step, let $A_0 = A$. Then trivially, $A_0 \subseteq A$ and $\mu(A_0) \geq \mu(A)$. For the recursive step, suppose that $A_n \in \mathcal{S}$ has been defined with $A_n \subseteq A$ and $\mu(A_n) \geq \mu(A)$. If A_n is positive for μ , let $P = A_n$. Otherwise let $a_n = \inf\{\mu(B) : B \in \mathcal{S}, B \subseteq A_n, \mu(B) < 0\}$. Note that since A_n is not positive for μ , the set in the infimum is nonempty and hence $a_n < 0$ (and possibly $-\infty$). Let $b_n = a_n/2$ if $-\infty < a_n < 0$ and let $b_n = -1$ if $a_n = -\infty$. Since $b_n > a_n$, by definition of the infimum, there exists $B_n \subseteq A$ with $\mu(B_n) \leq b_n$. Let $A_{n+1} = A_n \setminus B_n$. Then $A_{n+1} \subseteq A_n \subseteq A$ and

$$\mu(A_{n+1}) = \mu(A_n) - \mu(B_n) \geq \mu(A_n) - b_n \geq \mu(A_n) \geq \mu(A) \quad (3.12.10)$$

Now, if the recursive process terminates after a finite number of steps, P is well defined and is positive for μ . Otherwise, we have a disjoint sequence of sets (B_1, B_2, \dots) . Let $P = A \setminus (\bigcup_{i=1}^{\infty} B_i)$. Then $P \subseteq A$, and by countable additivity and the [difference rule](#),

$$\mu(P) = \mu(A) - \sum_{n=1}^{\infty} \mu(B_n) \geq \mu(A) - \sum_{n=1}^{\infty} b_n \geq \mu(A) \quad (3.12.11)$$

Suppose that $B \subseteq P$ and $\mu(B) < 0$. Then $B \subseteq A_n$ and by definition, $a_n \leq \mu(B)$ for every $n \in \mathbb{N}_+$. It follows that $b_n \leq \frac{1}{2}\mu(B)$ or $b_n = -1$ for every $n \in \mathbb{N}_+$. Hence $\sum_{n=1}^{\infty} b_n = -\infty$ and therefore $\mu(P) = \infty$, a contradiction since $\mu(A) < \infty$. Hence we must have $\mu(B) \geq 0$ and thus P is positive for μ .

The assumption that $\mu(A) < \infty$ is critical; a [counterexample](#) is given below. Our first decomposition result is the *Hahn decomposition theorem*, named for the Austrian mathematician Hans Hahn. It states that S can be partitioned into a positive set and a negative set, and this decomposition is essentially unique.

Hahn Decomposition Theorem. There exists $P \in \mathcal{S}$ such that P is positive for μ and P^c is negative for μ . The pair (P, P^c) is a *Hahn decomposition* of S . If (Q, Q^c) is another Hahn decomposition, then $P \triangle Q$ is null for μ .

Proof

Suppose first that μ does not take the value ∞ . As with the previous result, the proof is recursive. For the initialization step, let $P_0 = \emptyset$. Then trivially, P_0 is positive for μ . For the recursive step, suppose that $P_n \in \mathcal{S}$ is positive for μ . If P_n^c is negative for μ , let $P = P_n$. Otherwise let $a_n = \sup\{\mu(A) : A \in \mathcal{S}, A \subseteq P_n^c\}$. Since P_n^c is not negative for μ , it follows that $a_n > 0$ (and possibly ∞). Let $b_n = a_n/2$ if $0 < a_n < \infty$ and $b_n = 1$ if $a_n = \infty$. Then $b_n < a_n$ so there exists $B_n \in \mathcal{S}$ with $B_n \subseteq P_n^c$ and $\mu(B_n) \geq b_n > 0$. By the [previous lemma](#), there exists $A_n \in \mathcal{S}$ with $A_n \subseteq B_n$, A_n positive for μ , and $\mu(A_n) \geq \mu(B_n)$. Let $P_{n+1} = P_n \cup A_n$. Then $P_{n+1} \in \mathcal{S}$ is positive for μ .

If the recursive process ends after a finite number of steps, then P is well-defined and (P, P^c) is a Hahn decomposition. Otherwise we generate an infinite sequence (A_1, A_2, \dots) of disjoint sets in \mathcal{S} , each positive for μ . Let $P = \bigcup_{n=1}^{\infty} A_n$. Then $P \in \mathcal{S}$ is positive for μ by the [closure result](#) above. Let $A \subseteq P^c$. If $\mu(A) > 0$ then $\mu(A) \leq a_n$ for every $n \in \mathbb{N}_+$. Hence $b_n \geq \frac{1}{2}\mu(A)$ or $b_n = 1$ for every $n \in \mathbb{N}_+$. But then

$$\mu(P) = \sum_{n=1}^{\infty} \mu(A_n) \geq \sum_{n=1}^{\infty} \mu(B_n) \geq \sum_{n=1}^{\infty} b_n = \infty \quad (3.12.12)$$

a contradiction. Hence $\mu(A) \leq 0$ so P^c is negative for μ and thus (P, P^c) is a Hahn decomposition.

Suppose that (Q, Q^c) is another Hahn decomposition of S . Then $P \cap Q^c$ and $Q \cap P^c$ are both positive and negative for μ and hence are null for μ . Hence $P \triangle Q = (P \cap Q^c) \cup (Q \cap P^c)$ is null for μ .

Finally, suppose that μ takes the value ∞ . Then μ does not take the value $-\infty$ by the [infinity rule](#) and hence $-\mu$ does not take the value ∞ . By our proof so far, there exists a Hahn decomposition (P, P^c) for $-\mu$ that is essentially unique. But then (P^c, P) is a Hahn decomposition for μ .

It's easy to see the Hahn decomposition for a measure on a σ -algebra generated by a [countable partition](#).

Suppose that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of S into nonempty sets, and that $\mathcal{S} = \sigma(\mathcal{A})$. Suppose that μ is a measure on (S, \mathcal{S}) . Let $I_+ = \{i \in I : \mu(A_i) > 0\}$ and $I_0 = \{i \in I : \mu(A_i) = 0\}$. Then (P, P^c) is a Hahn decomposition of μ if and only if the positive set P has the form $P = \bigcup_{j \in J} A_j$ where $J = I_+ \cup K$ and $K \subseteq I_0$.

The Jordan Decomposition

The Hahn decomposition leads to another decomposition theorem called the *Jordan decomposition theorem*, named for the French mathematician Camille Jordan. This one shows that every measure is the [difference of positive measures](#). Once again we assume that μ is a measure on (S, \mathcal{S}) .

Jordan Decomposition Theorem. The measure μ can be written uniquely in the form $\mu = \mu_+ - \mu_-$ where μ_+ and μ_- are positive measures, at least one finite, and with the property that if (P, P^c) is any Hahn decomposition of S , then P^c is a null set of μ_+ and P is a null set of μ_- . The pair (μ_+, μ_-) is the *Jordan decomposition* of μ .

Proof

Let (P, P^c) be a Hahn decomposition of S relative to μ . Define $\mu_+(A) = \mu(A \cap P)$ and $\mu_-(A) = -\mu(A \cap P^c)$ for $A \in \mathcal{S}$. Then μ_+ and μ_- are positive measures and $\mu = \mu_+ - \mu_-$. Moreover, since μ cannot take both ∞ and $-\infty$ as values by the [infinity rule](#), one of these two positive measures is finite.

Suppose that (Q, Q^c) is an arbitrary Hahn decomposition. If $A \subseteq Q^c$, then $\mu_+(A) = \mu(P \cap A) = 0$ since $P \cap Q^c$ is a null set of μ by the [Hahn decomposition theorem](#). Similarly if $A \subseteq Q$ then $\mu_-(A) = \mu(P^c \cap A) = 0$ since $P^c \cap Q$ is a null set of μ .

Suppose that $\mu = \nu_+ - \nu_-$ is another decomposition with the same properties. If $A \in \mathcal{S}$ then $\mu_+(A) = \mu(A \cap P) = [\nu_+(A \cap P) - \nu_-(A \cap P)] = \nu_+(A \cap P)$. But also $\nu_+(A) = \nu_+(A \cap P) + \nu_+(A \cap P^c) = \nu_+(A \cap P)$. Hence $\nu_+ = \mu_+$ and therefore also $\nu_- = \mu_-$.

The Jordan decomposition leads to an important set of new definitions.

Suppose that μ has Jordan decomposition $\mu = \mu_+ - \mu_-$.

1. The positive measure μ_+ is called the *positive variation measure* of μ .
2. The positive measure μ_- is called the *negative variation measure* of μ .

3. The positive measure $|\mu| = \mu_+ + \mu_-$ is called the *total variation measure* of μ .
4. $\|\mu\| = |\mu|(S)$ is the *total variation* of μ .

Note that, in spite of the similarity in notation, $\mu_+(A)$ and $\mu_-(A)$ are *not* simply the positive and negative parts of the (extended) real number $\mu(A)$, nor is $|\mu|(A)$ the absolute value of $\mu(A)$. Also, be careful not to confuse the total variation of μ , a number in $[0, \infty]$, with the total variation *measure*. The positive, negative, and total variation measures can be written directly in terms of μ .

For $A \in \mathcal{S}$,

1. $\mu_+(A) = \sup\{\mu(B) : B \in \mathcal{S}, B \subseteq A\}$
2. $\mu_-(A) = -\inf\{\mu(B) : B \in \mathcal{S}, B \subseteq A\}$
3. $|\mu|(A) = \sup\left\{\sum_{i \in I} \mu(A_i) : \{A_i : i \in I\} \text{ is a finite, measurable partition of } A\right\}$
4. $\|\mu\| = \sup\left\{\sum_{i \in I} \mu(A_i) : \{A_i : i \in I\} \text{ is a finite, measurable partition of } S\right\}$

The total variation measure is related to sum and scalar multiples of measures in a natural way.

Suppose that μ and ν are measures of the same type and that $c \in \mathbb{R}$. Then

1. $|\mu| = 0$ if and only if $\mu = 0$ (the zero measure).
2. $|c\mu| = |c| |\mu|$
3. $|\mu + \nu| \leq |\mu| + |\nu|$

Proof

1. Since μ_+ , μ_- and $|\mu| = \mu_+ + \mu_-$ are positive measures, $|\mu| = 0$ if and only if $\mu_+ = \mu_- = 0$ if and only if $\mu = 0$.
2. If $c > 0$ then $(c\mu)_+ = c\mu_+$ and $(c\mu)_- = c\mu_-$. If $c < 0$ then $(c\mu)_+ = -c\mu_-$ and $(c\mu)_- = -c\mu_+$. Of course, if $c = 0$ then $(c\mu)_+ = (c\mu)_- = 0$. In all cases,

$$|c\mu| = (c\mu)_+ + (c\mu)_- = |c|(\mu_+ + \mu_-) = |c| |\mu| \quad (3.12.13)$$

3. From the [theorem above](#), $(\mu + \nu)_+ \leq \mu_+ + \nu_+$ and $(\mu + \nu)_- \leq \mu_- + \nu_-$. So

$$\begin{aligned} |\mu + \nu| &= (\mu + \nu)_+ + (\mu + \nu)_- \leq (\mu_+ + \nu_+) + (\mu_- + \nu_-) \\ &= (\mu_+ + \mu_-) + (\nu_+ + \nu_-) = |\mu| + |\nu| \end{aligned}$$

You may have noticed that the properties in the last result look a bit like norm properties. In fact, total variation really is a norm on the vector space of finite measures on (S, \mathcal{S}) :

Suppose that μ and ν are measures of the same type and that $c \in \mathbb{R}$. Then

1. $\|\mu\| = 0$ if and only if $\mu = 0$ (the *zero property*)
2. $\|c\mu\| = |c| \|\mu\|$ (the *scaling property*)
3. $\|\mu + \nu\| \leq \|\mu\| + \|\nu\|$ (the *triangle inequality*)

Proof

1. Since $|\mu|$ is a positive measure, $\|\mu\| = |\mu|(S) = 0$ if and only if $|\mu| = 0$. From part (a) of the [previous theorem](#), $|\mu| = 0$ if and only if $\mu = 0$.
2. From part (b) of the [previous theorem](#), $\|c\mu\| = |c\mu|(S) = |c| |\mu|(S) = |c| \|\mu\|$.
3. From part (c) of the [previous theorem](#), $\|\mu + \nu\| = |\mu + \nu|(S) \leq |\mu|(S) + |\nu|(S) = \|\mu\| + \|\nu\|$.

Every norm on a vector space leads to a corresponding measure of distance (a metric). Let \mathcal{M} denote the collection of finite measures on (S, \mathcal{S}) . Then \mathcal{M} , under the usual definition of addition and scalar multiplication of measures, is a vector space, and as the last theorem shows, $\|\cdot\|$ is a norm on \mathcal{M} . Here are the corresponding metric space properties:

Suppose that $\mu, \nu, \rho \in \mathcal{M}$ and $c \in \mathbb{R}$. Then

1. $\|\mu - \nu\| = \|\nu - \mu\|$, the *symmetric property*
2. $\|\mu\| = 0$ if and only if $\mu = 0$, the *zero property*
3. $\|\mu - \rho\| \leq \|\mu - \nu\| + \|\nu - \rho\|$, the *triangle inequality*

Now that we have a metric, we have a corresponding criterion for convergence.

Suppose that $\mu_n \in \mathcal{M}$ for $n \in \mathbb{N}_+$ and $\mu \in \mathcal{M}$. We say that $\mu_n \rightarrow \mu$ as $n \rightarrow \infty$ in *total variation* if $\|\mu_n - \mu\| \rightarrow 0$ as $n \rightarrow \infty$.

Of course, \mathcal{M} includes the probability measures on (S, \mathcal{S}) , so we have a new notion of convergence to go along with the others we have studied or will study. Here is a list:

- convergence with probability 1
- convergence in probability
- convergence in distribution
- convergence in k th mean
- convergence in total variation

The Integral

Armed with the Jordan decomposition, the integral can be extended to general measures in a natural way.

Suppose that μ is a measure on (S, \mathcal{S}) and that $f : S \rightarrow \mathbb{R}$ is measurable. We define

$$\int_S f d\mu = \int_S f d\mu_+ - \int_S f d\mu_- \quad (3.12.14)$$

assuming that the integrals on the right exist and that the right side is not of the form $\infty - \infty$.

We will not pursue this extension, but as you might guess, the essential properties of the integral hold.

Complex Measures

Again, suppose that (S, \mathcal{S}) is a measurable space. The same axioms that work for general measures can be used to define complex measures. Recall that $\mathbb{C} = \{x + iy : x, y \in \mathbb{R}\}$ denotes the set of complex numbers, where i is the imaginary unit.

A *complex measure* on (S, \mathcal{S}) is a function $\mu : \mathcal{S} \rightarrow \mathbb{C}$ that satisfies the following properties:

1. $\mu(\emptyset) = 0$
2. If $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} then $\mu(\bigcup_{i \in I} A_i) = \sum_{i \in I} \mu(A_i)$

Clearly a complex measure μ can be decomposed as $\mu = \nu + i\rho$ where ν and ρ are finite (real) measures on (S, \mathcal{S}) . We will have no use for complex measures in this text, but from the decomposition into finite measures, it's easy to see how to develop the theory.

Computational Exercises

Counterexamples

The [lemma](#) needed for the Hahn decomposition theorem can fail without the assumption that $\mu(A) < \infty$.

Let S be a set with subsets A and B satisfying $\emptyset \subset B \subset A \subset S$. Let $\mathcal{S} = \sigma\{A, B\}$ be the σ -algebra generated by $\{A, B\}$. Define $\mu(B) = -1$, $\mu(A \setminus B) = \infty$, $\mu(A^c) = 1$.

1. Draw the Venn diagram of A, B, S .
2. List the sets in \mathcal{S} .
3. Using additivity, give the value of μ on each set in \mathcal{S} .
4. Show that A does not have a positive subset $P \in \mathcal{S}$ with $\mu(P) \geq \mu(A)$.

This page titled [3.12: General Measures](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.13: Absolute Continuity and Density Functions

Basic Theory

Our starting point is a measurable space (S, \mathcal{S}) . That is S is a set and \mathcal{S} is a σ -algebra of subsets of S . In the last section, we discussed general measures on (S, \mathcal{S}) that can take positive and negative values. Special cases are positive measures, finite measures, and our favorite kind, probability measures. In particular, we studied properties of general measures, ways to construct them, special sets (positive, negative, and null), and the Hahn and Jordan decompositions.

In this section, we see how to construct a new measure from a given positive measure using a density function, and we answer the fundamental question of when a measure has a density function relative to the given positive measure.

Relations on Measures

The answer to the question involves two important relations on the collection of measures on (S, \mathcal{S}) that are defined in terms of null sets. Recall that $A \in \mathcal{S}$ is *null* for a measure μ on (S, \mathcal{S}) if $\mu(B) = 0$ for every $B \in \mathcal{S}$ with $B \subseteq A$. At the other extreme, $A \in \mathcal{S}$ is a *support set* for μ if A^c is a null set. Here are the basic definitions:

Suppose that μ and ν are measures on (S, \mathcal{S}) .

1. ν is *absolutely continuous* with respect to μ if every null set of μ is also a null set of ν . We write $\nu \ll \mu$.
2. μ and ν are *mutually singular* if there exists $A \in \mathcal{S}$ such that A is null for μ and A^c is null for ν . We write $\mu \perp \nu$.

Thus $\nu \ll \mu$ if every support set of μ is a support set of ν . At the opposite end, $\mu \perp \nu$ if μ and ν have disjoint support sets.

Suppose that μ, ν , and ρ are measures on (S, \mathcal{S}) . Then

1. $\mu \ll \mu$, the *reflexive property*.
2. If $\mu \ll \nu$ and $\nu \ll \rho$ then $\mu \ll \rho$, the *transitive property*.

Recall that every relation that is reflexive and transitive leads to an equivalence relation, and then in turn, the original relation can be extended to a partial order on the collection of equivalence classes. This general theorem on relations leads to the following two results.

Measures μ and ν on (S, \mathcal{S}) are *equivalent* if $\mu \ll \nu$ and $\nu \ll \mu$, and we write $\mu \equiv \nu$. The relation \equiv is an equivalence relation on the collection of measures on (S, \mathcal{S}) . That is, if μ, ν , and ρ are measures on (S, \mathcal{S}) then

1. $\mu \equiv \mu$, the *reflexive property*
2. If $\mu \equiv \nu$ then $\nu \equiv \mu$, the *symmetric property*
3. If $\mu \equiv \nu$ and $\nu \equiv \rho$ then $\mu \equiv \rho$, the *transitive property*

Thus, μ and ν are equivalent if they have the same null sets and thus the same support sets. This equivalence relation is rather weak: equivalent measures have the same support sets, but the values assigned to these sets can be very different. As usual, we will write $[\mu]$ for the equivalence class of a measure μ on (S, \mathcal{S}) , under the equivalence relation \equiv .

If μ and ν are measures on (S, \mathcal{S}) , we write $[\mu] \preceq [\nu]$ if $\mu \ll \nu$. The definition is consistent, and defines a partial order on the collection of equivalence classes. That is, if μ, ν , and ρ are measures on (S, \mathcal{S}) then

1. $[\mu] \preceq [\mu]$, the *reflexive property*.
2. If $[\mu] \preceq [\nu]$ and $[\nu] \preceq [\mu]$ then $[\mu] = [\nu]$, the *antisymmetric property*.
3. If $[\mu] \preceq [\nu]$ and $[\nu] \preceq [\rho]$ then $[\mu] \preceq [\rho]$, the *transitive property*

The singularity relation is trivially symmetric and is almost anti-reflexive.

Suppose that μ and ν are measures on (S, \mathcal{S}) . Then

1. If $\mu \perp \nu$ then $\nu \perp \mu$, the *symmetric property*.
2. $\mu \perp \mu$ if and only if $\mu = \mathbf{0}$, the *zero measure*.

Proof

Part (a) is trivial from the symmetry of the definition. For part (b), note that S is null for $\mathbf{0}$ and \emptyset is null for $\mathbf{0}$, so $\mathbf{0} \perp \mathbf{0}$. Conversely, suppose that μ is a measure and $\mu \perp \mu$. Then there exists $A \in \mathcal{S}$ such that A is null for μ and A^c is null for μ . But then $S = A \cup A^c$ is null for μ , so $\mu(B) = 0$ for every $B \in \mathcal{S}$.

Absolute continuity and singularity are preserved under multiplication by nonzero constants.

Suppose that μ and ν are measures on (S, \mathcal{S}) and that $a, b \in \mathbb{R} \setminus \{0\}$. Then

1. $\nu \ll \mu$ if and only if $a\nu \ll b\mu$.
2. $\nu \perp \mu$ if and only if $a\nu \perp b\mu$.

Proof

Recall that if $c \neq 0$, then $A \in \mathcal{S}$ is null for μ if and only if A is null for $c\mu$.

There is a corresponding result for sums of measures.

Suppose that μ is a measure on (S, \mathcal{S}) and that ν_i is a measure on (S, \mathcal{S}) for each i in a countable index set I . Suppose also that $\nu = \sum_{i \in I} \nu_i$ is a well-defined measure on (S, \mathcal{S}) .

1. If $\nu_i \ll \mu$ for every $i \in I$ then $\nu \ll \mu$.
2. If $\nu_i \perp \mu$ for every $i \in I$ then $\nu \perp \mu$.

Proof

Recall that if $A \in \mathcal{S}$ is null for ν_i for each $i \in I$, then A is null for $\nu = \sum_{i \in I} \nu_i$, assuming that this is a well-defined measure.

As before, note that $\nu = \sum_{i \in I} \nu_i$ is well-defined if ν_i is a positive measure for each $i \in I$ or if I is finite and ν_i is a finite measure for each $i \in I$. We close this subsection with a couple of results that involve both the absolute continuity relation and the singularity relation

Suppose that μ, ν , and ρ are measures on (S, \mathcal{S}) . If $\nu \ll \mu$ and $\mu \perp \rho$ then $\nu \perp \rho$.

Proof

Since $\mu \perp \rho$, there exists $A \in \mathcal{S}$ such that A is null for μ and A^c is null for ρ . But $\nu \ll \mu$ so A is null for ν . Hence $\nu \perp \rho$.

Suppose that μ and ν are measures on (S, \mathcal{S}) . If $\nu \ll \mu$ and $\nu \perp \mu$ then $\nu = \mathbf{0}$.

Proof

From the [previous theorem](#) (with $\rho = \nu$) we have $\nu \perp \nu$ and hence by (5), $\nu = \mathbf{0}$.

Density Functions

We are now ready for our study of density functions. Throughout this subsection, we assume that μ is a positive, σ -finite measure on our measurable space (S, \mathcal{S}) . Recall that if $f: S \rightarrow \mathbb{R}$ is measurable, then the integral of f with respect to μ may exist as a number in $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, \infty\}$ or may fail to exist.

Suppose that $f: S \rightarrow \mathbb{R}$ is a measurable function whose integral with respect to μ exists. Then function ν defined by

$$\nu(A) = \int_A f d\mu, \quad A \in \mathcal{S} \quad (3.13.1)$$

is a σ -finite measure on (S, \mathcal{S}) that is absolutely continuous with respect to μ . The function f is a *density function* of ν relative to μ .

Proof

To say that the integral exists means that either $\int_S f^+ d\mu < \infty$ or $\int_S f^- d\mu < \infty$, where as usual, f^+ and f^- are the positive and negative parts of f . So $\nu(A) = \nu_+(A) - \nu_-(A)$ for $A \in \mathcal{S}$ where $\nu_+(A) = \int_A f^+ d\mu$ and $\nu_-(A) = \int_A f^- d\mu$. Both ν_+ and ν_- are positive measures by basic properties of the integral: Generically, suppose $g: S \rightarrow [0, \infty)$ is measurable. The integral over the empty set is always 0, so $\int_\emptyset g d\mu = 0$. Next, if $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} and $A = \bigcup_{i \in I} A_i$, then by the additivity property of the integral over disjoint domains,

$$\int_A g d\mu = \sum_{i \in I} \int_{A_i} g d\mu \quad (3.13.2)$$

By the assumption that the integral exists, either ν_+ or ν_- is a finite positive measure, and hence ν is a measure. As you might guess, ν_+ and ν_- form the Jordan decomposition of ν , a point that we will revisit below.

Again, either ν_+ or ν_- is a finite measure. By symmetry, let's suppose that ν_- is finite. Then to show that ν is σ -finite, we just need to show that ν_+ is σ -finite. Since μ has this property, there exists a collection $\{A_n : n \in \mathbb{N}_+\}$ with $A_n \in \mathcal{S}$, $\mu(A_n) < \infty$, and $\bigcup_{n=1}^\infty A_n = S$. Let $B_n = \{x \in S : f^+(x) \leq n\}$ for $n \in \mathbb{N}_+$. Then $B_n \in \mathcal{S}$ for $n \in \mathbb{N}_+$ and $\bigcup_{n=1}^\infty B_n = S$. Hence $\{A_m \cap B_n : (m, n) \in \mathbb{N}_+^2\}$ is a countable collection of measurable sets whose union is also S . Moreover,

$$\nu_+(A_m \cap B_n) = \int_{A_m \cap B_n} f^+ d\mu \leq n\mu(A_m \cap B_n) < \infty \quad (3.13.3)$$

Finally, suppose $A \in \mathcal{S}$ is a null set of μ . If $B \in \mathcal{S}$ and $B \subseteq A$ then $\mu(B) = 0$ so $\nu(B) = \int_B f d\mu = 0$. Hence $\nu \ll \mu$.

The following three special cases are the most important:

1. If f is nonnegative (so that the integral exists in $\mathbb{R} \cup \{\infty\}$) then ν is a *positive measure* since $\nu(A) \geq 0$ for $A \in \mathcal{S}$.
2. If f is integrable (so that the integral exists in \mathbb{R}), then ν is a *finite measure* since $\nu(A) \in \mathbb{R}$ for $A \in \mathcal{S}$.
3. If f is nonnegative and $\int_S f d\mu = 1$ then ν is a *probability measure* since $\nu(A) \geq 0$ for $A \in \mathcal{S}$ and $\nu(S) = 1$.

In case 3, f is the *probability density function* of ν relative to μ , our favorite kind of density function. When they exist, density functions are essentially unique.

Suppose that ν is a σ -finite measure on (S, \mathcal{S}) and that ν has density function f with respect to μ . Then $g: S \rightarrow \mathbb{R}$ is a density function of ν with respect to μ if and only if $f = g$ almost everywhere on S with respect to μ .

Proof

These results also follow from basic properties of the integral. Suppose that $f, g: S \rightarrow \mathbb{R}$ are measurable functions whose integrals with respect to μ exist. If $g = f$ almost everywhere on S with respect to μ then $\int_A f d\mu = \int_A g d\mu$ for every $A \in \mathcal{S}$. Hence if f is a density function for ν with respect to μ then so is g . For the converse, if $\int_A f d\mu = \int_A g d\mu$ for every $A \in \mathcal{S}$, then since μ is σ -finite, it follows that $f = g$ almost everywhere on S with respect to μ .

The essential uniqueness of density functions can fail if the positive measure space (S, \mathcal{S}, μ) is not σ -finite. A simple example is given [below](#). Our next result answers the question of when a measure has a density function with respect to μ , and is the fundamental theorem of this section. The theorem is in two parts: Part (a) is the *Lebesgue decomposition theorem*, named for our old friend Henri Lebesgue. Part (b) is the *Radon-Nikodym theorem*, named for Johann Radon and Otto Nikodym. We combine the theorems because our proofs of the two results are inextricably linked.

Suppose that ν is a σ -finite measure on (S, \mathcal{S}) .

1. **Lebesgue Decomposition Theorem.** ν can be uniquely decomposed as $\nu = \nu_c + \nu_s$ where $\nu_c \ll \mu$ and $\nu_s \perp \mu$.
2. **Radon-Nikodym Theorem.** ν_c has a density function with respect to μ .

Proof

The proof proceeds in stages. we first prove the result for finite, positive measures, then for σ -finite, positive measures, and finally for general σ -finite measures. The first stage is the most complicated.

Part 1, suppose that μ and ν are positive, finite measures. Let \mathcal{F} denote the collection of measurable functions $g: S \rightarrow [0, \infty)$ with $\int_A g d\mu \leq \nu(A)$ for all $A \in \mathcal{S}$. Note that $\mathcal{F} \neq \emptyset$ since the constant function 0 is in \mathcal{F} . The proof works by finding a maximal element of \mathcal{F} and using this function as the density function of the absolutely continuous part of ν .

Our first step is to show that \mathcal{F} is closed under the max operator. Let $g_1, g_2 \in \mathcal{F}$. For $A \in \mathcal{S}$, let $A_1 = \{x \in A : g_1(x) \geq g_2(x)\}$ and $A_2 = \{x \in A : g_1(x) < g_2(x)\}$. Then $A_1, A_2 \in \mathcal{S}$ partition A so

$$\int_A \max\{g_1, g_2\} d\mu = \int_{A_1} \max\{g_1, g_2\} d\mu + \int_{A_2} \max\{g_1, g_2\} d\mu = \int_{A_1} g_1 d\mu + \int_{A_2} g_2 d\mu \leq \nu(A_1) + \nu(A_2) = \nu(A) \quad (3.13.4)$$

Hence $\max\{g_1, g_2\} \in \mathcal{F}$.

Our next step is to show that \mathcal{F} is closed with respect to increasing limits. Thus suppose that $g_n \in \mathcal{F}$ for $n \in \mathbb{N}_+$ and that g_n is increasing in n on S . Let $g = \lim_{n \rightarrow \infty} g_n$. Then $g: S \rightarrow [0, \infty]$ is measurable, and by the monotone convergence theorem, $\int_A g d\mu = \lim_{n \rightarrow \infty} \int_A g_n d\mu$ for every $A \in \mathcal{S}$. But $\int_A g_n d\mu \leq \nu(A)$ for every $n \in \mathbb{N}_+$ so $\int_A g d\mu \leq \nu(A)$. In particular, $\int_S g d\mu \leq \nu(S) < \infty$ so $g < \infty$ almost everywhere on S with respect to μ . Thus, by redefining g on a μ -null set if necessary, we can assume $g < \infty$ on S . Hence $g \in \mathcal{F}$.

Now let $\alpha = \sup \left\{ \int_S g d\mu : g \in \mathcal{F} \right\}$. Note that $\alpha \leq \nu(S) < \infty$. By definition of the supremum, for each $n \in \mathbb{N}_+$ there exist $g_n \in \mathcal{F}$ such that $\int_S g_n d\mu > \alpha - \frac{1}{n}$. Now let $f_n = \max\{g_1, g_2, \dots, g_n\}$ for $n \in \mathbb{N}_+$. Then $f_n \in \mathcal{F}$ and f_n is increasing in $n \in \mathbb{N}_+$ on S . Hence $f = \lim_{n \rightarrow \infty} f_n \in \mathcal{F}$ and $\int_S f d\mu = \lim_{n \rightarrow \infty} \int_S f_n d\mu$. But $\int_S f_n d\mu \geq \int_S g_n d\mu > \alpha - \frac{1}{n}$ for each $n \in \mathbb{N}_+$ and hence $\int_S f d\mu \geq \alpha$.

Define $\nu_c(A) = \int_A f d\mu$ and $\nu_s(A) = \nu(A) - \nu_c(A)$ for $A \in \mathcal{S}$. Then ν_c and ν_s are finite, positive measures and by our previous theorem, ν_c is absolutely continuous with respect to μ and has density function f . Our next step is to show that ν_s is singular with respect to μ . For $n \in \mathbb{N}$, let (P_n, P_n^c) denote a Hahn decomposition of the measure $\nu_s - \frac{1}{n}\mu$. Then

$$\int_A \left(f + \frac{1}{n} \mathbf{1}_{P_n} \right) d\mu = \nu_c(A) + \frac{1}{n} \mu(P_n \cap A) = \nu(A) - \left[\nu_s(A) - \frac{1}{n} \mu(P_n \cap A) \right] \quad (3.13.5)$$

But $\nu_s(A) - \frac{1}{n}\mu(P_n \cap A) \geq \nu_s(A \cap P_n) - \frac{1}{n}\mu(A \cap P_n) \geq 0$ since ν_s is a positive measure and P_n is positive for $\nu_s - \frac{1}{n}\mu$. Thus we have $\int_A (f + \frac{1}{n}\mathbf{1}_{P_n}) d\mu \leq \nu(A)$ for every $A \in \mathcal{S}$, so $f + \frac{1}{n}\mathbf{1}_{P_n} \in \mathcal{F}$ for every $n \in \mathbb{N}_+$. If $\mu(P_n) > 0$ then $\int_S (f + \frac{1}{n}\mathbf{1}_{P_n}) d\mu = \alpha + \frac{1}{n}\mu(P_n) > \alpha$, which contradicts the definition of α . Hence we must have $\mu(P_n) = 0$ for every $n \in \mathbb{N}_+$. Now let $P = \bigcup_{n=1}^{\infty} P_n$. Then $\mu(P) = 0$. If $\nu_s(P^c) > 0$ then $\nu_s(P^c) - \frac{1}{n}\mu(P^c) > 0$ for n sufficiently large. But this is a contradiction since $P^c \subseteq P_n^c$ which is negative for $\nu_s - \frac{1}{n}\mu$ for every $n \in \mathbb{N}_+$. Thus we must have $\nu_s(P^c) = 0$, so μ and ν_s are singular.

Part 2. Suppose that μ and ν are σ -finite, positive measures. Then there exists a countable partition $\{S_i : i \in I\}$ of S where $S_i \in \mathcal{S}$ for $i \in I$, and $\mu(S_i) < \infty$ and $\nu(S_i) < \infty$ for $i \in I$. Let $\mu_i(A) = \mu(A \cap S_i)$ and $\nu_i(A) = \nu(A \cap S_i)$ for $i \in I$. Then μ_i and ν_i are finite, positive measures for $i \in I$, and $\mu = \sum_{i \in I} \mu_i$ and $\nu = \sum_{i \in I} \nu_i$. By part 1, for each $i \in I$, there exists a measurable function $f_i : S \rightarrow [0, \infty)$ such that $\nu_i = \nu_{i,c} + \nu_{i,s}$ where $\nu_{i,c}(A) = \int_A f_i d\mu$ for $A \in \mathcal{S}$ and $\nu_{i,s} \perp \mu$. Let $f = \sum_{i \in I} \mathbf{1}_{A_i} f_i$. Then $f : S \rightarrow [0, \infty)$ is measurable. Define $\nu_c(A) = \int_A f d\mu$ and $\nu_s(A) = \nu(A) - \nu_c(A)$ for $A \in \mathcal{S}$. Note that $\nu_c = \sum_{i \in I} \nu_{i,c}$ and $\nu_s = \sum_{i \in I} \nu_{i,s}$. Then $\nu_c \ll \mu$ and has density function f and $\nu_s \perp \mu$.

Part 3. Suppose that ν is a σ -finite measure (not necessarily positive). By the Jordan decomposition theorem, $\nu = \nu_+ - \nu_-$ where ν_+ and ν_- are σ -finite, positive measures, and at least one is finite. By part 2, there exist measurable functions $f_+ : S \rightarrow [0, \infty)$ and $f_- : S \rightarrow [0, \infty)$ such that $\nu_+ = \nu_{+,c} + \nu_{+,s}$ and $\nu_- = \nu_{-,c} + \nu_{-,s}$ where $\nu_{+,c}(A) = \int_A f_+ d\mu$, $\nu_{-,c}(A) = \int_A f_- d\mu$ for $A \in \mathcal{S}$, and $\nu_{+,s} \perp \mu$, $\nu_{-,s} \perp \mu$. Let $f = f_+ - f_-$, $\nu_c(A) = \int_A f d\mu$, $\nu_s(A) = \nu(A) - \nu_c(A)$ for $A \in \mathcal{S}$. Then $\nu = \nu_c + \nu_s$ and $\nu_s = \nu_{+,s} - \nu_{-,s} \perp \mu$.

Uniqueness. Suppose that $\nu = \nu_{c,1} + \nu_{s,1} = \nu_{c,2} + \nu_{s,2}$ where $\nu_{c,i} \ll \mu$ and $\nu_{s,i} \perp \mu$ for $i \in \{1, 2\}$. Then $\nu_{c,1} - \nu_{c,2} = \nu_{s,2} - \nu_{s,1}$. But $\nu_{c,1} - \nu_{c,2} \ll \mu$ and $\nu_{s,2} - \nu_{s,1} \perp \mu$ so $\nu_{c,1} - \nu_{c,2} = \nu_{s,2} - \nu_{s,1} = \mathbf{0}$ by the theorem above.

In particular, a measure ν on (S, \mathcal{S}) has a density function with respect to μ if and only if $\nu \ll \mu$. The density function in this case is also referred to as the *Radon-Nikodym derivative* of ν with respect to μ and is sometimes written in derivative notation as $d\nu/d\mu$. This notation, however, can be a bit misleading because we need to remember that a density function is unique only up to a μ -null set. Also, the Radon-Nikodym theorem can fail if the positive measure space (S, \mathcal{S}, μ) is not σ -finite. A couple of examples are given below. Next we characterize the Hahn decomposition and the Jordan decomposition of ν in terms of the density function.

Suppose that ν is a measure on (S, \mathcal{S}) with $\nu \ll \mu$, and that ν has density function f with respect to μ . Let $P = \{x \in S : f(x) \geq 0\}$, and let f^+ and f^- denote the positive and negative parts of f .

1. A Hahn decomposition of ν is (P, P^c) .
2. The Jordan decomposition is $\nu = \nu_+ - \nu_-$ where $\nu_+(A) = \int_A f^+ d\mu$ and $\nu_-(A) = \int_A f^- d\mu$, for $A \in \mathcal{S}$.

Proof

Of course $P^c = \{x \in S : f(x) < 0\}$. The proofs are simple.

1. Suppose that $A \in \mathcal{S}$. If $A \subseteq P$ then $f(x) \geq 0$ for $x \in A$ and hence $\nu(A) = \int_A f d\mu \geq 0$. If $A \subseteq P^c$ then $\nu(A) = \int_A f d\mu \leq 0$.
2. This follows immediately from (a) and the Jordan decomposition theorem, since $\nu_+(A) = \nu(A \cap P)$ and $\nu_-(A) = -\nu(A \cap P^c)$ for $A \in \mathcal{S}$. Note that $f^+ = \mathbf{1}_P f$ and $f^- = -\mathbf{1}_{P^c} f$.

The following result is a basic *change of variables theorem* for integrals.

Suppose that ν is a positive measure on (S, \mathcal{S}) with $\nu \ll \mu$ and that ν has density function f with respect to μ . If $g : S \rightarrow \mathbb{R}$ is a measurable function whose integral with respect to ν exists, then

$$\int_S g d\nu = \int_S g f d\mu \quad (3.13.6)$$

Proof

The proof is a classical *bootstrapping argument*. Suppose first that $g = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ is a nonnegative simple function. That is, I is a finite index set, $a_i \in [0, \infty)$ for $i \in I$, and $\{A_i : i \in I\}$ is a disjoint collection of sets in \mathcal{S} . Then $\int_S g d\nu = \sum_{i \in I} a_i \nu(A_i)$. But $\nu(A_i) = \int_{A_i} f d\mu = \int_S \mathbf{1}_{A_i} f d\mu$ for each $i \in I$ so

$$\int_S g d\mu = \sum_{i \in I} a_i \int_S \mathbf{1}_{A_i} f d\mu = \int_S \left(\sum_{i \in I} a_i \mathbf{1}_{A_i} \right) f d\mu = \int_S g f d\mu \quad (3.13.7)$$

Suppose next that $g : S \rightarrow [0, \infty)$ is measurable. There exists a sequence of nonnegative simple functions (g_1, g_2, \dots) such that g_n is increasing in $n \in \mathbb{N}_+$ on S and $g_n \rightarrow g$ as $n \rightarrow \infty$ on S . Since f is nonnegative, $g_n f$ is increasing in $n \in \mathbb{N}_+$ on S and $g_n f \rightarrow g f$ as $n \rightarrow \infty$ on S . By the first step, $\int_S g_n d\nu = \int_S g_n f d\mu$ for each $n \in \mathbb{N}_+$. But by the monotone convergence theorem, $\int_S g_n d\nu \rightarrow \int_S g d\nu$ and $\int_S g_n f d\mu \rightarrow \int_S g f d\mu$ as $n \rightarrow \infty$. Hence $\int_S g d\nu = \int_S g f d\mu$.

Finally, suppose that $g: S \rightarrow \mathbb{R}$ is a measurable function whose integral with respect to ν exists. By the previous step, $\int_S g^+ d\nu = \int_S g^+ f d\mu$ and $\int_S g^- d\nu = \int_S g^- f d\mu$, and at least one of these integrals is finite. Hence by the additive property

$$\int_S g d\nu = \int_S g^+ d\nu - \int_S g^- d\nu = \int_S g^+ f d\mu - \int_S g^- f d\mu = \int_S (g^+ - g^-) f d\mu = \int_S g f d\mu \quad (3.13.8)$$

In *differential notation*, the change of variables theorem has the familiar form $d\nu = f d\mu$, and this is really the justification for the derivative notation $f = d\nu/d\mu$ in the first place. The following result gives the *scalar multiple rule* for density functions.

Suppose that ν is a measure on (S, \mathcal{S}) with $\nu \ll \mu$ and that ν has density function f with respect to μ . If $c \in \mathbb{R}$, then $c\nu$ has density function cf with respect to μ .

Proof

If $A \in \mathcal{S}$ then $\int_A cf d\mu = c \int_A f d\mu = c\nu(A)$.

Of course, we already knew that $\nu \ll \mu$ implies $c\nu \ll \mu$ for $c \in \mathbb{R}$, so the new information is the relation between the density functions. In derivative notation, the scalar multiple rule has the familiar form

$$\frac{d(c\nu)}{d\mu} = c \frac{d\nu}{d\mu} \quad (3.13.9)$$

The following result gives the sum rule for density functions. Recall that two measures are of the *same type* if neither takes the value ∞ or if neither takes the value $-\infty$.

Suppose that ν and ρ are measures on (S, \mathcal{S}) of the same type with $\nu \ll \mu$ and $\rho \ll \mu$, and that ν and ρ have density functions f and g with respect to μ , respectively. Then $\nu + \rho$ has density function $f + g$ with respect to μ .

Proof

If $A \in \mathcal{S}$ then

$$\int_A (f + g) d\mu = \int_A f d\mu + \int_A g d\mu = \nu(A) + \rho(A) \quad (3.13.10)$$

The additive property holds because we know that the integrals in the middle of the displayed equation are not of the form $\infty - \infty$.

Of course, we already knew that $\nu \ll \mu$ and $\rho \ll \mu$ imply $\nu + \rho \ll \mu$, so the new information is the relation between the density functions. In derivative notation, the sum rule has the familiar form

$$\frac{d(\nu + \rho)}{d\mu} = \frac{d\nu}{d\mu} + \frac{d\rho}{d\mu} \quad (3.13.11)$$

The following result is the *chain rule* for density functions.

Suppose that ν is a positive measure on (S, \mathcal{S}) with $\nu \ll \mu$ and that ν has density function f with respect to μ . Suppose ρ is a measure on (S, \mathcal{S}) with $\rho \ll \nu$ and that ρ has density function g with respect to ν . Then ρ has density function gf with respect to μ .

Proof

This is a simple consequence of the change of variables theorem above. If $A \in \mathcal{S}$ then $\rho(A) = \int_A g d\nu = \int_A gf d\mu$.

Of course, we already knew that $\nu \ll \mu$ and $\rho \ll \nu$ imply $\rho \ll \mu$, so once again the new information is the relation between the density functions. In derivative notation, the chain rule has the familiar form

$$\frac{d\rho}{d\mu} = \frac{d\rho}{d\nu} \frac{d\nu}{d\mu} \quad (3.13.12)$$

The following related result is the *inverse rule* for density functions.

Suppose that ν is a positive measure on (S, \mathcal{S}) with $\nu \ll \mu$ and $\mu \ll \nu$ (so that $\nu \equiv \mu$). If ν has density function f with respect to μ then μ has density function $1/f$ with respect to ν .

Proof

Let f be a density function of ν with respect to μ and let $Z = \{x \in S : f(x) = 0\}$. Then $\nu(Z) = \int_Z f d\mu = 0$ so Z is a null set of ν and hence is also a null set of μ . Thus, we can assume that $f \neq 0$ on S . Let g be a density of μ with respect to ν . Since $\mu \ll \nu \ll \mu$, it follows from the chain rule that fg is a density of μ with respect to μ . But of course the constant function 1 is also a density of μ with respect to itself so we have $fg = 1$ almost everywhere on S . Thus $1/f$ is a density of μ with respect to ν .

In derivative notation, the inverse rule has the familiar form

$$\frac{d\mu}{d\nu} = \frac{1}{d\nu/d\mu} \quad (3.13.13)$$

Examples and Special Cases

Discrete Spaces

Recall that a *discrete measure space* $(S, \mathcal{S}, \#)$ consists of a countable set S with the σ -algebra $\mathcal{S} = \mathcal{P}(S)$ of all subsets of S , and with counting measure $\#$. Of course $\#$ is a positive measure and is trivially σ -finite since S is countable. Note also that \emptyset is the only set that is null for $\#$. If ν is a measure on S , then by definition, $\nu(\emptyset) = 0$, so ν is absolutely continuous relative to μ . Thus, by the Radon-Nikodym theorem, ν can be written in the form

$$\nu(A) = \sum_{x \in A} f(x), \quad A \subseteq S \quad (3.13.14)$$

for a unique $f : S \rightarrow \mathbb{R}$. Of course, this is obvious by a direct argument. If we define $f(x) = \nu\{x\}$ for $x \in S$ then the displayed equation follows by the countable additivity of ν .

Spaces Generated by Countable Partitions

We can generalize the last discussion to *spaces generated by countable partitions*. Suppose that S is a set and that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of S into nonempty sets. Let $\mathcal{S} = \sigma(\mathcal{A})$ and recall that every $A \in \mathcal{S}$ has a unique representation of the form $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$. Suppose now that μ is a positive measure on \mathcal{S} with $0 < \mu(A_i) < \infty$ for every $i \in I$. Then once again, the measure space (S, \mathcal{S}, μ) is σ -finite and \emptyset is the only null set. Hence if ν is a measure on (S, \mathcal{S}) then ν is absolutely continuous with respect to μ and hence has unique density function f with respect to μ :

$$\nu(A) = \int_A f d\mu, \quad A \in \mathcal{S} \quad (3.13.15)$$

Once again, we can construct the density function explicitly.

In the setting above, define $f : S \rightarrow \mathbb{R}$ by $f(x) = \nu(A_i)/\mu(A_i)$ for $x \in A_i$ and $i \in I$. Then f is the density of ν with respect to μ .

Proof

Suppose that $A \in \mathcal{S}$ so that $A = \bigcup_{j \in J} A_j$ for some $J \subseteq I$. Then

$$\int_A f d\mu = \sum_{j \in J} \int_{A_j} f d\mu = \sum_{j \in J} \frac{\nu(A_j)}{\mu(A_j)} \mu(A_j) = \sum_{j \in J} \nu(A_j) = \nu(A) \quad (3.13.16)$$

Often positive measure spaces that occur in applications can be decomposed into spaces generated by countable partitions. In the section on Convergence in the chapter on Martingales, we show that more general density functions can be obtained as limits of density functions of the type in the [last theorem](#).

Probability Spaces

Suppose that $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space and that X is a random variable taking values in a measurable space (S, \mathcal{S}) . Recall that the *distribution* of X is the probability measure P_X on (S, \mathcal{S}) given by

$$P_X(A) = \mathbb{P}(X \in A), \quad A \in \mathcal{S} \quad (3.13.17)$$

If μ is a positive measure, σ -finite measure on (S, \mathcal{S}) , then the theory of this section applies, of course. The Radon-Nikodym theorem tells us precisely when (the distribution of) X has a probability density function with respect to μ : we need the distribution to be absolutely continuous with respect to μ : if $\mu(A) = 0$ then $P_X(A) = \mathbb{P}(X \in A) = 0$ for $A \in \mathcal{S}$.

Suppose that $r : S \rightarrow \mathbb{R}$ is measurable, so that $r(X)$ is a real-valued random variable. The integral of $r(X)$ (assuming that it exists) is of fundamental importance, and is known as the *expected value* of $r(X)$. We will study expected values in detail in the next chapter, but here we just note different ways to write the integral. By the change of variables theorem in the last section we have

$$\int_{\Omega} r[X(\omega)] d\mathbb{P}(\omega) = \int_S r(x) dP_X(x) \quad (3.13.18)$$

Assuming that P_X , the distribution of X , is absolutely continuous with respect to μ , with density function f , we can add to our chain of integrals using Theorem [\(14\)](#):

$$\int_{\Omega} r[X(\omega)] d\mathbb{P}(\omega) = \int_S r(x) dP_X(x) = \int_S r(x) f(x) d\mu(x) \quad (3.13.19)$$

Specializing, suppose that $(S, \mathcal{S}, \#)$ is a discrete measure space. Thus X has a discrete distribution and (as noted in the previous subsection), the distribution of X is absolutely continuous with respect to $\#$, with probability density function f given by $f(x) = \mathbb{P}(X = x)$ for $x \in S$. In this case the integral simplifies:

$$\int_{\Omega} r[X(\omega)] d\mathbb{P}(\omega) = \sum_{x \in S} r(x) f(x) \quad (3.13.20)$$

Recall next that for $n \in \mathbb{N}_+$, the n -dimensional Euclidean measure space is $(\mathbb{R}^n, \mathcal{R}_n, \lambda_n)$ where \mathcal{R}_n is the σ -algebra of Lebesgue measurable sets and λ_n is Lebesgue measure. Suppose now that $S \in \mathcal{R}_n$ and that \mathcal{S} is the σ -algebra of Lebesgue measurable subsets of S , and that once again, X is a random variable with values in S . By definition, X has a continuous distribution if $\mathbb{P}(X = x) = 0$ for $x \in S$. But we now know that this is not enough to ensure that the distribution of X has a density function with respect to λ_n . We need the distribution to be *absolutely* continuous, so that if $\lambda_n(A) = 0$ then $\mathbb{P}(X \in A) = 0$ for $A \in \mathcal{S}$. Of course $\lambda_n\{x\} = 0$ for $x \in S$, so absolute continuity implies continuity, but not conversely. Continuity of the distribution is a (much) weaker condition than absolute continuity of the distribution. If the distribution of X is continuous but not absolutely so, then the distribution will not have a density function with respect to λ_n .

For example, suppose that $\lambda_n(S) = 0$. Then the distribution of X and λ_n are mutually singular since $\mathbb{P}(X \in S) = 1$ and so X will not have a density function with respect to λ_n . This will always be the case if S is countable, so that the distribution of X is discrete. But it is also possible for X to have a continuous distribution on an uncountable set $S \in \mathcal{R}_n$ with $\lambda_n(S) = 0$. In such a case, the continuous distribution of X is said to be *degenerate*. There are a couple of natural ways in which this can happen that are illustrated in the following exercises.

Suppose that Θ is uniformly distributed on the interval $[0, 2\pi)$. Let $X = \cos \Theta$, $Y = \sin \Theta$.

1. (X, Y) has a continuous distribution on the circle $C = \{(x, y) : x^2 + y^2 = 1\}$.
2. The distribution of (X, Y) and λ_2 are mutually singular.
3. Find $\mathbb{P}(Y > X)$.

Solution

1. If $(x, y) \in C$ then there exist a unique $\theta \in [0, 2\pi)$ with $x = \cos \theta$ and $y = \sin \theta$. Hence $\mathbb{P}[(X, Y) = (x, y)] = \mathbb{P}(\Theta = \theta) = 0$.
2. $\mathbb{P}[(X, Y) \in C] = 1$ but $\lambda_2(C) = 0$.
3. $\frac{1}{2}$

The last example is artificial since (X, Y) has a one-dimensional distribution in a sense, in spite of taking values in \mathbb{R}^2 . And of course Θ has a probability density function f with respect to λ_1 given by $f(\theta) = 1/2\pi$ for $\theta \in [0, 2\pi)$.

Suppose that X is uniformly distributed on the set $\{0, 1, 2\}$, Y is uniformly distributed on the interval $[0, 2]$, and that X and Y are independent.

1. (X, Y) has a continuous distribution on the product set $S = \{0, 1, 2\} \times [0, 2]$.
2. The distribution of (X, Y) and λ_2 are mutually singular.
3. Find $\mathbb{P}(Y > X)$.

Solution

1. The variables are independent and Y has a continuous distribution so $\mathbb{P}[(X, Y) = (x, y)] = \mathbb{P}(X = 2)\mathbb{P}(Y = y) = 0$ for $(x, y) \in S$.
2. $\mathbb{P}[(X, Y) \in S] = 1$ but $\lambda_2(S) = 0$
3. $\frac{1}{2}$

The last exercise is artificial since X has a discrete distribution on $\{0, 1, 2\}$ (with all subsets measurable and with $\#$), and Y a continuous distribution on the Euclidean space $[0, 2]$ (with Lebesgue measurable subsets and with λ). Both are absolutely continuous; X has density function g given by $g(x) = 1/3$ for $x \in \{0, 1, 2\}$ and Y has density function h given by $h(y) = 1/2$ for $y \in [0, 2]$. So really, the proper measure space on S is the product measure space formed from these two spaces. Relative to this product space (X, Y) has a density f given by $f(x, y) = 1/6$ for $(x, y) \in S$.

It is also possible to have a continuous distribution on $S \subseteq \mathbb{R}^n$ with $\lambda_n(S) > 0$, yet still with no probability density function, a much more interesting situation. We will give a classical construction. Let (X_1, X_2, \dots) be a sequence of Bernoulli trials with success parameter $p \in (0, 1)$. We will indicate the dependence of the probability measure \mathbb{P} on the parameter p with a subscript. Thus, we have a sequence of independent indicator variables with

$$\mathbb{P}_p(X_i = 1) = p, \quad \mathbb{P}_p(X_i = 0) = 1 - p \quad (3.13.21)$$

We interpret X_i as the i th binary digit (*bit*) of a random variable X taking values in $(0, 1)$. That is, $X = \sum_{i=1}^{\infty} X_i/2^i$. Conversely, recall that every number $x \in (0, 1)$ can be written in binary form as $x = \sum_{i=1}^{\infty} x_i/2^i$ where $x_i \in \{0, 1\}$ for each $i \in \mathbb{N}_+$. This representation is unique

except when x is a *binary rational* of the form $x = k/2^n$ for $n \in \mathbb{N}_+$ and $k \in \{1, 3, \dots, 2^n - 1\}$. In this case, there are two representations, one in which the bits are eventually 0 and one in which the bits are eventually 1. Note, however, that the set of binary rationals is countable. Finally, note that the uniform distribution on $(0, 1)$ is the same as Lebesgue measure on $(0, 1)$.

X has a continuous distribution on $(0, 1)$ for every value of the parameter $p \in (0, 1)$. Moreover,

1. If $p, q \in (0, 1)$ and $p \neq q$ then the distribution of X with parameter p and the distribution of X with parameter q are mutually singular.
2. If $p = \frac{1}{2}$, X has the uniform distribution on $(0, 1)$.
3. If $p \neq \frac{1}{2}$, then the distribution of X is singular with respect to Lebesgue measure on $(0, 1)$, and hence has no probability density function in the usual sense.

Proof

If $x \in (0, 1)$ is not a binary rational, then

$$\mathbb{P}_p(X = x) = \mathbb{P}_p(X_i = x_i \text{ for all } i \in \mathbb{N}_+) = \lim_{n \rightarrow \infty} \mathbb{P}_p(X_i = x_i \text{ for } i = 1, 2, \dots, n) = \lim_{n \rightarrow \infty} p^y (1-p)^{n-y} \quad (3.13.22)$$

where $y = \sum_{i=1}^n x_i$. Let $q = \max\{p, 1-p\}$. Then $p^y (1-p)^{n-y} \leq q^n \rightarrow 0$ as $n \rightarrow \infty$. Hence, $\mathbb{P}_p(X = x) = 0$. If $x \in (0, 1)$ is a binary rational, then there are two bit strings that represent x , say (x_1, x_2, \dots) (with bits eventually 0) and (y_1, y_2, \dots) (with bits eventually 1). Hence $\mathbb{P}_p(X = x) = \mathbb{P}_p(X_i = x_i \text{ for all } i \in \mathbb{N}_+) + \mathbb{P}_p(X_i = y_i \text{ for all } i \in \mathbb{N}_+)$. But both of these probabilities are 0 by the same argument as before.

Next, we define the set of numbers for which the limiting relative frequency of 1's is p . Let $C_p = \{x \in (0, 1) : \frac{1}{n} \sum_{i=1}^n x_i \rightarrow p \text{ as } n \rightarrow \infty\}$. Note that since limits are unique, $C_p \cap C_q = \emptyset$ for $p \neq q$. Next, by the strong law of large numbers, $\mathbb{P}_p(X \in C_p) = 1$. Although we have not yet studied the law of large numbers, the basic idea is simple: in a sequence of Bernoulli trials with success probability p , the long-term relative frequency of successes is p . Thus the distributions of X , as p varies from 0 to 1, are mutually singular; that is, as p varies, X takes values with probability 1 in mutually disjoint sets.

Let F denote the distribution function of X , so that $F(x) = \mathbb{P}_p(X \leq x) = \mathbb{P}_p(X < x)$ for $x \in (0, 1)$. If $x \in (0, 1)$ is not a binary rational, then $X < x$ if and only if there exists $n \in \mathbb{N}_+$ such that $X_i = x_i$ for $i \in \{1, 2, \dots, n-1\}$ and $X_n = 0$ while $x_n = 1$. Hence $\mathbb{P}_{1/2}(X < x) = \sum_{n=1}^{\infty} \frac{x_n}{2^n} = x$. Since the distribution function of a continuous distribution is continuous, it follows that $F(x) = x$ for all $x \in [0, 1]$. This means that X has the uniform distribution on $(0, 1)$. If $p \neq \frac{1}{2}$, the distribution of X and the uniform distribution are mutually singular, so in particular, X does not have a probability density function with respect to Lebesgue measure.

For an application of some of the ideas in this example, see Bold Play in the game of Red and Black.

Counterexamples

The essential uniqueness of density functions can fail if the underlying positive measure μ is not σ -finite. Here is a trivial counterexample:

Suppose that S is a nonempty set and that $\mathcal{S} = \{S, \emptyset\}$ is the trivial σ -algebra. Define the positive measure μ on (S, \mathcal{S}) by $\mu(\emptyset) = 0$, $\mu(S) = \infty$. Let ν_c denote the measure on (S, \mathcal{S}) with constant density function $c \in \mathbb{R}$ with respect to μ .

1. (S, \mathcal{S}, μ) is not σ -finite.
2. $\nu_c = \mu$ for every $c \in (0, \infty)$.

The Radon-Nikodym theorem can fail if the measure μ is not σ -finite, even if ν is finite. Here are a couple of standard counterexamples:

Suppose that S is an uncountable set and \mathcal{S} is the σ -algebra of *countable and co-countable sets*:

$$\mathcal{S} = \{A \subseteq S : A \text{ is countable or } A^c \text{ is countable}\} \quad (3.13.23)$$

As usual, let $\#$ denote counting measure on \mathcal{S} , and define ν on \mathcal{S} by $\nu(A) = 0$ if A is countable and $\nu(A) = 1$ if A^c is countable. Then

1. $(S, \mathcal{S}, \#)$ is not σ -finite.
2. ν is a finite, positive measure on (S, \mathcal{S}) .
3. ν is absolutely continuous with respect to $\#$.
4. ν does not have a density function with respect to $\#$.

Proof

1. Recall that a countable union of countable sets is countable, and so S cannot be written as such a union.
2. Note that $\nu(\emptyset) = 0$. Suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} . If A_i is countable for every $i \in I$ then $\bigcup_{i \in I} A_i$ is countable. Hence $\nu(\bigcup_{i \in I} A_i) = 0$ and $\nu(A_i) = 0$ for every $i \in I$. Next suppose that A_j^c and A_k^c are countable for distinct $j, k \in I$. Since $A_j \cap A_k = \emptyset$, we have $A_j^c \cup A_k^c = S$. But then S would be countable, which is a contradiction. Hence it is only possible for to have A_j^c countable for a single $j \in I$. In this case, $\nu(A_j) = 1$ and $\nu(A_i) = 0$ for $i \neq j$. But also $(\bigcup_{i \in I} A_i)^c = \bigcap_{i \in I} A_i^c$

is countable, so $\nu(\bigcup_{i \in I} A_i) = 1$. Hence in all cases, $\nu(\bigcup_{i \in I} A_i) = \sum_{i \in I} \nu(A_i)$ so ν is a measure on (S, \mathcal{S}) . It is clearly positive and finite.

3. Recall that any measure is absolutely continuous with respect to counting measure, since $\#(A) = 0$ if and only if $A = \emptyset$.
4. Suppose that ν has density function f with respect to $\#$. Then $0 = \nu\{x\} = \int_{\{x\}} f d\# = f(x)$ for every $x \in S$. But then $\nu(S) = \int_S f d\# = 0$, which is a contradiction.

Let \mathcal{R} denote the standard Borel σ -algebra on \mathbb{R} . Let $\#$ and λ denote counting measure and Lebesgue measure on $(\mathbb{R}, \mathcal{R})$, respectively. Then

1. $(\mathbb{R}, \mathcal{R}, \#)$ is not σ -finite.
2. λ is absolutely continuous with respect to $\#$.
3. λ does not have a density function with respect to $\#$.

Proof

1. \mathbb{R} is uncountable and hence cannot be written as a countable union of finite sets.
2. Since \emptyset is the only null set of $\#$, $\lambda \ll \#$.
3. Suppose that λ has density function f with respect to $\#$. Then

$$0 = \lambda\{x\} = \int_{\{x\}} f d\# = f(x), \quad x \in \mathbb{R} \quad (3.13.24)$$

But then also $\lambda(\mathbb{R}) = \int_{\mathbb{R}} f d\# = 0$, a contradiction.

This page titled [3.13: Absolute Continuity and Density Functions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

3.14: Function Spaces

Basic Theory

Our starting point is a positive measure space (S, \mathcal{S}, μ) . That is S is a set, \mathcal{S} is a σ -algebra of subsets of S , and μ is a positive measure on (S, \mathcal{S}) . As usual, the most important special cases are

- *Euclidean space*: S is a Lebesgue measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$, \mathcal{S} is the σ -algebra of Lebesgue measurable subsets of S , and $\mu = \lambda_n$ is n -dimensional Lebesgue measure.
- *Discrete space*: S is a countable set, $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S , and $\mu = \#$ is counting measure.
- *Probability space*: S is the set of outcomes of a random experiment, \mathcal{S} is the σ -algebra of events, and $\mu = \mathbb{P}$ is a probability measure.

In previous sections, we defined the integral of certain measurable functions $f: S \rightarrow \mathbb{R}$ with respect to μ , and we studied properties of the integral. In this section, we will study vector spaces of functions that are defined in terms of certain integrability conditions. These function spaces are of fundamental importance in all areas of analysis, including probability. In particular, the results of this section will reappear in the form of spaces of random variables in our study of expected value.

Definitions and Basic Properties

Consider a statement on the elements of S , for example an equation or an inequality with $x \in S$ as a free variable. (Technically such a statement is a *predicate* on S .) For $A \in \mathcal{S}$, we say that the statement holds *on* A if it is true for every $x \in A$. We say that the statement holds *almost everywhere on* A (with respect to μ) if there exists $B \in \mathcal{S}$ with $B \subseteq A$ such that the statement holds on B and $\mu(A \setminus B) = 0$.

Measurable functions $f, g: S \rightarrow \mathbb{R}$ are *equivalent* if $f = g$ almost everywhere on S , in which case we write $f \equiv g$. The relation \equiv is an equivalence relation on the collection of measurable functions from S to \mathbb{R} . That is, if $f, g, h: S \rightarrow \mathbb{R}$ are measurable then

1. $f \equiv f$, the *reflexive property*.
2. If $f \equiv g$ then $g \equiv f$, the *symmetric property*.
3. If $f \equiv g$ and $g \equiv h$ then $f \equiv h$, the *transitive property*.

Thus, equivalent functions are indistinguishable from the point of view of the measure μ . As with any equivalence relation, \equiv partitions the underlying set (in this case the collection of real-valued measurable functions on S) into equivalence classes of mutually equivalent elements. As we will see, we often view these equivalence classes as the basic objects of study. Our next task is to define measures of the “size” of a function; these will become norms in our spaces.

Suppose that $f: S \rightarrow \mathbb{R}$ is measurable. For $p \in (0, \infty)$ we define

$$\|f\|_p = \left(\int_S |f|^p d\mu \right)^{1/p} \quad (3.14.1)$$

We also define $\|f\|_\infty = \inf \{b \in [0, \infty] : |f| \leq b \text{ almost everywhere on } S\}$.

Since $|f|^p$ is a nonnegative, measurable function for $p \in (0, \infty)$, $\int_S |f|^p d\mu$ exists in $[0, \infty]$, and hence so does $\|f\|_p$. Clearly $\|f\|_\infty$ also exists in $[0, \infty]$ and is known as the *essential supremum* of f . A number $b \in [0, \infty]$ such that $|f| \leq b$ almost everywhere on S is an *essential bound* of f and so, appropriately enough, the essential supremum of f is the infimum of the essential bounds of f . Thus, we have defined $\|f\|_p$ for all $p \in (0, \infty]$. The definition for $p = \infty$ is special, but we will see that it's the appropriate one.

For $p \in (0, \infty]$, let L^p denote the collection of measurable functions $f: S \rightarrow \mathbb{R}$ such that $\|f\|_p < \infty$.

So for $p \in (0, \infty)$, $f \in L^p$ if and only if $|f|^p$ is integrable. The symbol L is in honor of Henri Lebesgue, who first developed the theory. If we want to indicate the dependence on the underlying measure space, we write $L^p(S, \mathcal{S}, \mu)$. Of course, L^1 is simply the

collection of functions that are integrable with respect to μ . Our goal is to study the spaces L^p for $p \in (0, \infty]$. We start with some simple properties.

Suppose that $f : S \rightarrow \mathbb{R}$ is measurable. Then for $p \in (0, \infty]$,

1. $\|f\|_p \geq 0$
2. $\|f\|_p = 0$ if and only if $f = 0$ almost everywhere on S , so that $f \equiv 0$.

Proof

1. This is obvious from the definitions.
2. For $p \in (0, \infty)$, this follows from properties of the integral that we already have. First of course, $\int_S 0^p d\mu = \int_S 0 d\mu = 0$ so $\|0\|_p = 0$. Conversely if $\|f\|_p = 0$ then $\int_S |f|^p d\mu = 0$ and hence $|f|^p = 0$ almost everywhere on S and so $f = 0$ almost everywhere on S . Suppose $p = \infty$. Clearly $\|0\|_\infty = 0$. Conversely suppose that $\|f\|_\infty = 0$. Then for each $n \in \mathbb{N}_+$ there exists $b_n \in [0, \infty)$ with $b_n \rightarrow 0$ as $n \rightarrow \infty$ and $|f| \leq b_n$ almost everywhere on S . Hence $f = 0$ almost everywhere on S .

Suppose that $f : S \rightarrow \mathbb{R}$ is measurable and $c \in \mathbb{R}$. Then $\|cf\|_p = |c| \|f\|_p$ for $p \in (0, \infty]$.

Proof

Again, when $p \in (0, \infty)$, this result follows easily from properties of the integral that we already have:

$$\int_S |cf|^p d\mu = |c|^p \int_S |f|^p d\mu \quad (3.14.2)$$

Taking the p th root of both sides gives the result. For $p = \infty$, the result is trivially true if $c = 0$. For $c \neq 0$, note that $b \in [0, \infty]$ is an essential bound of $|f|$ if and only if $|c|b$ is an essential bound of $|cf|$.

In particular, if $f \in L^p$ and $c \in \mathbb{R}$ then $cf \in L^p$.

Conjugate Indices and Hölder's inequality

Certain pairs of our function spaces turn out to be *dual* or *complementary* to one another in a sense. To understand this, we need the following definition.

Indices $p, q \in (1, \infty)$ are said to be *conjugate* if $1/p + 1/q = 1$. In addition, 1 and ∞ are conjugate indices.

For justification of the last case, note that if $p \in (1, \infty)$, then the index conjugate to p is

$$q = \frac{1}{1 - 1/p} \quad (3.14.3)$$

and $q \uparrow \infty$ as $p \downarrow 1$. Note that $p = q = 2$ are conjugate indices, and this is the only case where the indices are the same. Ultimately, the importance of conjugate indices stems from the following inequality:

If $x, y \in (0, \infty)$ and if $p, q \in (1, \infty)$ are conjugate indices, then

$$xy \leq \frac{1}{p} x^p + \frac{1}{q} y^q \quad (3.14.4)$$

Moreover, equality occurs if and only if $x^p = y^q$.

Proof 1

From properties of the natural logarithm function,

$$\ln(xy) = \ln(x) + \ln(y) = \frac{1}{p} \ln(x^p) + \frac{1}{q} \ln(y^q) \quad (3.14.5)$$

But the natural logarithm function is concave and $1/p + 1/q = 1$ so

$$\ln(xy) = \frac{1}{p} \ln(x^p) + \frac{1}{q} \ln(y^q) \leq \ln\left(\frac{1}{p} x^p + \frac{1}{q} y^q\right) \quad (3.14.6)$$

Taking exponentials we have

$$xy \leq \frac{1}{p}x^p + \frac{1}{q}y^q \quad (3.14.7)$$

Proof 2

Fix $y \in (0, \infty)$ and define $f : (0, \infty) \rightarrow \mathbb{R}$ by

$$f(x) = \frac{1}{p}x^p + \frac{1}{q}y^q - xy, \quad x \in (0, \infty) \quad (3.14.8)$$

Then $f'(x) = x^{p-1} - y$ and $f''(x) = (p-1)x^{p-2}$ for $x \in (0, \infty)$. Hence f has a single critical point at $x = y^{1/(p-1)} = y^{q/p}$ and $f''(x) > 0$ for $x \in (0, \infty)$. It follows that the minimum value of f on $(0, \infty)$ occurs at $y^{q/p}$ and $f(y^{q/p}) = 0$. Hence $f(x) \geq 0$ for $x \in (0, \infty)$ with equality only at $x = y^{q/p}$ (that is, $x^p = y^q$).

our next major result is *Hölder's inequality*, named for Otto Hölder, which clearly indicates the importance of conjugate indices.

Suppose that $f, g : S \rightarrow \mathbb{R}$ are measurable and that p and q are conjugate indices. Then

$$\|fg\|_1 \leq \|f\|_p \|g\|_q \quad (3.14.9)$$

Proof

The result is obvious if $\|f\|_p = \infty$ or $\|g\|_q = \infty$, so suppose that $f \in L^p$ and $g \in L^q$. For our first case, suppose that $p = 1$ and $q = \infty$. Note that $|g| \leq \|g\|_\infty$ almost everywhere on S . Hence

$$\int_S |fg| d\mu = \int_S |f| |g| d\mu \leq \|g\|_\infty \int_S |f| d\mu = \|f\|_1 \|g\|_\infty \quad (3.14.10)$$

For the second case, suppose $p, q \in (1, \infty)$. By part (b) of the [positive property](#), the result holds if $\|f\|_p = 0$ or $\|g\|_q = 0$, so assume that $\|f\|_p > 0$ and $\|g\|_q > 0$. By the additivity of the integral over disjoint domains, we can restrict the integrals to the set $\{x \in S : f(x) \neq 0, g(x) \neq 0\}$, or simply assume that $f \neq 0$ and $g \neq 0$ on S . From the [basic inequality](#),

$$|fg| \leq \frac{1}{p}|f|^p + \frac{1}{q}|g|^q \quad (3.14.11)$$

Suppose first that $\|f\|_p = \|g\|_q = 1$. From the increasing and linearity properties of the integral,

$$\int_S |fg| d\mu \leq \frac{1}{p} \int_S |f|^p d\mu + \frac{1}{q} \int_S |g|^q d\mu = \frac{1}{p} + \frac{1}{q} = 1 \quad (3.14.12)$$

For the general case where $\|f\|_p > 0$ and $\|g\|_q > 0$, let $f_1 = f/\|f\|_p$ and $g_1 = g/\|g\|_q$. Then $\|f_1\|_p = \|g_1\|_q = 1$ so $\|f_1 g_1\|_1 \leq 1$. So by the [scaling property](#),

$$\|f_1 g_1\|_1 = \frac{\|fg\|_1}{\|f\|_p \|g\|_q} \leq 1 \quad (3.14.13)$$

In particular, if $f \in L^p$ and $g \in L^q$ then $fg \in L^1$. The most important special case of Hölder's inequality is when $p = q = 2$, in which case we have the *Cauchy-Schwartz inequality*, named for Augustin Louis Cauchy and Karl Hermann Schwarz:

$$\|fg\|_1 \leq \|f\|_2 \|g\|_2 \quad (3.14.14)$$

Minkowski's Inequality

Our next major result is *Minkowski's inequality*, named for Hermann Minkowski. This inequality will help show that L^p is a vector space and that $\|\cdot\|_p$ is a norm (up to equivalence) when $p \geq 1$.

Suppose that $f, g : S \rightarrow \mathbb{R}$ are measurable and that $p \in [1, \infty]$. Then

$$\|f + g\|_p \leq \|f\|_p + \|g\|_p \quad (3.14.15)$$

Proof

Again, the result is trivial if $\|f\|_p = \infty$ or $\|g\|_p = \infty$, so assume that $f, g \in L^p$. When $p = 1$, the result is the simple triangle inequality for the integral:

$$\|f + g\|_1 = \int_S |f + g| d\mu \leq \int_S (|f| + |g|) d\mu = \int_S |f| d\mu + \int_S |g| d\mu = \|f\|_1 + \|g\|_1 \quad (3.14.16)$$

For the case $p = \infty$, note that if $a \in [0, \infty]$ is an essential bound for f and $b \in [0, \infty]$ is an essential bound for g then $a + b$ is an essential bound for $f + g$. Hence $\|f + g\|_\infty \leq \|f\|_\infty + \|g\|_\infty$. For the last case, suppose that $p \in (1, \infty)$ and let q be the index conjugate to p . Then

$$|f + g|^p = |f + g|^{p-1} |f + g| \leq |f + g|^{p-1} (|f| + |g|) = |f + g|^{p-1} |f| + |f + g|^{p-1} |g| \quad (3.14.17)$$

Integrating over S and using the additive and increasing properties of the integral gives

$$\|f + g\|_p^p \leq \int_S |f + g|^{p-1} |f| d\mu + \int_S |f + g|^{p-1} |g| d\mu \quad (3.14.18)$$

But by Höder's inequality,

$$\int_S |f + g|^{p-1} |f| d\mu \leq \| |f + g|^{p-1} \|_q \|f\|_p, \quad \int_S |f + g|^{p-1} |g| d\mu \leq \| |f + g|^{p-1} \|_q \|g\|_p \quad (3.14.19)$$

Combining this with the previous inequality we have

$$\|f + g\|_p^p \leq \| |f + g|^{p-1} \|_q (\|f\|_p + \|g\|_p) \quad (3.14.20)$$

But $(p-1)q = p$ and $1/q = (p-1)/p$ so

$$\| |f + g|^{p-1} \|_q = \left(\int_S |f + g|^{(p-1)q} d\mu \right)^{1/q} = \left(\int_S |f + g|^p d\mu \right)^{(p-1)/p} = \|f + g\|_p^{p-1} \quad (3.14.21)$$

Hence we have

$$\|f + g\|_p^p \leq \|f + g\|_p^{p-1} (\|f\|_p + \|g\|_p) \quad (3.14.22)$$

and therefore $\|f + g\|_p \leq \|f\|_p + \|g\|_p$.

Vector Spaces

We can now discuss various vector spaces of functions. First, we know from our previous work with measure spaces, that the set \mathcal{V} of all measurable functions $f : S \rightarrow \mathbb{R}$ is a vector space under our standard (pointwise) definitions of sum and scalar multiple. The spaces we are studying in this section are subspaces:

L^p is a subspace of \mathcal{V} for every $p \in [1, \infty]$.

Proof

We just need to show that L^p is closed under addition and scalar multiplication. From the [positive property](#), if $f \in L^p$ and $c \in \mathbb{R}$ then $cf \in L^p$. From [Minkowski's inequality](#), if $f, g \in L^p$ then $f + g \in L^p$.

However, we usually want to *identify* functions that are equal almost everywhere on S (with respect to μ). Recalling the equivalence relation \equiv [defined above](#), here are the definitions:

Let $[f]$ denote the equivalence class of $f \in \mathcal{V}$ under the equivalence relation \equiv , and let $\mathcal{U} = \{[f] : f \in \mathcal{V}\}$. If $f, g \in \mathcal{V}$ and $c \in \mathbb{R}$ we define

1. $[f] + [g] = [f + g]$
2. $c[f] = [cf]$

Then \mathcal{U} is a vector space.

Proof

we know from our previous work that these definitions are consistent in the sense that they do not depend on the particular representatives of the equivalence classes. That is if $f_1 \equiv f$ and $g_1 \equiv g$ then $f_1 + g_1 \equiv f + g$ and $cf_1 \equiv cf$. That \mathcal{U} is a vector space then follows from the fact that \mathcal{V} is a vector space.

Now we can define the Lebesgue vector spaces precisely.

For $p \in [1, \infty]$, let $\mathcal{L}^p = \{[f] : f \in L^p\}$. For $f \in \mathcal{V}$ define $\|[f]\|_p = \|f\|_p$. Then \mathcal{L}^p is a subspace of \mathcal{U} and $\|\cdot\|_p$ is a norm on \mathcal{L}^p . That is, for $f, g \in L^p$ and $c \in \mathbb{R}$

1. $\|f\|_p \geq 0$ and $\|f\|_p = 0$ if and only if $f \equiv 0$, the *positive property*
2. $\|cf\|_p = |c| \|f\|_p$, the *scaling property*
3. $\|f + g\|_p \leq \|f\|_p + \|g\|_p$, the *triangle inequality*

Proof

That \mathcal{L}^p is a subspace of \mathcal{U} follows immediately from the fact that L^p is a subspace of \mathcal{V} . The fact that $\|\cdot\|_p$ is a norm on \mathcal{L}^p also follows from our previous work.

We have stated these results precisely, but on the other hand, we don't want to be overly pedantic. It's more natural and intuitive to simply work with the space \mathcal{V} and the subspaces L^p for $p \in [1, \infty]$, and just remember that functions that are equal almost everywhere on S are regarded as the same vector. This will be our point of view for the rest of this section.

Every norm on a vector space naturally leads to a metric. That is, we measure the distance between vectors as the norm of their difference. Stated in terms of the norm $\|\cdot\|_p$, here are the properties of the metric on L^p .

For $f, g, h \in L^p$,

1. $\|f - g\|_p \geq 0$ and $\|f - g\|_p = 0$ if and only if $f \equiv g$, the *positive property*
2. $\|f - g\|_p = \|g - f\|_p$, the *symmetric property*
3. $\|f - h\|_p \leq \|f - g\|_p + \|g - h\|_p$, the *triangle inequality*

Once we have a metric, we naturally have a criterion for *convergence*.

Suppose that $f_n \in L^p$ for $n \in \mathbb{N}_+$ and $f \in L^p$. Then by definition, $f_n \rightarrow f$ as $n \rightarrow \infty$ in L^p if and only if $\|f_n - f\|_p \rightarrow 0$ as $n \rightarrow \infty$.

Limits are unique, up to equivalence. (That is, limits *are* unique in \mathcal{L}^p .)

Suppose again that $f_n \in L^p$ for $n \in \mathbb{N}_+$. Recall that this sequence is said to be a *Cauchy sequence* if for every $\epsilon > 0$ there exists $N \in \mathbb{N}_+$ such that if $n > N$ and $m > N$ then $\|f_n - f_m\|_p < \epsilon$. Needless to say, the Cauchy criterion is named for our ubiquitous friend Augustin Cauchy. A metric space in which every Cauchy sequence converges (to an element of the space) is said to be *complete*. Intuitively, one expects a Cauchy sequence to converge, so a complete space is literally one that is not missing any elements that should be there. A complete, normed vector space is called a *Banach space*, after the Polish mathematician Stefan Banach. Banach spaces are of fundamental importance in analysis, in large part because of the following result:

L^p is a Banach space for every $p \in [1, \infty]$.

The Space L^2

The norm $\|\cdot\|_2$ is special because it corresponds to an *inner product*.

For $f, g \in L^2$, define

$$\langle f, g \rangle = \int_S fg d\mu \quad (3.14.23)$$

Note that the integral is well-defined by the Cauchy-Schwarz inequality. As with all of our other definitions, this one is consistent with the equivalence relation. That is, if $f \equiv f_1$ and $g \equiv g_1$ then $fg \equiv f_1g_1$ so $\int_S fg d\mu = \int_S f_1g_1 d\mu$ and hence $\langle f, g \rangle = \langle f_1, g_1 \rangle$. Note also that $\langle f, f \rangle = \|f\|_2^2$ for $f \in L^2$, so this definition generates the 2-norm.

L^2 is an inner product space. That is, if $f, g, h \in L^2$ and $c \in \mathbb{R}$ then

1. $\langle f, f \rangle \geq 0$ and $\langle f, f \rangle = 0$ if and only if $f \equiv 0$, the *positive property*
2. $\langle f, g \rangle = \langle g, f \rangle$, the *symmetric property*
3. $\langle cf, g \rangle = c\langle f, g \rangle$, the *scaling property*
4. $\langle f + g, h \rangle = \langle f, h \rangle + \langle g, h \rangle$, the *additive property*

Proof

Part (a) is a restatement of the positive property of the norm $\|\cdot\|_2$. Part (b) is obvious and parts (c) and (d) follow from the linearity of the integral.

From parts (c) and (d), the inner product is linear in the first argument, with the second argument fixed. By the symmetric property (b), it follows that the inner product is also linear in the second argument with the first argument fixed. That is, the inner product is *bi-linear*. A complete inner product space is known as a *Hilbert space*, named for the German mathematician David Hilbert. Thus, the following result follows immediately from the previous two.

L^2 is a Hilbert space.

All inner product spaces lead naturally to the concept of orthogonality; L^2 is no exception.

Functions $f, g \in L^2$ are *orthogonal* if $\langle f, g \rangle = 0$, in which case we write $f \perp g$. Equivalently $f \perp g$ if

$$\int_S fg d\mu = 0 \quad (3.14.24)$$

Of course, all of the basic theorems of general inner product spaces hold in L^2 . For example, the following result is the *Pythagorean theorem*, named of course for Pythagoras.

If $f, g \in L^2$ and $f \perp g$ then $\|f + g\|_2^2 = \|f\|_2^2 + \|g\|_2^2$.

Proof

The proof just uses basic properties of inner product in (17). No special properties of L^2 are used. If $f, g \in L^2$ and $f \perp g$ then

$$\|f + g\|^2 = \langle f + g, f + g \rangle = \langle f, f \rangle + 2\langle f, g \rangle + \langle g, g \rangle = \|f\|^2 + \|g\|^2 \quad (3.14.25)$$

Examples and Special Cases

Discrete Spaces

Recall again that the measure space $(S, \mathcal{S}, \#)$ is *discrete* if S is countable, $\mathcal{S} = \mathcal{P}(S)$ is the σ -algebra of all subsets of S , and of course, $\#$ is counting measure. In this case, recall that integrals are sums. The exposition will look more familiar if we use the notation of sequences rather than functions. Thus, let $x : S \rightarrow \mathbb{R}$, and denote the value of x at $i \in S$ by x_i rather than $x(i)$. For $p \in [1, \infty)$, the p -norm is

$$\|x\|_p = \left(\sum_{i \in S} |x_i|^p \right)^{1/p} \quad (3.14.26)$$

On the other hand, $\|x\|_\infty = \sup\{x_i : i \in S\}$. The only null set for $\#$ is \emptyset , so the equivalence relation \equiv is simply equality, and so the spaces L^p and \mathcal{L}^p are the same. For $p \in [1, \infty)$, $x \in L^p$ if and only if

$$\sum_{i \in S} |x_i|^p < \infty \quad (3.14.27)$$

When $p \in \mathbb{N}_+$ (as is often the case), this condition means that $\sum_{i \in S} x_i^p$ is absolutely convergent. On the other hand, $x \in L^\infty$ if and only if x is bounded. When $S = \mathbb{N}_+$, the space L^p is often denoted ℓ^p . The inner product on L^2 is

$$\langle x, y \rangle = \sum_{i \in S} x_i y_i, \quad x, y \in L^2 \quad (3.14.28)$$

When $S = \{1, 2, \dots, n\}$, L^2 is simply the vector space \mathbb{R}^n with the usual addition, scalar multiplication, inner product, and norm that we study in elementary linear algebra. Orthogonal vectors are perpendicular in the usual sense.

Probability Spaces

Suppose that $(S, \mathcal{S}, \mathbb{P})$ is a probability space, so that S is the set of outcomes of a random experiment, \mathcal{S} is the σ -algebra of events, and \mathbb{P} is a probability measure on the sample space (S, \mathcal{S}) . Of course, a measurable function $X : S \rightarrow \mathbb{R}$ is simply a real-valued random variable. For $p \in [1, \infty)$, the integral $\int_S |x|^p d\mathbb{P}$ is the *expected value* of $|X|^p$, and is denoted $\mathbb{E}(|X|^p)$. Thus in this case, L^p is the collection of real-valued random variables X with $\mathbb{E}(|X|^p) < \infty$. We will study these spaces in more detail in the chapter on expected value.

This page titled [3.14: Function Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

4: Expected Value

Expected value is one of the fundamental concepts in probability, in a sense more general than probability itself. The expected value of a real-valued random variable gives a measure of the center of the distribution of the variable. More importantly, by taking the expected value of various functions of a general random variable, we can measure many interesting features of its distribution, including spread, skewness, kurtosis, and correlation. Generating functions are certain types of expected value that completely determine the distribution of the variable. Conditional expected value, which incorporates known information in the computation, is one of the fundamental concepts in probability.

In the advanced topics, we define expected value as an integral with respect to the underlying probability measure. We also revisit conditional expected value from a measure-theoretic point of view. We study vector spaces of random variables with certain expected values as the norms of the spaces, which in turn leads to modes of convergence for random variables.

[4.1: Definitions and Basic Properties](#)

[4.2: Additional Properties](#)

[4.3: Variance](#)

[4.4: Skewness and Kurtosis](#)

[4.5: Covariance and Correlation](#)

[4.6: Generating Functions](#)

[4.7: Conditional Expected Value](#)

[4.8: Expected Value and Covariance Matrices](#)

[4.9: Expected Value as an Integral](#)

[4.10: Conditional Expected Value Revisited](#)

[4.11: Vector Spaces of Random Variables](#)

[4.12: Uniformly Integrable Variables](#)

[4.13: Kernels and Operators](#)

This page titled [4: Expected Value](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.1: Definitions and Basic Properties

Expected value is one of the most important concepts in probability. The expected value of a real-valued random variable gives the center of the distribution of the variable, in a special sense. Additionally, by computing expected values of various real transformations of a general random variable, we can extract a number of interesting characteristics of the distribution of the variable, including measures of spread, symmetry, and correlation. In a sense, expected value is a more general concept than probability itself.

Basic Concepts

Definitions

As usual, we start with a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} the collection of events and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . In the following definitions, we assume that X is a random variable for the experiment, taking values in $S \subseteq \mathbb{R}$.

If X has a discrete distribution with probability density function f (so that S is countable), then the *expected value* of X is defined as follows (assuming that the sum is well defined):

$$\mathbb{E}(X) = \sum_{x \in S} x f(x) \quad (4.1.1)$$

The sum defining the expected value makes sense if either the sum over the positive $x \in S$ is finite or the sum over the negative $x \in S$ is finite (or both). This ensures that the entire sum exists (as an extended real number) and does not depend on the order of the terms. So as we will see, it's possible for $\mathbb{E}(X)$ to be a real number or ∞ or $-\infty$ or to simply not exist. Of course, if S is finite the expected value always exists as a real number.

If X has a continuous distribution with probability density function f (and so S is typically an interval or a union of disjoint intervals), then the *expected value* of X is defined as follows (assuming that the integral is well defined):

$$\mathbb{E}(X) = \int_S x f(x) dx \quad (4.1.2)$$

The probability density functions in basic applied probability that describe continuous distributions are piecewise continuous. So the integral above makes sense if the integral over positive $x \in S$ is finite or the integral over negative $x \in S$ is finite (or both). This ensures that the entire integral exists (as an extended real number). So as in the discrete case, it's possible for $\mathbb{E}(X)$ to exist as a real number or as ∞ or as $-\infty$ or to not exist at all. As you might guess, the definition for a mixed distribution is a combination of the definitions for the discrete and continuous cases.

If X has a mixed distribution, with partial discrete density g on D and partial continuous density h on C , where D and C are disjoint, D is countable, C is typically an interval, and $S = D \cup C$. The expected value of X is defined as follows (assuming that the expression on the right is well defined):

$$\mathbb{E}(X) = \sum_{x \in D} x g(x) + \int_C x h(x) dx \quad (4.1.3)$$

For the expected value above to make sense, the sum must be well defined, as in the discrete case, the integral must be well defined, as in the continuous case, and we must avoid the dreaded indeterminate form $\infty - \infty$. In the next section on additional properties, we will see that the various definitions given here can be unified into a single definition that works regardless of the type of distribution of X . An even more general definition is given in the advanced section on expected value as an integral.

Interpretation

The expected value of X is also called the *mean* of the distribution of X and is frequently denoted μ . The mean is the center of the probability distribution of X in a special sense. Indeed, if we think of the distribution as a mass distribution (with total mass 1), then the mean is the *center of mass* as defined in physics. The two pictures below show discrete and continuous probability density functions; in each case the mean μ is the center of mass, the balance point.




Figure 4.1.1: The mean μ as the center of mass of a discrete distribution.




Figure 4.1.2: The mean μ as the center of mass of a continuous distribution.

Recall the other measures of the center of a distribution that we have studied:

- A mode is any $x \in S$ that maximizes f .
- A median is any $x \in \mathbb{R}$ that satisfies $\mathbb{P}(X < x) \leq \frac{1}{2}$ and $\mathbb{P}(X \leq x) \geq \frac{1}{2}$.

To understand expected value in a probabilistic way, suppose that we create a new, compound experiment by repeating the basic experiment over and over again. This gives a sequence of independent random variables (X_1, X_2, \dots) , each with the same distribution as X . In statistical terms, we are sampling from the distribution of X . The average value, or *sample mean*, after n runs is

$$M_n = \frac{1}{n} \sum_{i=1}^n X_i \quad (4.1.4)$$

Note that M_n is a random variable in the compound experiment. The important fact is that the *average value* M_n converges to the *expected value* $\mathbb{E}(X)$ as $n \rightarrow \infty$. The precise statement of this is the law of large numbers, one of the fundamental theorems of probability. You will see the law of large numbers at work in many of the simulation exercises given below.

Extensions

If $a \in \mathbb{R}$ and $n \in \mathbb{N}$, the *moment of X about a of order n* is defined to be

$$\mathbb{E}[(X - a)^n] \quad (4.1.5)$$

(assuming of course that this expected value exists).

The moments about 0 are simply referred to as *moments* (or sometimes *raw moments*). The moments about μ are the *central moments*. The second central moment is particularly important, and is studied in detail in the section on variance. In some cases, if we know *all* of the moments of X , we can determine the entire distribution of X . This idea is explored in the section on generating functions.

The expected value of a random variable X is based, of course, on the probability measure \mathbb{P} for the experiment. This probability measure could be a conditional probability measure, conditioned on a given event $A \in \mathcal{F}$ for the experiment (with $\mathbb{P}(A) > 0$). The usual notation is $\mathbb{E}(X | A)$, and this expected value is computed by the definitions given above, except that the conditional probability density function $x \mapsto f(x | A)$ replaces the ordinary probability density function f . It is very important to realize that, except for notation, no new concepts are involved. All results that we obtain for expected value in general have analogues for these conditional expected values. On the other hand, we will study a more general notion of conditional expected value in a later section.

Basic Properties

The purpose of this subsection is to study some of the essential properties of expected value. Unless otherwise noted, we will assume that the indicated expected values exist, and that the various sets and functions that we use are measurable. We start with two simple but still essential results.

Simple Variables

First, recall that a constant $c \in \mathbb{R}$ can be thought of as a random variable (on any probability space) that takes only the value c with probability 1. The corresponding distribution is sometimes called *point mass* at c .

If c is a constant random variable, then $\mathbb{E}(c) = c$.

Proof

As a random variable, c has a discrete distribution, so $\mathbb{E}(c) = c \cdot 1 = c$.

Next recall that an indicator variable is a random variable that takes only the values 0 and 1.

If X is an indicator variable then $\mathbb{E}(X) = \mathbb{P}(X = 1)$.

Proof

X is discrete so by definition, $\mathbb{E}(X) = 1 \cdot \mathbb{P}(X = 1) + 0 \cdot \mathbb{P}(X = 0) = \mathbb{P}(X = 1)$.

In particular, if $\mathbf{1}_A$ is the indicator variable of an event A , then $\mathbb{E}(\mathbf{1}_A) = \mathbb{P}(A)$, so in a sense, expected value subsumes probability. For a book that takes expected value, rather than probability, as the fundamental starting concept, see the book [Probability via Expectation](#), by Peter Whittle.

Change of Variables Theorem

The expected value of a real-valued random variable gives the center of the distribution of the variable. This idea is much more powerful than might first appear. By finding expected values of various *functions* of a general random variable, we can measure many interesting features of its distribution.

Thus, suppose that X is a random variable taking values in a general set S , and suppose that r is a function from S into \mathbb{R} . Then $r(X)$ is a real-valued random variable, and so it makes sense to compute $\mathbb{E}[r(X)]$ (assuming as usual that this expected value exists). However, to compute this expected value from the definition would require that we know the probability density function of the transformed variable $r(X)$ (a difficult problem, in general). Fortunately, there is a much better way, given by the *change of variables theorem* for expected value. This theorem is sometimes referred to as the *law of the unconscious statistician*, presumably because it is so basic and natural that it is often used without the realization that it is a theorem, and not a definition.

If X has a discrete distribution on a countable set S with probability density function f , then

$$\mathbb{E}[r(X)] = \sum_{x \in S} r(x)f(x) \quad (4.1.6)$$

Proof




Figure 4.1.3: The change of variables theorem when X has a discrete distribution.

The next result is the change of variables theorem when X has a continuous distribution. We will prove the continuous version in stages, first when r has discrete range below and then in the next section in full generality. Even though the complete proof is delayed, however, we will use the change of variables theorem in the proofs of many of the other properties of expected value.

Suppose that X has a continuous distribution on $S \subseteq \mathbb{R}^n$ with probability density function f , and that $r : S \rightarrow \mathbb{R}$. Then

$$\mathbb{E}[r(X)] = \int_S r(x)f(x) dx \quad (4.1.7)$$

Proof when r has discrete range




Figure 4.1.4: The change of variables theorem when X has a continuous distribution and r has countable range.

The results below gives basic properties of expected value. These properties are true in general, but we will restrict the proofs primarily to the continuous case. The proofs for the discrete case are analogous, with sums replacing integrals. The change of variables theorem is the main tool we will need. In these theorems X and Y are real-valued random variables for an experiment (that is, defined on an underlying probability space) and c is a constant. As usual, we assume that the indicated expected values exist. Be sure to try the proofs yourself before reading the ones in the text.

Linearity

Our first property is the *additive property*.

$$\mathbb{E}(X + Y) = \mathbb{E}(X) + \mathbb{E}(Y)$$

Proof

We apply the [change of variables theorem](#) with the function $r(x, y) = x + y$. Suppose that (X, Y) has a continuous distribution with PDF f , and that X takes values in $S \subseteq \mathbb{R}$ and Y takes values in $T \subseteq \mathbb{R}$. Recall that X has PDF g given by $g(x) = \int_T f(x, y) dy$ for $x \in S$ and Y has PDF h given by $h(y) = \int_S f(x, y) dx$ for $y \in T$. Thus

$$\mathbb{E}(X + Y) = \int_{S \times T} (x + y) f(x, y) d(x, y) = \int_{S \times T} x f(x, y) d(x, y) + \int_{S \times T} y f(x, y) d(x, y) \quad (4.1.8)$$

$$= \int_S x \left(\int_T f(x, y) dy \right) dx + \int_T y \left(\int_S f(x, y) dx \right) dy = \int_S x g(x) dx + \int_T y h(y) dy = \mathbb{E}(X) + \mathbb{E}(Y) \quad (4.1.9)$$

Writing the double integrals as iterated integrals is a special case of Fubini's theorem. The proof in the discrete case is the same, with sums replacing integrals.

Our next property is the *scaling property*.

$$\mathbb{E}(cX) = c \mathbb{E}(X)$$

Proof

We apply the [change of variables formula](#) with the function $r(x) = cx$. Suppose that X has a continuous distribution on $S \subseteq \mathbb{R}$ with PDF f . Then

$$\mathbb{E}(cX) = \int_S c x f(x) dx = c \int_S x f(x) dx = c \mathbb{E}(X) \quad (4.1.10)$$

Again, the proof in the discrete case is the same, with sums replacing integrals.

Here is the linearity of expected value in full generality. It's a simple corollary of the previous two results.

Suppose that (X_1, X_2, \dots) is a sequence of real-valued random variables defined on the underlying probability space and that (a_1, a_2, \dots, a_n) is a sequence of constants. Then

$$\mathbb{E} \left(\sum_{i=1}^n a_i X_i \right) = \sum_{i=1}^n a_i \mathbb{E}(X_i) \quad (4.1.11)$$

Thus, expected value is a *linear operation* on the collection of real-valued random variables for the experiment. The linearity of expected value is so basic that it is important to understand this property on an intuitive level. Indeed, it is implied by the interpretation of expected value given in the law of large numbers.

Suppose that (X_1, X_2, \dots, X_n) is a sequence of real-valued random variables with common mean μ .

1. Let $Y = \sum_{i=1}^n X_i$, the sum of the variables. Then $\mathbb{E}(Y) = n\mu$.
2. Let $M = \frac{1}{n} \sum_{i=1}^n X_i$, the average of the variables. Then $\mathbb{E}(M) = \mu$.

Proof

1. By the [additive property](#),

$$\mathbb{E}(Y) = \mathbb{E} \left(\sum_{i=1}^n X_i \right) = \sum_{i=1}^n \mathbb{E}(X_i) = \sum_{i=1}^n \mu = n\mu \quad (4.1.12)$$

2. Note that $M = Y/n$. Hence from the [scaling property](#) and part (a), $\mathbb{E}(M) = \mathbb{E}(Y)/n = \mu$.

If the random variables in the previous result are also independent and identically distributed, then in statistical terms, the sequence is a *random sample* of size n from the common distribution, and M is the *sample mean*.

In several important cases, a random variable from a special distribution can be decomposed into a sum of simpler random variables, and then part (a) of the [last theorem](#) can be used to compute the expected value.

Inequalities

The following exercises give some basic inequalities for expected value. The first, known as the *positive property* is the most obvious, but is also the main tool for proving the others.

Suppose that $\mathbb{P}(X \geq 0) = 1$. Then

1. $\mathbb{E}(X) \geq 0$
2. If $\mathbb{P}(X > 0) > 0$ then $\mathbb{E}(X) > 0$.

Proof

1. This result follows from the definition, since we can take the set of values S of X to be a subset of $[0, \infty)$.
2. Suppose that $\mathbb{P}(X > 0) > 0$ (in addition to $\mathbb{P}(X \geq 0) = 1$). By the continuity theorem for increasing events, there exists $\epsilon > 0$ such that $\mathbb{P}(X \geq \epsilon) > 0$. Therefore $X - \epsilon \mathbf{1}_{(X \geq \epsilon)} \geq 0$ (with probability 1). By part (a), linearity, and Theorem 2, $\mathbb{E}(X) - \epsilon \mathbb{P}(X \geq \epsilon) > 0$ so $\mathbb{E}(X) \geq \epsilon \mathbb{P}(X \geq \epsilon) > 0$.

Next is the *increasing property*, perhaps the most important property of expected value, after linearity.

Suppose that $\mathbb{P}(X \leq Y) = 1$. Then

1. $\mathbb{E}(X) \leq \mathbb{E}(Y)$
2. If $\mathbb{P}(X < Y) > 0$ then $\mathbb{E}(X) < \mathbb{E}(Y)$.

Proof

1. The assumption is equivalent to $\mathbb{P}(Y - X \geq 0) = 1$. Thus $\mathbb{E}(Y - X) \geq 0$ by part (a) of the [positive property](#). But then $\mathbb{E}(Y) - \mathbb{E}(X) \geq 0$ by the linearity of expected value.
2. Similarly, this result follows from part (b) of the [positive property](#).

Absolute value inequalities:

1. $|\mathbb{E}(X)| \leq \mathbb{E}(|X|)$
2. If $\mathbb{P}(X > 0) > 0$ and $\mathbb{P}(X < 0) > 0$ then $|\mathbb{E}(X)| < \mathbb{E}(|X|)$.

Proof

1. Note that $-|X| \leq X \leq |X|$ (with probability 1) so by part (a) of the [increasing property](#), $\mathbb{E}(-|X|) \leq \mathbb{E}(X) \leq \mathbb{E}(|X|)$. By linearity, $-\mathbb{E}(|X|) \leq \mathbb{E}(X) \leq \mathbb{E}(|X|)$ which implies $|\mathbb{E}(X)| \leq \mathbb{E}(|X|)$.
2. If $\mathbb{P}(X > 0) > 0$ then $\mathbb{P}(-|X| < X) > 0$, and if $\mathbb{P}(X < 0) > 0$ then $\mathbb{P}(X < |X|) > 0$. Hence by part (b) of the [increasing property](#), $-\mathbb{E}(|X|) < \mathbb{E}(X) < \mathbb{E}(|X|)$ and therefore $|\mathbb{E}(X)| < \mathbb{E}(|X|)$.

Only in [Lake Wobegone](#) are all of the children above average:

If $\mathbb{P}[X \neq \mathbb{E}(X)] > 0$ then

1. $\mathbb{P}[X > \mathbb{E}(X)] > 0$
2. $\mathbb{P}[X < \mathbb{E}(X)] > 0$

Proof

1. We prove the contrapositive. Thus suppose that $\mathbb{P}[X > \mathbb{E}(X)] = 0$ so that $\mathbb{P}[X \leq \mathbb{E}(X)] = 1$. If $\mathbb{P}[X < \mathbb{E}(X)] > 0$ then by the [increasing property](#) we have $\mathbb{E}(X) < \mathbb{E}(X)$, a contradiction. Thus $\mathbb{P}[X = \mathbb{E}(X)] = 1$.
2. Similarly, if $\mathbb{P}[X < \mathbb{E}(X)] = 0$ then $\mathbb{P}[X = \mathbb{E}(X)] = 1$.

Thus, if X is not a constant (with probability 1), then X must take values greater than its mean with positive probability and values less than its mean with positive probability.

Symmetry

Again, suppose that X is a random variable taking values in \mathbb{R} . The distribution of X is *symmetric* about $a \in \mathbb{R}$ if the distribution of $a - X$ is the same as the distribution of $X - a$.

Suppose that the distribution of X is symmetric about $a \in \mathbb{R}$. If $\mathbb{E}(X)$ exists, then $\mathbb{E}(X) = a$.

Proof

By assumption, the distribution of $X - a$ is the same as the distribution of $a - X$. Since $\mathbb{E}(X)$ exists we have $\mathbb{E}(a - X) = \mathbb{E}(X - a)$ so by linearity $a - \mathbb{E}(X) = \mathbb{E}(X) - a$. Equivalently $2\mathbb{E}(X) = 2a$.

The previous result applies if X has a continuous distribution on \mathbb{R} with a probability density f that is symmetric about a ; that is, $f(a + x) = f(a - x)$ for $x \in \mathbb{R}$.

Independence

If X and Y are independent real-valued random variables then $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y)$.

Proof

Suppose that X has a continuous distribution on $S \subseteq \mathbb{R}$ with PDF g and that Y has a continuous distribution on $T \subseteq \mathbb{R}$ with PDF h . Then (X, Y) has PDF $f(x, y) = g(x)h(y)$ on $S \times T$. We apply the [change of variables theorem](#) with the function $r(x, y) = xy$.

$$\mathbb{E}(XY) = \int_{S \times T} xyf(x, y) d(x, y) = \int_{S \times T} xyg(x)h(y) d(x, y) = \int_S xg(x) dx \int_T yh(y) dy = \mathbb{E}(X)\mathbb{E}(Y) \quad (4.1.13)$$

The proof in the discrete case is similar with sums replacing integrals.

It follows from the last result that independent random variables are uncorrelated (a concept that we will study in a later section). Moreover, this result is more powerful than might first appear. Suppose that X and Y are independent random variables taking values in general spaces S and T respectively, and that $u : S \rightarrow \mathbb{R}$ and $v : T \rightarrow \mathbb{R}$. Then $u(X)$ and $v(Y)$ are independent, real-valued random variables and hence

$$\mathbb{E}[u(X)v(Y)] = \mathbb{E}[u(X)]\mathbb{E}[v(Y)] \quad (4.1.14)$$

Examples and Applications

As always, be sure to try the proofs and computations yourself before reading the proof and answers in the text.

Uniform Distributions

Discrete uniform distributions are widely used in combinatorial probability, and model a point chosen *at random* from a finite set.

Suppose that X has the discrete uniform distribution on a finite set $S \subseteq \mathbb{R}$.

1. $\mathbb{E}(X)$ is the arithmetic average of the numbers in S .
2. If the points in S are evenly spaced with endpoints a, b , then $\mathbb{E}(X) = \frac{a+b}{2}$, the average of the endpoints.

Proof

1. Let $n = \#(S)$, the number of points in S . Then X has PDF $f(x) = 1/n$ for $x \in S$ so

$$\mathbb{E}(X) = \sum_{x \in S} x \frac{1}{n} = \frac{1}{n} \sum_{x \in S} x \quad (4.1.15)$$

2. Suppose that $S = \{a, a + h, a + 2h, \dots, a + (n - 1)h\}$ and let $b = a + (n - 1)h$, the right endpoint. As in (a), S has n points so using (a) and the formula for the sum of the first $n - 1$ positive integers, we have

$$\mathbb{E}(X) = \frac{1}{n} \sum_{i=0}^{n-1} (a + ih) = \frac{1}{n} \left(na + h \frac{(n-1)n}{2} \right) = a + \frac{(n-1)h}{2} = \frac{a+b}{2} \quad (4.1.16)$$

The previous results are easy to see if we think of $\mathbb{E}(X)$ as the center of mass, since the discrete uniform distribution corresponds to a finite set of points with equal mass.

Open the special distribution simulator, and select the discrete uniform distribution. This is the uniform distribution on n points, starting at a , evenly spaced at distance h . Vary the parameters and note the location of the mean in relation to the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean to the distribution mean.

Next, recall that the continuous uniform distribution on a bounded interval corresponds to selecting a point *at random* from the interval. Continuous uniform distributions arise in geometric probability and a variety of other applied problems.

Suppose that X has the continuous uniform distribution on an interval $[a, b]$, where $a, b \in \mathbb{R}$ and $a < b$.

1. $\mathbb{E}(X) = \frac{a+b}{2}$, the midpoint of the interval.
2. $\mathbb{E}(X^n) = \frac{1}{n+1} (a^n + a^{n-1}b + \dots + ab^{n-1} + b^n)$ for $n \in \mathbb{N}$.

Proof

1. Recall that X has PDF $f(x) = \frac{1}{b-a}$. Hence

$$\mathbb{E}(X) = \int_a^b x \frac{1}{b-a} dx = \frac{1}{b-a} \frac{b^2 - a^2}{2} = \frac{a+b}{2} \quad (4.1.17)$$

2. By the [change of variables formula](#),

$$\mathbb{E}(X^n) = \int_a^b \frac{1}{b-a} x^n dx = \frac{b^{n+1} - a^{n+1}}{(n+1)(b-a)} = \frac{1}{n+1} (a^n + a^{n-1}b + \dots + ab^{n-1} + b^n) \quad (4.1.18)$$

Part (a) is easy to see if we think of the mean as the center of mass, since the uniform distribution corresponds to a uniform distribution of mass on the interval.

Open the special distribution simulator, and select the continuous uniform distribution. This is the uniform distribution the interval $[a, a+w]$. Vary the parameters and note the location of the mean in relation to the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean to the distribution mean.

Next, the *average value* of a function on an interval, as defined in calculus, has a nice interpretation in terms of the uniform distribution.

Suppose that X is uniformly distributed on the interval $[a, b]$, and that g is an integrable function from $[a, b]$ into \mathbb{R} . Then $\mathbb{E}[g(X)]$ is the average value of g on $[a, b]$:

$$\mathbb{E}[g(X)] = \frac{1}{b-a} \int_a^b g(x) dx \quad (4.1.19)$$

Proof

This result follows immediately from the change of variables theorem, since X has PDF $f(x) = 1/(b-a)$ for $a \leq x \leq b$.

Find the average value of the following functions on the given intervals:

1. $f(x) = x$ on $[2, 4]$
2. $g(x) = x^2$ on $[0, 1]$
3. $h(x) = \sin(x)$ on $[0, \pi]$.

Answer

1. 3
2. $\frac{1}{3}$
3. $\frac{2}{\pi}$

The next exercise illustrates the value of the change of variables theorem in computing expected values.

Suppose that X is uniformly distributed on $[-1, 3]$.

1. Give the probability density function of X .
2. Find the probability density function of X^2 .
3. Find $\mathbb{E}(X^2)$ using the probability density function in (b).
4. Find $\mathbb{E}(X^2)$ using the [change of variables theorem](#).

Answer

1. $f(x) = \frac{1}{4}$ for $-1 \leq x \leq 3$
2. $g(y) = \begin{cases} \frac{1}{4}y^{-1/2}, & 0 < y < 1 \\ \frac{1}{8}y^{-1/2}, & 1 < y < 9 \end{cases}$

$$3. \int_0^9 yg(y) dy = \frac{7}{3}$$

$$4. \int_{-1}^3 x^2 f(x) dx = \frac{7}{3}$$

The discrete uniform distribution and the continuous uniform distribution are studied in more detail in the chapter on Special Distributions.

Dice

Recall that a *standard die* is a six-sided die. A *fair die* is one in which the faces are equally likely. An *ace-six flat die* is a standard die in which faces 1 and 6 have probability $\frac{1}{4}$ each, and faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each.

Two standard, fair dice are thrown, and the scores (X_1, X_2) recorded. Find the expected value of each of the following variables.

1. $Y = X_1 + X_2$, the sum of the scores.
2. $M = \frac{1}{2}(X_1 + X_2)$, the average of the scores.
3. $Z = X_1 X_2$, the product of the scores.
4. $U = \min\{X_1, X_2\}$, the minimum score
5. $V = \max\{X_1, X_2\}$, the maximum score.

Answer

1. 7
2. $\frac{7}{2}$
3. $\frac{49}{4}$
4. $\frac{101}{36}$
5. $\frac{19}{4}$

In the dice experiment, select two fair die. Note the shape of the probability density function and the location of the mean for the sum, minimum, and maximum variables. Run the experiment 1000 times and compare the sample mean and the distribution mean for each of these variables.

Two standard, ace-six flat dice are thrown, and the scores (X_1, X_2) recorded. Find the expected value of each of the following variables.

1. $Y = X_1 + X_2$, the sum of the scores.
2. $M = \frac{1}{2}(X_1 + X_2)$, the average of the scores.
3. $Z = X_1 X_2$, the product of the scores.
4. $U = \min\{X_1, X_2\}$, the minimum score
5. $V = \max\{X_1, X_2\}$, the maximum score.

Answer

1. 7
2. $\frac{7}{2}$
3. $\frac{49}{4}$
4. $\frac{77}{32}$
5. $\frac{147}{32}$

In the dice experiment, select two ace-six flat die. Note the shape of the probability density function and the location of the mean for the sum, minimum, and maximum variables. Run the experiment 1000 times and compare the sample mean and the distribution mean for each of these variables.

Bernoulli Trials

Recall that a *Bernoulli trials process* is a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent, identically distributed indicator random variables. In the usual language of reliability, X_i denotes the outcome of trial i , where 1 denotes success and 0 denotes failure. The probability of success $p = \mathbb{P}(X_i = 1) \in [0, 1]$ is the basic parameter of the process. The process is named for Jacob Bernoulli. A separate chapter on the Bernoulli Trials explores this process in detail.

For $n \in \mathbb{N}_+$, the number of successes in the first n trials is $Y = \sum_{i=1}^n X_i$. Recall that this random variable has the binomial distribution with parameters n and p , and has probability density function f given by

$$f(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (4.1.20)$$

If Y has the binomial distribution with parameters n and p then $\mathbb{E}(Y) = np$

Proof from the definition

The critical tools that we need involve binomial coefficients: the identity $y \binom{n}{y} = n \binom{n-1}{y-1}$ for $y, n \in \mathbb{N}_+$, and the binomial theorem:

$$\mathbb{E}(Y) = \sum_{y=0}^n y \binom{n}{y} p^y (1-p)^{n-y} = \sum_{y=1}^n n \binom{n-1}{y-1} p^y (1-p)^{n-y} \quad (4.1.21)$$

$$= np \sum_{y=1}^{n-1} \binom{n-1}{y-1} p^{y-1} (1-p)^{(n-1)-(y-1)} = np [p + (1-p)]^{n-1} = np \quad (4.1.22)$$

Proof using the additive property

Since $Y = \sum_{i=1}^n X_i$, the result follows immediately from the [expected value of an indicator variable](#) and the [additive property](#), since $\mathbb{E}(X_i) = p$ for each $i \in \mathbb{N}_+$.

Note the superiority of the second proof to the first. The result also makes intuitive sense: in n trials with success probability p , we expect np successes.

In the binomial coin experiment, vary n and p and note the shape of the probability density function and the location of the mean. For selected values of n and p , run the experiment 1000 times and compare the sample mean to the distribution mean.

Suppose that $p \in (0, 1]$, and let N denote the trial number of the first success. This random variable has the geometric distribution on \mathbb{N}_+ with parameter p , and has probability density function g given by

$$g(n) = p(1-p)^{n-1}, \quad n \in \mathbb{N}_+ \quad (4.1.23)$$

If N has the geometric distribution on \mathbb{N}_+ with parameter $p \in (0, 1]$ then $\mathbb{E}(N) = 1/p$.

Proof

The key is the formula for the derivative of a geometric series:

$$\mathbb{E}(N) = \sum_{n=1}^{\infty} np(1-p)^{n-1} = -p \frac{d}{dp} \sum_{n=0}^{\infty} (1-p)^n = -p \frac{d}{dp} \frac{1}{p} = p \frac{1}{p^2} = \frac{1}{p} \quad (4.1.24)$$

Again, the result makes intuitive sense. Since p is the probability of success, we expect a success to occur after $1/p$ trials.

In the negative binomial experiment, select $k = 1$ to get the geometric distribution. Vary p and note the shape of the probability density function and the location of the mean. For selected values of p , run the experiment 1000 times and compare the sample mean to the distribution mean.

The Hypergeometric Distribution

Suppose that a population consists of m objects; r of the objects are type 1 and $m-r$ are type 0. A sample of n objects is chosen at random, without replacement. The parameters $m, r, n \in \mathbb{N}$ with $r \leq m$ and $n \leq m$. Let X_i denote the type of the i th object selected. Recall that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of identically distributed (but *not* independent) indicator random variable with $\mathbb{P}(X_i = 1) = r/m$ for each $i \in \{1, 2, \dots, n\}$.

Let Y denote the number of type 1 objects in the sample, so that $Y = \sum_{i=1}^n X_i$. Recall that Y has the hypergeometric distribution, which has probability density function f given by

$$f(y) = \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}}, \quad y \in \{0, 1, \dots, n\} \quad (4.1.25)$$

If Y has the hypergeometric distribution with parameters m, n , and r then $\mathbb{E}(Y) = n \frac{r}{m}$.

Proof from the definition

Using the hypergeometric PDF,

$$\mathbb{E}(Y) = \sum_{y=0}^n y \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}} \quad (4.1.26)$$

Note that the $y = 0$ term is 0. For the other terms, we can use the identity $y \binom{r}{y} = r \binom{r-1}{y-1}$ to get

$$\mathbb{E}(Y) = \frac{r}{\binom{m}{n}} \sum_{y=1}^n \binom{r-1}{y-1} \binom{m-r}{n-y} \quad (4.1.27)$$

But substituting $k = y - 1$ and using another fundamental identity,

$$\sum_{y=1}^n \binom{r-1}{y-1} \binom{m-r}{n-y} = \sum_{k=0}^{n-1} \binom{r-1}{k} \binom{m-r}{n-1-k} = \binom{m-1}{n-1} \quad (4.1.28)$$

So substituting and doing a bit of algebra gives $\mathbb{E}(Y) = n \frac{r}{m}$.

Proof using the additive property

A much better proof uses the [additive property](#) and the representation of Y as a sum of indicator variables. The result follows immediately since $\mathbb{E}(X_i) = r/m$ for each $i \in \{1, 2, \dots, n\}$.

In the ball and urn experiment, vary n , r , and m and note the shape of the probability density function and the location of the mean. For selected values of the parameters, run the experiment 1000 times and compare the sample mean to the distribution mean.

Note that if we select the objects *with replacement*, then \mathbf{X} would be a sequence of Bernoulli trials, and hence Y would have the binomial distribution with parameters n and $p = \frac{r}{m}$. Thus, the mean would still be $\mathbb{E}(Y) = n \frac{r}{m}$.

The Poisson Distribution

Recall that the *Poisson distribution* has probability density function f given by

$$f(n) = e^{-a} \frac{a^n}{n!}, \quad n \in \mathbb{N} \quad (4.1.29)$$

where $a \in (0, \infty)$ is a parameter. The Poisson distribution is named after Simeon Poisson and is widely used to model the number of “random points” in a region of time or space; the parameter a is proportional to the size of the region. The Poisson distribution is studied in detail in the chapter on the Poisson Process.

If N has the Poisson distribution with parameter a then $\mathbb{E}(N) = a$. Thus, the parameter of the Poisson distribution is the mean of the distribution.

Proof

The proof depends on the standard series for the exponential function

$$\mathbb{E}(N) = \sum_{n=0}^{\infty} n e^{-a} \frac{a^n}{n!} = e^{-a} \sum_{n=1}^{\infty} \frac{a^n}{(n-1)!} = e^{-a} a \sum_{n=1}^{\infty} \frac{a^{n-1}}{(n-1)!} = e^{-a} a e^a = a. \quad (4.1.30)$$

In the Poisson experiment, the parameter is $a = rt$. Vary the parameter and note the shape of the probability density function and the location of the mean. For various values of the parameter, run the experiment 1000 times and compare the sample mean to the distribution mean.

The Exponential Distribution

Recall that the *exponential distribution* is a continuous distribution with probability density function f given by

$$f(t) = r e^{-rt}, \quad t \in [0, \infty) \quad (4.1.31)$$

where $r \in (0, \infty)$ is the *rate parameter*. This distribution is widely used to model failure times and other “arrival times”; in particular, the distribution governs the time between arrivals in the Poisson model. The exponential distribution is studied in detail in the chapter on the Poisson Process.

Suppose that T has the exponential distribution with rate parameter r . Then $\mathbb{E}(T) = 1/r$.

Proof

This result follows from the definition and an integration by parts:

$$\mathbb{E}(T) = \int_0^{\infty} t r e^{-rt} dt = -te^{-rt} \Big|_0^{\infty} + \int_0^{\infty} e^{-rt} dt = 0 - \frac{1}{r} e^{-rt} \Big|_0^{\infty} = \frac{1}{r} \quad (4.1.32)$$

Recall that the mode of T is 0 and the median of T is $\ln 2/r$. Note how these measures of center are ordered: $0 < \ln 2/r < 1/r$

In the gamma experiment, set $n = 1$ to get the exponential distribution. This app simulates the first arrival in a Poisson process. Vary r with the scroll bar and note the position of the mean relative to the graph of the probability density function. For selected values of r , run the experiment 1000 times and compare the sample mean to the distribution mean.

Suppose again that T has the exponential distribution with rate parameter r and suppose that $t > 0$. Find $\mathbb{E}(T \mid T > t)$.

Answer

$$t + \frac{1}{r}$$

The Gamma Distribution

Recall that the *gamma distribution* is a continuous distribution with probability density function f given by

$$f(t) = r^n \frac{t^{n-1}}{(n-1)!} e^{-rt}, \quad t \in [0, \infty) \quad (4.1.33)$$

where $n \in \mathbb{N}_+$ is the *shape parameter* and $r \in (0, \infty)$ is the *rate parameter*. This distribution is widely used to model failure times and other “arrival times”, and in particular, models the n th arrival in the Poisson process. Thus it follows that if (X_1, X_2, \dots, X_n) is a sequence of independent random variables, each having the exponential distribution with rate parameter r , then $T = \sum_{i=1}^n X_i$ has the gamma distribution with shape parameter n and rate parameter r . The gamma distribution is studied in more generality, with non-integer shape parameters, in the chapter on the Special Distributions.

Suppose that T has the gamma distribution with shape parameter n and rate parameter r . Then $\mathbb{E}(T) = n/r$.

Proof from the definition

The proof is by induction on n , so let μ_n denote the mean when the shape parameter is $n \in \mathbb{N}_+$. When $n = 1$, we have the exponential distribution with rate parameter r , so we know $\mu_1 = 1/r$ by our [result above](#). Suppose that $\mu_n = r/n$ for a given $n \in \mathbb{N}_+$. Then

$$\mu_{n+1} = \int_0^{\infty} t r^{n+1} \frac{t^n}{n!} e^{-rt} dt = \int_0^{\infty} r^{n+1} \frac{t^{n+1}}{n!} e^{-rt} dt \quad (4.1.34)$$

Integrate by parts with $u = \frac{t^{n+1}}{n!}$, $dv = r^{n+1} e^{-rt} dt$ so that $du = (n+1) \frac{t^n}{n!} dt$ and $v = -r^n e^{-rt}$. Then

$$\mu_{n+1} = (n+1) \int_0^{\infty} r^n \frac{t^n}{n!} e^{-rt} dt = \frac{n+1}{n} \int_0^{\infty} t r^n \frac{t^{n-1}}{(n-1)!} e^{-rt} dt \quad (4.1.35)$$

But the last integral is μ_n , so by the induction hypothesis, $\mu_{n+1} = \frac{n+1}{n} \frac{n}{r} = \frac{n+1}{r}$.

Proof using the additive property

The result follows immediately from the [additive property](#) and the fact that T can be represented in the form $T = \sum_{i=1}^n X_i$ where X_i has the exponential distribution with parameter r for each $i \in \{1, 2, \dots, n\}$.

Note again how much easier and more intuitive the second proof is than the first.

Open the gamma experiment, which simulates the arrival times in the Poisson process. Vary the parameters and note the position of the mean relative to the graph of the probability density function. For selected parameter values, run the experiment 1000 times and compare the sample mean to the distribution mean.

Beta Distributions

The distributions in this subsection belong to the family of *beta distributions*, which are widely used to model random proportions and probabilities. The beta distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has probability density function f given by $f(x) = 3x^2$ for $x \in [0, 1]$.

1. Find the mean of X .

2. Find the mode of X .
3. Find the median of X .
4. Sketch the graph of f and show the location of the mean, median, and mode on the x -axis.

Answer

1. $\frac{3}{4}$
2. 1
3. $(\frac{1}{2})^{1/3}$

In the special distribution simulator, select the beta distribution and set $a = 3$ and $b = 1$ to get the distribution in the last exercise. Run the experiment 1000 times and compare the sample mean to the distribution mean.

Suppose that a sphere has a random radius R with probability density function f given by $f(r) = 12r^2(1-r)$ for $r \in [0, 1]$. Find the expected value of each of the following:

1. The circumference $C = 2\pi R$
2. The surface area $A = 4\pi R^2$
3. The volume $V = \frac{4}{3}\pi R^3$

Answer

1. $\frac{6}{5}\pi$
2. $\frac{8}{5}\pi$
3. $\frac{8}{21}\pi$

Suppose that X has probability density function f given by $f(x) = \frac{1}{\pi\sqrt{x(1-x)}}$ for $x \in (0, 1)$.

1. Find the mean of X .
2. Find median of X .
3. Note that f is unbounded, so X does not have a mode.
4. Sketch the graph of f and show the location of the mean and median on the x -axis.

Answer

1. $\frac{1}{2}$
2. $\frac{1}{2}$

The particular beta distribution in the last exercise is also known as the (standard) *arcsine distribution*. It governs the last time that the Brownian motion process hits 0 during the time interval $[0, 1]$. The arcsine distribution is studied in more generality in the chapter on Special Distributions.

Open the Brownian motion experiment and select the last zero. Run the simulation 1000 times and compare the sample mean to the distribution mean.

Suppose that the grades on a test are described by the random variable $Y = 100X$ where X has the beta distribution with probability density function f given by $f(x) = 12x(1-x)^2$ for $x \in [0, 1]$. The grades are generally low, so the teacher decides to “curve” the grades using the transformation $Z = 10\sqrt{Y} = 100\sqrt{X}$. Find the expected value of each of the following variables

1. X
2. Y
3. Z

Answer

1. $\mathbb{E}(X) = \frac{2}{5}$
2. $\mathbb{E}(Y) = 40$
3. $\mathbb{E}(Z) = \frac{1280}{21} \approx 60.95$

The Pareto Distribution

Recall that the *Pareto distribution* is a continuous distribution with probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (4.1.36)$$

where $a \in (0, \infty)$ is a parameter. The Pareto distribution is named for Vilfredo Pareto. It is a heavy-tailed distribution that is widely used to model certain financial variables. The Pareto distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with shape parameter a . Then

1. $\mathbb{E}(X) = \infty$ if $0 < a \leq 1$
2. $\mathbb{E}(X) = \frac{a}{a-1}$ if $a > 1$

Proof

1. If $0 < a < 1$,

$$\mathbb{E}(X) = \int_1^\infty x \frac{a}{x^{a+1}} dx = \int_1^\infty \frac{a}{x^a} dx = \frac{a}{-a+1} x^{-a+1} \Big|_1^\infty = \infty \quad (4.1.37)$$

since the exponent $-a+1 > 0$. If $a = 1$, $\mathbb{E}(X) = \int_1^\infty x \frac{1}{x^2} dx = \int_1^\infty \frac{1}{x} dx = \ln x \Big|_1^\infty = \infty$.

2. If $a > 1$ then

$$\mathbb{E}(X) = \int_1^\infty x \frac{a}{x^{a+1}} dx = \int_1^\infty \frac{a}{x^a} dx = \frac{a}{-a+1} x^{-a+1} \Big|_1^\infty = \frac{a}{a-1} \quad (4.1.38)$$

The previous exercise gives us our first example of a distribution whose mean is infinite.

In the special distribution simulator, select the Pareto distribution. Note the shape of the probability density function and the location of the mean. For the following values of the shape parameter a , run the experiment 1000 times and note the behavior of the empirical mean.

1. $a = 1$
2. $a = 2$
3. $a = 3$.

The Cauchy Distribution

Recall that the (standard) *Cauchy distribution* has probability density function f given by

$$f(x) = \frac{1}{\pi(1+x^2)}, \quad x \in \mathbb{R} \quad (4.1.39)$$

This distribution is named for Augustin Cauchy. The Cauchy distributions is studied in detail in the chapter on Special Distributions.

If X has the Cauchy distribution then $\mathbb{E}(X)$ does not exist.

Proof

By definition,

$$\mathbb{E}(X) = \int_{-\infty}^\infty x \frac{1}{\pi(1+x^2)} dx = \frac{1}{2\pi} \ln(1+x^2) \Big|_{-\infty}^\infty \quad (4.1.40)$$

which evaluates to the meaningless expression $\infty - \infty$.

Note that the graph of f is symmetric about 0 and is unimodal. Thus, the mode and median of X are both 0. By the [symmetry result](#), if X had a mean, the mean would be 0 also, but alas the mean does not exist. Moreover, the non-existence of the mean is not just a pedantic technicality. If we think of the probability distribution as a mass distribution, then the moment to the right of a is $\int_a^\infty (x-a)f(x) dx = \infty$ and the moment to the left of a is $\int_{-\infty}^a (x-a)f(x) dx = -\infty$ for every $a \in \mathbb{R}$. The center of mass simply does not exist. Probabilistically, the law of large numbers fails, as you can see in the following simulation exercise:

In the Cauchy experiment (with the default parameter values), a light source is 1 unit from position 0 on an infinite straight wall. The angle that the light makes with the perpendicular is uniformly distributed on the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$, so that the position of the light beam on the wall has the Cauchy distribution. Run the simulation 1000 times and note the behavior of the empirical mean.

The Normal Distribution

Recall that the *standard normal distribution* is a continuous distribution with density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (4.1.41)$$

Normal distributions are widely used to model physical measurements subject to small, random errors and are studied in detail in the chapter on Special Distributions.

If Z has the standard normal distribution then $\mathbb{E}(Z) = 0$.

Proof

Using a simple change of variables, we have

$$\mathbb{E}(Z) = \int_{-\infty}^{\infty} z \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} dz = -\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} \Big|_{-\infty}^{\infty} = 0 - 0 \quad (4.1.42)$$

The standard normal distribution is unimodal and symmetric about 0. Thus, the median, mean, and mode all agree. More generally, for $\mu \in (-\infty, \infty)$ and $\sigma \in (0, \infty)$, recall that $X = \mu + \sigma Z$ has the normal distribution with *location parameter* μ and *scale parameter* σ . X has probability density function f given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (4.1.43)$$

The location parameter is the mean of the distribution:

If X has the normal distribution with location parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$, then $\mathbb{E}(X) = \mu$

Proof

Of course we could use the definition, but a proof using [linearity](#) and the representation in terms of the standard normal distribution is trivial: $\mathbb{E}(X) = \mu + \sigma \mathbb{E}(Z) = \mu$.

In the special distribution simulator, select the normal distribution. Vary the parameters and note the location of the mean. For selected parameter values, run the simulation 1000 times and compare the sample mean to the distribution mean.

Additional Exercises

Suppose that (X, Y) has probability density function f given by $f(x, y) = x + y$ for $(x, y) \in [0, 1] \times [0, 1]$. Find the following expected values:

1. $\mathbb{E}(X)$
2. $\mathbb{E}(X^2 Y)$
3. $\mathbb{E}(X^2 + Y^2)$
4. $\mathbb{E}(XY \mid Y > X)$

Answer

1. $\frac{7}{12}$
2. $\frac{17}{72}$
3. $\frac{5}{6}$
4. $\frac{1}{3}$

Suppose that N has a discrete distribution with probability density function f given by $f(n) = \frac{1}{50}n^2(5-n)$ for $n \in \{1, 2, 3, 4\}$. Find each of the following:

1. The median of N .
2. The mode of N
3. $\mathbb{E}(N)$.
4. $\mathbb{E}(N^2)$
5. $\mathbb{E}(1/N)$.
6. $\mathbb{E}(1/N^2)$.

Answer

1. 3
2. 3
3. $\frac{73}{25}$
4. $\frac{47}{5}$
5. $\frac{2}{5}$
6. $\frac{1}{5}$

Suppose that X and Y are real-valued random variables with $\mathbb{E}(X) = 5$ and $\mathbb{E}(Y) = -2$. Find $\mathbb{E}(3X + 4Y - 7)$.

Answer

0

Suppose that X and Y are real-valued, independent random variables, and that $\mathbb{E}(X) = 5$ and $\mathbb{E}(Y) = -2$. Find $\mathbb{E}[(3X - 4)(2Y + 7)]$.

Answer

33

Suppose that there are 5 duck hunters, each a perfect shot. A flock of 10 ducks fly over, and each hunter selects one duck at random and shoots. Find the expected number of ducks killed.

Solution

Number the ducks from 1 to 10. For $k \in \{1, 2, \dots, 10\}$ let X_k be the indicator variable that takes the value 1 if duck k is killed and 0 otherwise. Duck k is killed if at least one of the hunters selects her, so $\mathbb{E}(X_k) = \mathbb{P}(X_k = 1) = 1 - \left(\frac{9}{10}\right)^5$. The number of ducks killed is $N = \sum_{k=1}^{10} X_k$ so $\mathbb{E}(N) = 10 \left[1 - \left(\frac{9}{10}\right)^5\right] = 4.095$

For a more complete analysis of the duck hunter problem, see The Number of Distinct Sample Values in the chapter on Finite Sampling Models.

Consider the following game: An urn initially contains one red and one green ball. A ball is selected at random, and if the ball is green, the game is over. If the ball is red, the ball is returned to the urn, another red ball is added, and the game continues. At each stage, a ball is selected at random, and if the ball is green, the game is over. If the ball is red, the ball is returned to the urn, another red ball is added, and the game continues. Let X denote the length of the game (that is, the number of selections required to obtain a green ball). Find $\mathbb{E}(X)$.

Solution

The probability density function f of X was found in the section on discrete distributions: $f(x) = \frac{1}{x(x+1)}$ for $x \in \mathbb{N}_+$. The expected length of the game is infinite:

$$\mathbb{E}(X) = \sum_{x=1}^{\infty} x \frac{1}{x(x+1)} = \sum_{x=1}^{\infty} \frac{1}{x+1} = \infty \quad (4.1.44)$$

This page titled 4.1: Definitions and Basic Properties is shared under a CC BY 2.0 license and was authored, remixed, and/or curated by Kyle Siegrist (Random Services) via source content that was edited to the style and standards of the LibreTexts platform.

4.2: Additional Properties

In this section, we study some properties of expected value that are a bit more specialized than the basic properties considered in the previous section. Nonetheless, the new results are also very important. They include two fundamental inequalities as well as special formulas for the expected value of a nonnegative variable. As usual, unless otherwise noted, we assume that the referenced expected values exist.

Basic Theory

Markov's Inequality

Our first result is known as *Markov's inequality* (named after Andrei Markov). It gives an upper bound for the tail probability of a nonnegative random variable in terms of the expected value of the variable.

If X is a nonnegative random variable, then

$$\mathbb{P}(X \geq x) \leq \frac{\mathbb{E}(X)}{x}, \quad x > 0 \quad (4.2.1)$$

Proof

For $x > 0$, note that $x \cdot \mathbf{1}(X \geq x) \leq X$. Taking expected values through this inequality gives $x\mathbb{P}(X \geq x) \leq \mathbb{E}(X)$.

The upper bound in Markov's inequality may be rather crude. In fact, it's quite possible that $\mathbb{E}(X)/x \geq 1$, in which case the bound is worthless. However, the real value of Markov's inequality lies in the fact that it holds with no assumptions whatsoever on the distribution of X (other than that X be nonnegative). Also, as an [example below](#) shows, the inequality is *tight* in the sense that equality can hold for a given x . Here is a simple corollary of Markov's inequality.

If X is a real-valued random variable and $k \in (0, \infty)$ then

$$\mathbb{P}(|X| \geq x) \leq \frac{\mathbb{E}(|X|^k)}{x^k} \quad x > 0 \quad (4.2.2)$$

Proof

Since $k \geq 0$, the function $x \mapsto x^k$ is strictly increasing on $[0, \infty)$. Hence using Markov's inequality,

$$\mathbb{P}(|X| \geq x) = \mathbb{P}(|X|^k \geq x^k) \leq \frac{\mathbb{E}(|X|^k)}{x^k} \quad (4.2.3)$$

In this corollary of Markov's inequality, we could try to find $k > 0$ so that $\mathbb{E}(|X|^k)/x^k$ is minimized, thus giving the tightest bound on $\mathbb{P}(|X| \geq x)$.

Right Distribution Function

Our next few results give alternative ways to compute the expected value of a nonnegative random variable by means of the right-tail distribution function. This function also known as the *reliability function* if the variable represents the lifetime of a device.

If X is a nonnegative random variable then

$$\mathbb{E}(X) = \int_0^\infty \mathbb{P}(X > x) dx \quad (4.2.4)$$

Proof

A proof can be constructed by expressing $\mathbb{P}(X > x)$ in terms of the probability density function of X , as a sum in the discrete case or an integral in the continuous case. Then in the expression $\int_0^\infty \mathbb{P}(X > x) dx$ interchange the integral and the sum (in the discrete case) or the two integrals (in the continuous case). There is a much more elegant proof if we use the fact that we can interchange expected values and integrals when the integrand is nonnegative:

$$\int_0^\infty \mathbb{P}(X > x) dx = \int_0^\infty \mathbb{E}[\mathbf{1}(X > x)] dx = \mathbb{E}\left(\int_0^\infty \mathbf{1}(X > x) dx\right) = \mathbb{E}\left(\int_0^X 1 dx\right) = \mathbb{E}(X) \quad (4.2.5)$$

This interchange is a special case of *Fubini's theorem*, named for the Italian mathematician Guido Fubini. See the advanced section on expected value as an integral for more details.

Here is a slightly more general result:

If X is a nonnegative random variable and $k \in (0, \infty)$ then

$$\mathbb{E}(X^k) = \int_0^\infty kx^{k-1} \mathbb{P}(X > x) dx \quad (4.2.6)$$

Proof

The same basic proof works:

$$\int_0^\infty kx^{k-1} \mathbb{P}(X > x) dx = \int_0^\infty kx^{k-1} \mathbb{E}[\mathbf{1}(X > x)] dx = \mathbb{E}\left(\int_0^\infty kx^{k-1} \mathbf{1}(X > x) dx\right) = \mathbb{E}\left(\int_0^X kx^{k-1} dx\right) = \mathbb{E}(X^k) \quad (4.2.7)$$

The following result is similar to the [theorem above](#), but is specialized to nonnegative integer valued variables:

Suppose that N has a discrete distribution, taking values in \mathbb{N} . Then

$$\mathbb{E}(N) = \sum_{n=0}^\infty \mathbb{P}(N > n) = \sum_{n=1}^\infty \mathbb{P}(N \geq n) \quad (4.2.8)$$

Proof

First, the two sums on the right are equivalent by a simple change of variables. A proof can be constructed by expressing $\mathbb{P}(N > n)$ as a sum in terms of the probability density function of N . Then in the expression $\sum_{n=0}^\infty \mathbb{P}(N > n)$ interchange the two sums. Here is a more elegant proof:

$$\sum_{n=1}^\infty \mathbb{P}(N \geq n) = \sum_{n=1}^\infty \mathbb{E}[\mathbf{1}(N \geq n)] = \mathbb{E}\left(\sum_{n=1}^\infty \mathbf{1}(N \geq n)\right) = \mathbb{E}\left(\sum_{n=1}^N 1\right) = \mathbb{E}(N) \quad (4.2.9)$$

This interchange is a special case of a general rule that allows the interchange of expected value and an infinite series, when the terms are nonnegative. See the advanced section on expected value as an integral for more details.

A General Definition

The [special expected value formula](#) for nonnegative variables can be used as the basis of a general formulation of expected value that would work for discrete, continuous, or even mixed distributions, and would not require the assumption of the existence of probability density functions. First, the special formula is taken as the definition of $\mathbb{E}(X)$ if X is nonnegative.

If X is a nonnegative random variable, define

$$\mathbb{E}(X) = \int_0^\infty \mathbb{P}(X > x) dx \quad (4.2.10)$$

Next, for $x \in \mathbb{R}$, recall that the *positive and negative parts* of x are $x^+ = \max\{x, 0\}$ and $x^- = \max\{0, -x\}$.

For $x \in \mathbb{R}$,

1. $x^+ \geq 0, x^- \geq 0$
2. $x = x^+ - x^-$
3. $|x| = x^+ + x^-$

Now, if X is a real-valued random variable, then X^+ and X^- , the positive and negative parts of X , are nonnegative random variables, so their expected values are defined as above. The definition of $\mathbb{E}(X)$ is then natural, anticipating of course the linearity property.

If X is a real-valued random variable, define $\mathbb{E}(X) = \mathbb{E}(X^+) - \mathbb{E}(X^-)$, assuming that at least one of the expected values on the right is finite.

The usual formulas for expected value in terms of the probability density function, for discrete, continuous, or mixed distributions, would now be proven as theorems. We will not go further in this direction, however, since the most complete and general definition of expected value is given in the advanced section on expected value as an integral.

The Change of Variables Theorem

Suppose that X takes values in S and has probability density function f . Suppose also that $r: S \rightarrow \mathbb{R}$, so that $r(X)$ is a real-valued random variable. The *change of variables theorem* gives a formula for computing $\mathbb{E}[r(X)]$ without having to first find the probability density function of $r(X)$. If S is countable, so that X has a discrete distribution, then

$$\mathbb{E}[r(X)] = \sum_{x \in S} r(x)f(x) \quad (4.2.11)$$

If $S \subseteq \mathbb{R}^n$ and X has a continuous distribution on S then

$$\mathbb{E}[r(X)] = \int_S r(x)f(x) dx \quad (4.2.12)$$

In both cases, of course, we assume that the expected values exist. In the previous section on basic properties, we proved the change of variables theorem when X has a discrete distribution and when X has a continuous distribution but r has countable range. Now we can finally finish our proof in the continuous case.

Suppose that X has a continuous distribution on S with probability density function f , and $r : S \rightarrow \mathbb{R}$. Then

$$\mathbb{E}[r(X)] = \int_S r(x)f(x) dx \quad (4.2.13)$$

Proof

Suppose first that r is nonnegative. From the [theorem above](#),

$$\mathbb{E}[r(X)] = \int_0^\infty \mathbb{P}[r(X) > t] dt = \int_0^\infty \int_{r^{-1}(t, \infty)} f(x) dx dt = \int_S \int_0^{r(x)} f(x) dt dx = \int_S r(x)f(x) dx \quad (4.2.14)$$

For general r , we decompose into positive and negative parts, and use the result just established.

$$\mathbb{E}[r(X)] = \mathbb{E}[r^+(X) - r^-(X)] = \mathbb{E}[r^+(X)] - \mathbb{E}[r^-(X)] \quad (4.2.15)$$

$$= \int_S r^+(x)f(x) dx - \int_S r^-(x)f(x) dx = \int_S [r^+(x) - r^-(x)] f(x) dx = \int_S r(x)f(x) dx \quad (4.2.16)$$

Jensens's Inequality

Our next sequence of exercises will establish an important inequality known as *Jensen's inequality*, named for Johan Jensen. First we need a definition.

A real-valued function g defined on an interval $S \subseteq \mathbb{R}$ is said to be *convex* (or *concave upward*) on S if for each $t \in S$, there exist numbers a and b (that may depend on t), such that

1. $a + bt = g(t)$
2. $a + bx \leq g(x)$ for all $x \in S$

The graph of $x \mapsto a + bx$ is called a *supporting line* for g at t .

Thus, a convex function has *at least one* supporting line at each point in the domain




Figure 4.2.1: A convex function and several supporting lines

You may be more familiar with convexity in terms of the following theorem from calculus: If g has a continuous, non-negative second derivative on S , then g is convex on S (since the tangent line at t is a supporting line at t for each $t \in S$). The next result is the single variable version of Jensen's inequality

If X takes values in an interval S and $g : S \rightarrow \mathbb{R}$ is convex on S , then

$$\mathbb{E}[g(X)] \geq g[\mathbb{E}(X)] \quad (4.2.17)$$

Proof

Note that $\mathbb{E}(X) \in S$ so let $y = a + bx$ be a supporting line for g at $\mathbb{E}(X)$. Thus $a + b\mathbb{E}(X) = g[\mathbb{E}(X)]$ and $a + bX \leq g(X)$. Taking expected values through the inequality gives

$$a + b\mathbb{E}(X) = g[\mathbb{E}(X)] \leq \mathbb{E}[g(X)] \quad (4.2.18)$$

Jensens's inequality extends easily to higher dimensions. The 2-dimensional version is particularly important, because it will be used to derive several special inequalities in the section on vector spaces of random variables. We need two definitions.

A set $S \subseteq \mathbb{R}^n$ is *convex* if for every pair of points in S , the line segment connecting those points also lies in S . That is, if $\mathbf{x}, \mathbf{y} \in S$ and $p \in [0, 1]$ then $p\mathbf{x} + (1-p)\mathbf{y} \in S$.




Figure 4.2.2: A convex subset of \mathbb{R}^2

Suppose that $S \subseteq \mathbb{R}^n$ is convex. A function $g: S \rightarrow \mathbb{R}$ on S is *convex* (or *concave upward*) if for each $\mathbf{t} \in S$, there exist $\mathbf{a} \in \mathbb{R}$ and $\mathbf{b} \in \mathbb{R}^n$ (depending on \mathbf{t}) such that

1. $\mathbf{a} + \mathbf{b} \cdot \mathbf{t} = g(\mathbf{t})$
2. $\mathbf{a} + \mathbf{b} \cdot \mathbf{x} \leq g(\mathbf{x})$ for all $\mathbf{x} \in S$

The graph of $\mathbf{x} \mapsto \mathbf{a} + \mathbf{b} \cdot \mathbf{x}$ is called a *supporting hyperplane* for g at \mathbf{t} .

In \mathbb{R}^2 a supporting hyperplane is an ordinary plane. From calculus, if g has continuous second derivatives on S and has a *positive non-definite* second derivative matrix, then g is convex on S . Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ takes values in $S \subseteq \mathbb{R}^n$, and let $\mathbb{E}(\mathbf{X}) = (\mathbb{E}(X_1), \mathbb{E}(X_2), \dots, \mathbb{E}(X_n))$. The following result is the general version of Jensen's inequality.

If S is convex and $g: S \rightarrow \mathbb{R}$ is convex on S then

$$\mathbb{E}[g(\mathbf{X})] \geq g(\mathbb{E}(\mathbf{X})) \quad (4.2.19)$$

Proof

First $\mathbb{E}(\mathbf{X}) \in S$, so let $\mathbf{y} = \mathbf{a} + \mathbf{b} \cdot \mathbf{x}$ be a supporting hyperplane for g at $\mathbb{E}(\mathbf{X})$. Thus $\mathbf{a} + \mathbf{b} \cdot \mathbb{E}(\mathbf{X}) = g(\mathbb{E}(\mathbf{X}))$ and $\mathbf{a} + \mathbf{b} \cdot \mathbf{X} \leq g(\mathbf{X})$. Taking expected values through the inequality gives

$$\mathbf{a} + \mathbf{b} \cdot \mathbb{E}(\mathbf{X}) = g(\mathbb{E}(\mathbf{X})) \leq \mathbb{E}[g(\mathbf{X})] \quad (4.2.20)$$

We will study the expected value of random vectors and matrices in more detail in a later section. In both the one and n -dimensional cases, a function $g: S \rightarrow \mathbb{R}$ is *concave* (or *concave downward*) if the inequality in the definition is reversed. Jensen's inequality also reverses.

Expected Value in Terms of the Quantile Function

If X has a continuous distribution with support on an interval of \mathbb{R} , then there is a simple (but not well known) formula for the expected value of X as the integral the quantile function of X . Here is the general result:

Suppose that X has a continuous distribution with support on an interval $(a, b) \subseteq \mathbb{R}$. Let F denote the cumulative distribution function of X so that F^{-1} is the quantile function of X . If $g: (a, b) \rightarrow \mathbb{R}$ then (assuming that the expected value exists),

$$\mathbb{E}[g(X)] = \int_0^1 g[F^{-1}(p)] dp, \quad n \in \mathbb{N} \quad (4.2.21)$$

Proof

Suppose that X has probability density function f , although the theorem is true without this assumption. Under the assumption that X has a continuous distribution with support on the interval (a, b) , the distribution function F is strictly increasing on (a, b) , and the quantile function F^{-1} is the ordinary inverse of F . Substituting $p = F(x)$, $dp = F'(x) dx = f(x) dx$ we have

$$\int_0^1 g[F^{-1}(p)] dp = \int_a^b g(F^{-1}[F(x)]) f(x) dx = \int_a^b g(x) f(x) dx = \mathbb{E}[g(X)] \quad (4.2.22)$$

So in particular, $\mathbb{E}(X) = \int_0^1 F^{-1}(p) dp$.

Examples and Applications

Let $a \in (0, \infty)$ and let $\mathbb{P}(X = a) = 1$, so that X is a constant random variable. Show that Markov's inequality is in fact equality at $x = a$.

Solution

Of course $\mathbb{E}(X) = a$. Hence $\mathbb{P}(X \geq a) = 1$ and $\mathbb{E}(X)/a = 1$.

The Exponential Distribution

Recall that the *exponential distribution* is a continuous distribution with probability density function f given by

$$f(t) = re^{-rt}, \quad t \in [0, \infty) \quad (4.2.23)$$

where $r \in (0, \infty)$ is the *rate parameter*. This distribution is widely used to model failure times and other “arrival times”; in particular, the distribution governs the time between arrivals in the Poisson model. The exponential distribution is studied in detail in the chapter on the Poisson Process.

Suppose that X has exponential distribution with rate parameter r .

1. Find $\mathbb{E}(X)$ using the [right distribution formula](#).
2. Find $\mathbb{E}(X)$ using the [quantile function formula](#).
3. Compute both sides of [Markov's inequality](#).

Answer

1. $\int_0^\infty e^{-rt} dt = \frac{1}{r}$
2. $\int_0^1 -\frac{1}{r} \ln(1-p) dp = \frac{1}{r}$
3. $e^{-rt} < \frac{1}{rt}$ for $t > 0$

Open the gamma experiment. Keep the default value of the stopping parameter ($n = 1$), which gives the exponential distribution. Vary the rate parameter r and note the shape of the probability density function and the location of the mean. For various values of the rate parameter, run the experiment 1000 times and compare the sample mean with the distribution mean.

The Geometric Distribution

Recall that *Bernoulli trials* are independent trials each with two outcomes, which in the language of reliability, are called *success* and *failure*. The probability of success on each trial is $p \in [0, 1]$. A separate chapter on Bernoulli Trials explores this random process in more detail. It is named for Jacob Bernoulli. If $p \in (0, 1)$, the trial number N of the first success has the geometric distribution on \mathbb{N}_+ with success parameter p . The probability density function f of N is given by

$$f(n) = p(1-p)^{n-1}, \quad n \in \mathbb{N}_+ \quad (4.2.24)$$

Suppose that N has the geometric distribution on \mathbb{N}_+ with parameter $p \in (0, 1)$.

1. Find $\mathbb{E}(N)$ using the [right distribution function formula](#).
2. Compute both sides of [Markov's inequality](#).
3. Find $\mathbb{E}(N \mid N \text{ is even})$.

Answer

1. $\sum_{n=0}^\infty (1-p)^n = \frac{1}{p}$
2. $(1-p)^{n-1} < \frac{1}{np}, \quad n \in \mathbb{N}_+$
3. $\frac{2(1-p)^2}{p(2-p)^2}$

Open the negative binomial experiment. Keep the default value of the stopping parameter ($k = 1$), which gives the geometric distribution. Vary the success parameter p and note the shape of the probability density function and the location of the mean. For various values of the success parameter, run the experiment 1000 times and compare the sample mean with the distribution mean.

The Pareto Distribution

Recall that the *Pareto distribution* is a continuous distribution with probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (4.2.25)$$

where $a \in (0, \infty)$ is a parameter. The Pareto distribution is named for Vilfredo Pareto. It is a heavy-tailed distribution that is widely used to model certain financial variables. The Pareto distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with parameter $a > 1$.

1. Find $\mathbb{E}(X)$ using the [right distribution function formula](#).
2. Find $\mathbb{E}(X)$ using the [quantile function formula](#).
3. Find $\mathbb{E}(1/X)$.

4. Show that $x \mapsto 1/x$ is convex on $(0, \infty)$.
5. Verify [Jensen's inequality](#) by comparing $\mathbb{E}(1/X)$ and $1/\mathbb{E}(X)$.

Answer

1. $\int_0^1 1 \, dx + \int_1^\infty x^{-a} \, dx = \frac{a}{a-1}$
2. $\int_0^1 (1-p)^{-1/a} \, dp = \frac{a}{a-1}$
3. $\frac{a}{a+1}$
4. The convexity of $1/x$ is clear from the graph. Note also that $\frac{d^2}{dx^2} \frac{1}{x} = \frac{2}{x^3} > 0$ for $x > 0$.
5. $\frac{a}{a+1} > \frac{a-1}{a}$

Open the special distribution simulator and select the Pareto distribution. Keep the default value of the scale parameter. Vary the shape parameter and note the shape of the probability density function and the location of the mean. For various values of the shape parameter, run the experiment 1000 times and compare the sample mean with the distribution mean.

A Bivariate Distribution

Suppose that (X, Y) has probability density function f given by $f(x, y) = 2(x + y)$ for $0 \leq x \leq y \leq 1$.

1. Show that the domain of f is a convex set.
2. Show that $(x, y) \mapsto x^2 + y^2$ is convex on the domain of f .
3. Compute $\mathbb{E}(X^2 + Y^2)$.
4. Compute $[\mathbb{E}(X)]^2 + [\mathbb{E}(Y)]^2$.
5. Verify [Jensen's inequality](#) by comparing (b) and (c).

Answer

1. Note that the domain is a triangular region.
2. The second derivative matrix is $\begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$.
3. $\frac{5}{6}$
4. $\frac{53}{72}$
5. $\frac{5}{6} > \frac{53}{72}$

The Arithmetic and Geometric Means

Suppose that $\{x_1, x_2, \dots, x_n\}$ is a set of positive numbers. The *arithmetic mean* is at least as large as the *geometric mean*:

$$\left(\prod_{i=1}^n x_i \right)^{1/n} \leq \frac{1}{n} \sum_{i=1}^n x_i \quad (4.2.26)$$

Proof

Let X be uniformly distributed on $\{x_1, x_2, \dots, x_n\}$. We apply [Jensen's inequality](#) with the natural logarithm function, which is concave on $(0, \infty)$:

$$\mathbb{E}(\ln X) = \frac{1}{n} \sum_{i=1}^n \ln x_i = \ln \left[\left(\prod_{i=1}^n x_i \right)^{1/n} \right] \leq \ln[\mathbb{E}(X)] = \ln \left(\frac{1}{n} \sum_{i=1}^n x_i \right) \quad (4.2.27)$$

Taking exponentials of each side gives the inequality.

This page titled [4.2: Additional Properties](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.3: Variance

Recall the **expected value** of a real-valued random variable is the mean of the variable, and is a measure of the center of the distribution. Recall also that by taking the expected value of various transformations of the variable, we can measure other interesting characteristics of the distribution. In this section, we will study expected values that measure the spread of the distribution about the mean.

Basic Theory

Definitions and Interpretations

As usual, we start with a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} the collection of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . Suppose that X is a random variable for the experiment, taking values in $S \subseteq \mathbb{R}$. Recall that $\mathbb{E}(X)$, the expected value (or mean) of X gives the center of the distribution of X .

The *variance* and *standard deviation* of X are defined by

1. $\text{var}(X) = \mathbb{E}([X - \mathbb{E}(X)]^2)$
2. $\text{sd}(X) = \sqrt{\text{var}(X)}$

Implicit in the definition is the assumption that the mean $\mathbb{E}(X)$ exists, as a real number. If this is not the case, then $\text{var}(X)$ (and hence also $\text{sd}(X)$) are undefined. Even if $\mathbb{E}(X)$ does exist as a real number, it's possible that $\text{var}(X) = \infty$. For the remainder of our discussion of the basic theory, we will assume that expected values that are mentioned exist as real numbers.

The variance and standard deviation of X are both measures of the spread of the distribution about the mean. Variance (as we will see) has nicer mathematical properties, but its physical unit is the square of that of X . Standard deviation, on the other hand, is not as nice mathematically, but has the advantage that its physical unit is the same as that of X . When the random variable X is understood, the standard deviation is often denoted by σ , so that the variance is σ^2 .

Recall that the *second moment* of X about $a \in \mathbb{R}$ is $\mathbb{E}[(X - a)^2]$. Thus, the variance is the second moment of X about the mean $\mu = \mathbb{E}(X)$, or equivalently, the *second central moment* of X . In general, the second moment of X about $a \in \mathbb{R}$ can also be thought of as the *mean square error* if the constant a is used as an *estimate* of X . In addition, second moments have a nice interpretation in physics. If we think of the distribution of X as a mass distribution in \mathbb{R} , then the second moment of X about $a \in \mathbb{R}$ is the *moment of inertia* of the mass distribution about a . This is a measure of the resistance of the mass distribution to any change in its rotational motion about a . In particular, the variance of X is the moment of inertia of the mass distribution about the center of mass μ .



Figure 4.3.1: The moment of inertia about a .

The mean square error (or equivalently the moment of inertia) about a is minimized when $a = \mu$:

Let $\text{mse}(a) = \mathbb{E}[(X - a)^2]$ for $a \in \mathbb{R}$. Then mse is minimized when $a = \mu$, and the minimum value is σ^2 .

Proof




Figure 4.3.2: $\text{mse}(a)$ is minimized when $a = \mu$.

The relationship between measures of center and measures of spread is studied in more detail in the advanced section on vector spaces of random variables.

Properties

The following exercises give some basic properties of variance, which in turn rely on basic properties of expected value. As usual, be sure to try the proofs yourself before reading the ones in the text. Our first results are computational formulas based on the change of variables formula for expected value

Let $\mu = \mathbb{E}(X)$.

1. If X has a discrete distribution with probability density function f , then $\text{var}(X) = \sum_{x \in S} (x - \mu)^2 f(x)$.
2. If X has a continuous distribution with probability density function f , then $\text{var}(X) = \int_S (x - \mu)^2 f(x) dx$

Proof

1. This follows from the discrete version of the change of variables formula.
2. Similarly, this follows from the continuous version of the change of variables formula.

Our next result is a variance formula that is usually better than the definition for computational purposes.

$$\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2.$$

Proof

Let $\mu = \mathbb{E}(X)$. Using the linearity of expected value we have

$$\text{var}(X) = \mathbb{E}[(X - \mu)^2] = \mathbb{E}(X^2 - 2\mu X + \mu^2) = \mathbb{E}(X^2) - 2\mu \mathbb{E}(X) + \mu^2 = \mathbb{E}(X^2) - 2\mu^2 + \mu^2 = \mathbb{E}(X^2) - \mu^2 \quad (4.3.1)$$

Of course, by the change of variables formula, $\mathbb{E}(X^2) = \sum_{x \in S} x^2 f(x)$ if X has a discrete distribution, and $\mathbb{E}(X^2) = \int_S x^2 f(x) dx$ if X has a continuous distribution. In both cases, f is the probability density function of X .

Variance is always nonnegative, since it's the expected value of a nonnegative random variable. Moreover, any random variable that really is *random* (not a constant) will have strictly positive variance.

The nonnegative property.

1. $\text{var}(X) \geq 0$
2. $\text{var}(X) = 0$ if and only if $\mathbb{P}(X = c) = 1$ for some constant c (and then of course, $\mathbb{E}(X) = c$).

Proof

These results follow from the basic positive property of expected value. Let $\mu = \mathbb{E}(X)$. First $(X - \mu)^2 \geq 0$ with probability 1 so $\mathbb{E}[(X - \mu)^2] \geq 0$. In addition, $\mathbb{E}[(X - \mu)^2] = 0$ if and only if $\mathbb{P}(X = \mu) = 1$.

Our next result shows how the variance and standard deviation are changed by a linear transformation of the random variable. In particular, note that variance, unlike general expected value, is not a linear operation. This is not really surprising since the variance is the expected value of a nonlinear function of the variable: $x \mapsto (x - \mu)^2$.

If $a, b \in \mathbb{R}$ then

1. $\text{var}(a + bX) = b^2 \text{var}(X)$
2. $\text{sd}(a + bX) = |b| \text{sd}(X)$

Proof

1. Let $\mu = \mathbb{E}(X)$. By linearity, $\mathbb{E}(a + bX) = a + b\mu$. Hence $\text{var}(a + bX) = \mathbb{E}([(a + bX) - (a + b\mu)]^2) = \mathbb{E}[b^2(X - \mu)^2] = b^2 \text{var}(X)$.
2. This result follows from (a) by taking square roots.

Recall that when $b > 0$, the linear transformation $x \mapsto a + bx$ is called a *location-scale* transformation and often corresponds to a change of location and change of scale in the physical units. For example, the change from inches to centimeters in a measurement of length is a scale transformation, and the change from Fahrenheit to Celsius in a measurement of temperature is both a location and scale transformation. The previous result shows that when a location-scale transformation is applied to a random variable, the standard deviation does not depend on the location parameter, but is multiplied by the scale factor. There is a particularly important location-scale transformation.

Suppose that X is a random variable with mean μ and variance σ^2 . The random variable Z defined as follows is the *standard score* of X .

$$Z = \frac{X - \mu}{\sigma} \quad (4.3.2)$$

1. $\mathbb{E}(Z) = 0$
2. $\text{var}(Z) = 1$

Proof

1. From the linearity of expected value, $\mathbb{E}(Z) = \frac{1}{\sigma} [\mathbb{E}(X) - \mu] = 0$
2. From the [scaling property](#), $\text{var}(Z) = \frac{1}{\sigma^2} \text{var}(X) = 1$.

Since X and its mean and standard deviation all have the same physical units, the standard score Z is dimensionless. It measures the directed distance from $\mathbb{E}(X)$ to X in terms of standard deviations.

Let Z denote the standard score of X , and suppose that $Y = a + bX$ where $a, b \in \mathbb{R}$ and $b \neq 0$.

1. If $b > 0$, the standard score of Y is Z .
2. If $b < 0$, the standard score of Y is $-Z$.

Proof

$E(Y) = a + bE(X)$ and $sd(Y) = |b| sd(X)$. Hence

$$\frac{Y - E(Y)}{sd(Y)} = \frac{b}{|b|} \frac{X - E(X)}{sd(X)} \quad (4.3.3)$$

As just noted, when $b > 0$, the variable $Y = a + bX$ is a location-scale transformation and often corresponds to a change of physical units. Since the standard score is dimensionless, it's reasonable that the standard scores of X and Y are the same. Here is another standardized measure of dispersion:

Suppose that X is a random variable with $E(X) \neq 0$. The *coefficient of variation* is the ratio of the standard deviation to the mean:

$$cv(X) = \frac{sd(X)}{E(X)} \quad (4.3.4)$$

The coefficient of variation is also dimensionless, and is sometimes used to compare variability for random variables with different means. We will learn how to compute the variance of the sum of two random variables in the section on covariance.

Chebyshev's Inequality

Chebyshev's inequality (named after Pafnuty Chebyshev) gives an upper bound on the probability that a random variable will be more than a specified distance from its mean. This is often useful in applied problems where the distribution is unknown, but the mean and variance are known (at least approximately). In the following two results, suppose that X is a real-valued random variable with mean $\mu = E(X) \in \mathbb{R}$ and standard deviation $\sigma = sd(X) \in (0, \infty)$.

Chebyshev's inequality 1.

$$\mathbb{P}(|X - \mu| \geq t) \leq \frac{\sigma^2}{t^2}, \quad t > 0 \quad (4.3.5)$$

Proof




Figure 4.3.3: Chebyshev's inequality

Here's an alternate version, with the distance in terms of standard deviation.

Chebyshev's inequality 2.

$$\mathbb{P}(|X - \mu| \geq k\sigma) \leq \frac{1}{k^2}, \quad k > 0 \quad (4.3.6)$$

Proof

Let $t = k\sigma$ in the first version of [Chebyshev's inequality](#).

The usefulness of the Chebyshev inequality comes from the fact that it holds for any distribution (assuming only that the mean and variance exist). The tradeoff is that for many specific distributions, the Chebyshev bound is rather crude. Note in particular that the [first inequality](#) is useless when $t \leq \sigma$, and the [second inequality](#) is useless when $k \leq 1$, since 1 is an upper bound for the probability of any event. On the other hand, it's easy to construct a distribution for which [Chebyshev's inequality](#) is sharp for a specified value of $t \in (0, \infty)$. Such a distribution is given in an [exercise below](#).

Examples and Applications

As always, be sure to try the problems yourself before looking at the solutions and answers.

Indicator Variables

Suppose that X is an indicator variable with $p = \mathbb{P}(X = 1)$, where $p \in [0, 1]$. Then

1. $\mathbb{E}(X) = p$
2. $\text{var}(X) = p(1 - p)$

Proof

1. We proved this in the section on basic properties, although the result is so simple that we can do it again: $\mathbb{E}(X) = 1 \cdot p + 0 \cdot (1 - p) = p$.
2. Note that $X^2 = X$ since X only takes values 0 and 1. Hence $\mathbb{E}(X^2) = p$ and therefore $\text{var}(X) = p - p^2 = p(1 - p)$.

The graph of $\text{var}(X)$ as a function of p is a parabola, opening downward, with roots at 0 and 1. Thus the minimum value of $\text{var}(X)$ is 0, and occurs when $p = 0$ and $p = 1$ (when X is deterministic, of course). The maximum value is $\frac{1}{4}$ and occurs when $p = \frac{1}{2}$.




Figure 4.3.4: The variance of an indicator variable as a function of p .

Uniform Distributions

Discrete uniform distributions are widely used in combinatorial probability, and model a point chosen *at random* from a finite set. The mean and variance have simple forms for the discrete uniform distribution on a set of evenly spaced points (sometimes referred to as a *discrete interval*):

Suppose that X has the discrete uniform distribution on $\{a, a + h, \dots, a + (n - 1)h\}$ where $a \in \mathbb{R}$, $h \in (0, \infty)$, and $n \in \mathbb{N}_+$. Let $b = a + (n - 1)h$, the right endpoint. Then

1. $\mathbb{E}(X) = \frac{1}{2}(a + b)$.
2. $\text{var}(X) = \frac{1}{12}(b - a)(b - a + 2h)$.

Proof

1. We proved this in the section on basic properties. Here it is again, using the formula for the sum of the first $n - 1$ positive integers:

$$\mathbb{E}(X) = \frac{1}{n} \sum_{i=0}^{n-1} (a + ih) = \frac{1}{n} \left(na + h \frac{(n-1)n}{2} \right) = a + \frac{(n-1)h}{2} = \frac{a+b}{2} \quad (4.3.7)$$

2. Note that

$$\mathbb{E}(X^2) = \frac{1}{n} \sum_{i=0}^{n-1} (a + ih)^2 = \frac{1}{n-1} \sum_{i=0}^{n-1} (a^2 + 2ahi + h^2 i^2) \quad (4.3.8)$$

Using the formulas for the sum of the first $n - 1$ positive integers, and the sum of the squares of the first $n - 1$ positive integers, we have

$$\mathbb{E}(X^2) = \frac{1}{n} \left[na^2 + 2ah \frac{(n-1)n}{2} + h^2 \frac{(n-1)n(2n-1)}{6} \right] \quad (4.3.9)$$

Using [computational formula](#) and simplifying gives the result.

Note that mean is simply the average of the endpoints, while the variance depends only on difference between the endpoints and the step size.

Open the special distribution simulator, and select the discrete uniform distribution. Vary the parameters and note the location and size of the mean \pm standard deviation bar in relation to the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Next, recall that the continuous uniform distribution on a bounded interval corresponds to selecting a point *at random* from the interval. Continuous uniform distributions arise in geometric probability and a variety of other applied problems.

Suppose that X has the continuous uniform distribution on the interval $[a, b]$ where $a, b \in \mathbb{R}$ with $a < b$. Then

1. $\mathbb{E}(X) = \frac{1}{2}(a + b)$
2. $\text{var}(X) = \frac{1}{12}(b - a)^2$

Proof

1. $\mathbb{E}(X) = \int_a^b x \frac{1}{b-a} dx = \frac{b^2 - a^2}{2(b-a)} = \frac{a+b}{2}$
2. $\mathbb{E}(X^2) = \int_a^b x^2 \frac{1}{b-a} dx = \frac{b^3 - a^3}{3(b-a)}$. The variance result then follows from (a), the [computational formula](#) and simple algebra.

Note that the mean is the midpoint of the interval and the variance depends only on the length of the interval. Compare this with the results in the discrete case.

Open the special distribution simulator, and select the continuous uniform distribution. This is the uniform distribution the interval $[a, a + w]$. Vary the parameters and note the location and size of the mean \pm standard deviation bar in relation to the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Dice

Recall that a *fair* die is one in which the faces are equally likely. In addition to fair dice, there are various types of crooked dice. Here are three:

- An *ace-six flat die* is a six-sided die in which faces 1 and 6 have probability $\frac{1}{4}$ each while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each.
- A *two-five flat die* is a six-sided die in which faces 2 and 5 have probability $\frac{1}{4}$ each while faces 1, 3, 4, and 6 have probability $\frac{1}{8}$ each.
- A *three-four flat die* is a six-sided die in which faces 3 and 4 have probability $\frac{1}{4}$ each while faces 1, 2, 5, and 6 have probability $\frac{1}{8}$ each.

A flat die, as the name suggests, is a die that is not a cube, but rather is shorter in one of the three directions. The particular probabilities that we use ($\frac{1}{4}$ and $\frac{1}{8}$) are fictitious, but the essential property of a flat die is that the opposite faces on the shorter axis have slightly larger probabilities than the other four faces. Flat dice are sometimes used by gamblers to cheat. In the following problems, you will compute the mean and variance for each of the various types of dice. Be sure to compare the results.

A standard, fair die is thrown and the score X is recorded. Sketch the graph of the probability density function and compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{35}{12}$

An ace-six flat die is thrown and the score X is recorded. Sketch the graph of the probability density function and compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{15}{4}$

A two-five flat die is thrown and the score X is recorded. Sketch the graph of the probability density function and compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{11}{4}$

A three-four flat die is thrown and the score X is recorded. Sketch the graph of the probability density function and compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{9}{4}$

In the dice experiment, select one die. For each of the following cases, note the location and size of the mean \pm standard deviation bar in relation to the probability density function. Run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

1. Fair die
2. Ace-six flat die
3. Two-five flat die
4. Three-four flat die

The Poisson Distribution

Recall that the *Poisson distribution* is a discrete distribution on \mathbb{N} with probability density function f given by

$$f(n) = e^{-a} \frac{a^n}{n!}, \quad n \in \mathbb{N} \quad (4.3.10)$$

where $a \in (0, \infty)$ is a parameter. The Poisson distribution is named after Simeon Poisson and is widely used to model the number of “random points” in a region of time or space; the parameter a is proportional to the size of the region. The Poisson distribution is studied in detail in the chapter on the Poisson Process.

Suppose that N has the Poisson distribution with parameter a . Then

1. $\mathbb{E}(N) = a$
2. $\text{var}(N) = a$

Proof

1. We did this computation in the previous section. Here it is again:

$$\mathbb{E}(N) = \sum_{n=0}^{\infty} n e^{-a} \frac{a^n}{n!} = e^{-a} \sum_{n=1}^{\infty} \frac{a^n}{(n-1)!} = e^{-a} a \sum_{n=1}^{\infty} \frac{a^{n-1}}{(n-1)!} = e^{-a} a e^a = a. \quad (4.3.11)$$

2. First we compute the *second factorial moment*:

$$\mathbb{E}[N(N-1)] = \sum_{n=1}^{\infty} n(n-1) e^{-a} \frac{a^n}{n!} = \sum_{n=2}^{\infty} e^{-a} \frac{a^n}{(n-2)!} = e^{-a} a^2 \sum_{n=2}^{\infty} \frac{a^{n-2}}{(n-2)!} = a^2 e^{-a} e^a = a^2 \quad (4.3.12)$$

Hence, $E(N^2) = \mathbb{E}[N(N-1)] + \mathbb{E}(N) = a^2 + a$ and so $\text{var}(N) = (a^2 + a) - a^2 = a$.

Thus, the parameter of the Poisson distribution is both the mean and the variance of the distribution.

In the Poisson experiment, the parameter is $a = rt$. Vary the parameter and note the size and location of the mean \pm standard deviation bar in relation to the probability density function. For selected values of the parameter, run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The Geometric Distribution

Recall that *Bernoulli trials* are independent trials each with two outcomes, which in the language of reliability, are called *success* and *failure*. The probability of success on each trial is $p \in [0, 1]$. A separate chapter on Bernoulli Trials explores this random process in more detail. It is named for Jacob Bernoulli. If $p \in (0, 1]$, the trial number N of the first success has the geometric distribution on \mathbb{N}_+ with success parameter p . The probability density function f of N is given by

$$f(n) = p(1-p)^{n-1}, \quad n \in \mathbb{N}_+ \quad (4.3.13)$$

Suppose that N has the geometric distribution on \mathbb{N}_+ with success parameter $p \in (0, 1]$. Then

1. $\mathbb{E}(N) = \frac{1}{p}$
2. $\text{var}(N) = \frac{1-p}{p^2}$

Proof

1. We proved this in the section on basic properties. Here it is again:

$$\mathbb{E}(N) = \sum_{n=1}^{\infty} n p (1-p)^{n-1} = -p \frac{d}{dp} \sum_{n=0}^{\infty} (1-p)^n = -p \frac{d}{dp} \frac{1}{p} = \frac{1}{p} \quad (4.3.14)$$

2. First we compute the *second factorial moment*:

$$\mathbb{E}[N(N-1)] = \sum_{n=2}^{\infty} n(n-1) (1-p)^{n-1} p = p(1-p) \frac{d^2}{dp^2} \sum_{n=0}^{\infty} (1-p)^n = p(1-p) \frac{d^2}{dp^2} \frac{1}{p} = p(1-p) \frac{2}{p^3} = \frac{2(1-p)}{p^2} \quad (4.3.15)$$

Hence $\mathbb{E}(N^2) = \mathbb{E}[N(N-1)] + \mathbb{E}(N) = 2/p^2 - 1/p$ and hence $\text{var}(X) = 2/p^2 - 1/p - 1/p^2 = 1/p^2 - 1/p$.

Note that the variance is 0 when $p = 1$, not surprising since X is deterministic in this case.

In the negative binomial experiment, set $k = 1$ to get the geometric distribution. Vary p with the scroll bar and note the size and location of the mean \pm standard deviation bar in relation to the probability density function. For selected values of p , run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Suppose that N has the geometric distribution with parameter $p = \frac{3}{4}$. Compute the true value and the Chebyshev bound for the probability that N is at least 2 standard deviations away from the mean.

Answer

1. $\frac{1}{16}$
2. $\frac{1}{4}$

The Exponential Distribution

Recall that the *exponential distribution* is a continuous distribution on $[0, \infty)$ with probability density function f given by

$$f(t) = re^{-rt}, \quad t \in [0, \infty) \quad (4.3.16)$$

where $r \in (0, \infty)$ is the *rate parameter*. This distribution is widely used to model failure times and other “arrival times”. The exponential distribution is studied in detail in the chapter on the Poisson Process.

Suppose that T has the exponential distribution with rate parameter r . Then

1. $\mathbb{E}(T) = \frac{1}{r}$.
2. $\text{var}(T) = \frac{1}{r^2}$.

Proof

1. We proved this in the section on basic properties. Here it is again, using integration by parts:

$$\mathbb{E}(T) = \int_0^\infty t r e^{-rt} dt = -t e^{-rt} \Big|_0^\infty + \int_0^\infty e^{-rt} dt = 0 - \frac{1}{r} e^{-rt} \Big|_0^\infty = \frac{1}{r} \quad (4.3.17)$$

2. Integrating by parts again and using (a), we have

$$\mathbb{E}(T^2) = \int_0^\infty t^2 r e^{-rt} dt = -t^2 e^{-rt} \Big|_0^\infty + \int_0^\infty 2t e^{-rt} dt = 0 + \frac{2}{r^2} \quad (4.3.18)$$

$$\text{Hence } \text{var}(T) = \frac{2}{r^2} - \frac{1}{r^2} = \frac{1}{r^2}$$

Thus, for the exponential distribution, the mean and standard deviation are the same.

In the gamma experiment, set $k = 1$ to get the exponential distribution. Vary r with the scroll bar and note the size and location of the mean \pm standard deviation bar in relation to the probability density function. For selected values of r , run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Suppose that X has the exponential distribution with rate parameter $r > 0$. Compute the true value and the Chebyshev bound for the probability that X is at least k standard deviations away from the mean.

Answer

1. $e^{-(k+1)}$
2. $\frac{1}{k^2}$

The Pareto Distribution

Recall that the *Pareto distribution* is a continuous distribution on $[1, \infty)$ with probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (4.3.19)$$

where $a \in (0, \infty)$ is a parameter. The Pareto distribution is named for Vilfredo Pareto. It is a heavy-tailed distribution that is widely used to model financial variables such as income. The Pareto distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with shape parameter a . Then

1. $\mathbb{E}(X) = \infty$ if $0 < a \leq 1$ and $\mathbb{E}(X) = \frac{a}{a-1}$ if $1 < a < \infty$

2. $\text{var}(X)$ is undefined if $0 < a \leq 1$, $\text{var}(X) = \infty$ if $1 < a \leq 2$, and $\text{var}(X) = \frac{a}{(a-1)^2(a-2)}$ if $2 < a < \infty$

Proof

1. We proved this in the section on basic properties. Here it is again:

$$\mathbb{E}(X) = \int_1^\infty x \frac{a}{x^{a+1}} dx = \int_1^\infty \frac{a}{x^a} dx = \frac{a}{-a+1} x^{-a+1} \Big|_1^\infty = \begin{cases} \infty, & 0 < a < 1 \\ \frac{a}{a-1}, & a > 1 \end{cases} \quad (4.3.20)$$

When $a = 1$, $\mathbb{E}(X) = \int_1^\infty \frac{1}{x} dx = \ln x \Big|_1^\infty = \infty$

2. If $0 < a \leq 1$ then $\mathbb{E}(X) = \infty$ and so $\text{var}(X)$ is undefined. On the other hand,

$$\mathbb{E}(X^2) = \int_1^\infty x^2 \frac{a}{x^{a+1}} dx = \int_1^\infty \frac{a}{x^{a-1}} dx = ax^{-a+2} \Big|_1^\infty = \begin{cases} \infty, & 0 < a < 2 \\ \frac{a}{a-2}, & a > 2 \end{cases} \quad (4.3.21)$$

When $a = 2$, $\mathbb{E}(X^2) = \int_1^\infty \frac{2}{x} dx = \infty$. Hence $\text{var}(X) = \infty$ if $1 < a \leq 2$ and $\text{var}(X) = \frac{a}{a-2} - \left(\frac{a}{a-1}\right)^2$ if $a > 2$.

In the special distribution simulator, select the Pareto distribution. Vary a with the scroll bar and note the size and location of the mean \pm standard deviation bar. For each of the following values of a , run the experiment 1000 times and note the behavior of the empirical mean and standard deviation.

1. $a = 1$
2. $a = 2$
3. $a = 3$

The Normal Distribution

Recall that the *standard normal distribution* is a continuous distribution on \mathbb{R} with probability density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (4.3.22)$$

Normal distributions are widely used to model physical measurements subject to small, random errors and are studied in detail in the chapter on Special Distributions.

Suppose that Z has the standard normal distribution. Then

1. $\mathbb{E}(Z) = 0$
2. $\text{var}(Z) = 1$

Proof

1. We proved this in the section on basic properties. Here it is again:

$$\mathbb{E}(Z) = \int_{-\infty}^\infty z \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} dz = -\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} \Big|_{-\infty}^\infty = 0 - 0 \quad (4.3.23)$$

2. From (a), $\text{var}(Z) = \mathbb{E}(Z^2) = \int_{-\infty}^\infty z^2 \phi(z) dz$. Integrate by parts with $u = z$ and $dv = z\phi(z) dz$. Thus, $du = dz$ and $v = -\phi(z)$. Hence

$$\text{var}(Z) = -z\phi(z) \Big|_{-\infty}^\infty + \int_{-\infty}^\infty \phi(z) dz = 0 + 1 \quad (4.3.24)$$

More generally, for $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, recall that the normal distribution with *location parameter* μ and *scale parameter* σ is a continuous distribution on \mathbb{R} with probability density function f given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (4.3.25)$$

Moreover, if Z has the standard normal distribution, then $X = \mu + \sigma Z$ has the normal distribution with location parameter μ and scale parameter σ . As the notation suggests, the location parameter is the mean of the distribution and the scale parameter is the standard deviation.

Suppose that X has the normal distribution with location parameter μ and scale parameter σ . Then

1. $\mathbb{E}(X) = \mu$
2. $\text{var}(X) = \sigma^2$

Proof

We could use the probability density function, of course, but it's much better to use the representation of X in terms of the standard normal variable Z , and use properties of expected value and variance.

1. $\mathbb{E}(X) = \mu + \sigma \mathbb{E}(Z) = \mu + 0 = \mu$
2. $\text{var}(X) = \sigma^2 \text{var}(Z) = \sigma^2 \cdot 1 = \sigma^2$.

So to summarize, if X has a normal distribution, then its standard score Z has the standard normal distribution.

In the special distribution simulator, select the normal distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar in relation to the probability density function. For selected parameter values, run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Beta Distributions

The distributions in this subsection belong to the family of *beta distributions*, which are widely used to model random proportions and probabilities. The beta distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has a beta distribution with probability density function f . In each case below, graph f below and compute the mean and variance.

1. $f(x) = 6x(1-x)$ for $x \in [0, 1]$
2. $f(x) = 12x^2(1-x)$ for $x \in [0, 1]$
3. $f(x) = 12x(1-x)^2$ for $x \in [0, 1]$

Answer

1. $\mathbb{E}(X) = \frac{1}{2}$, $\text{var}(X) = \frac{1}{20}$
2. $\mathbb{E}(X) = \frac{3}{5}$, $\text{var}(X) = \frac{1}{25}$
3. $\mathbb{E}(X) = \frac{2}{6}$, $\text{var}(X) = \frac{1}{25}$

In the special distribution simulator, select the beta distribution. The parameter values below give the distributions in the previous exercise. In each case, note the location and size of the mean \pm standard deviation bar. Run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

1. $a = 2$, $b = 2$
2. $a = 3$, $b = 2$
3. $a = 2$, $b = 3$

Suppose that a sphere has a random radius R with probability density function f given by $f(r) = 12r^2(1-r)$ for $r \in [0, 1]$. Find the mean and standard deviation of each of the following:

1. The circumference $C = 2\pi R$
2. The surface area $A = 4\pi R^2$
3. The volume $V = \frac{4}{3}\pi R^3$

Answer

1. $\frac{6}{5}\pi$, $\frac{2}{5}\pi$
2. $\frac{8}{5}\pi$, $\frac{2}{5}\sqrt{\frac{38}{7}}\pi$
3. $\frac{8}{21}\pi$, $\frac{8}{3}\sqrt{\frac{19}{1470}}\pi$

Suppose that X has probability density function f given by $f(x) = \frac{1}{\pi\sqrt{x(1-x)}}$ for $x \in (0, 1)$. Find

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

Answer

1. $\frac{1}{2}$
2. $\frac{1}{8}$

The particular beta distribution in the last exercise is also known as the (standard) *arcsine distribution*. It governs the last time that the Brownian motion process hits 0 during the time interval $[0, 1]$. The arcsine distribution is studied in more generality in the chapter on Special Distributions.

Open the Brownian motion experiment and select the last zero. Note the location and size of the mean \pm standard deviation bar in relation to the probability density function. Run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Suppose that the grades on a test are described by the random variable $Y = 100X$ where X has the beta distribution with probability density function f given by $f(x) = 12x(1-x)^2$ for $x \in [0, 1]$. The grades are generally low, so the teacher decides to “curve” the grades using the transformation $Z = 10\sqrt{Y} = 100\sqrt{X}$. Find the mean and standard deviation of each of the following variables:

1. X
2. Y
3. Z

Answer

1. $\mathbb{E}(X) = \frac{2}{5}$, $\text{sd}(X) = \frac{1}{5}$
2. $\mathbb{E}(Y) = 40$, $\text{sd}(Y) = 20$
3. $\mathbb{E}(Z) = 60.95$, $\text{sd}(Z) = 16.88$

Exercises on Basic Properties

Suppose that X is a real-valued random variable with $\mathbb{E}(X) = 5$ and $\text{var}(X) = 4$. Find each of the following:

1. $\text{var}(3X - 2)$
2. $\mathbb{E}(X^2)$

Answer

1. 36
2. 29

Suppose that X is a real-valued random variable with $\mathbb{E}(X) = 2$ and $\mathbb{E}[X(X-1)] = 8$. Find each of the following:

1. $\mathbb{E}(X^2)$
2. $\text{var}(X)$

Answer

1. 10
2. 6

The expected value $\mathbb{E}[X(X-1)]$ is an example of a *factorial moment*.

Suppose that X_1 and X_2 are independent, real-valued random variables with $\mathbb{E}(X_i) = \mu_i$ and $\text{var}(X_i) = \sigma_i^2$ for $i \in \{1, 2\}$. Then

1. $\mathbb{E}(X_1 X_2) = \mu_1 \mu_2$
2. $\text{var}(X_1 X_2) = \sigma_1^2 \sigma_2^2 + \sigma_1^2 \mu_2^2 + \sigma_2^2 \mu_1^2$

Proof

1. This is an important, basic result that was proved in the section on basic properties.
2. Since X_1^2 and X_2^2 are also independent, we have $\mathbb{E}(X_1^2 X_2^2) = \mathbb{E}(X_1^2) \mathbb{E}(X_2^2) = (\sigma_1^2 + \mu_1^2)(\sigma_2^2 + \mu_2^2)$. The result then follows from the [computational formula](#) and algebra.

Marilyn Vos Savant has an IQ of 228. Assuming that the distribution of IQ scores has mean 100 and standard deviation 15, find Marilyn's standard score.

Answer

$$z = 8.53$$

Fix $t \in (0, \infty)$. Suppose that X is the discrete random variable with probability density function defined by $\mathbb{P}(X = t) = \mathbb{P}(X = -t) = p$, $\mathbb{P}(X = 0) = 1 - 2p$, where $p \in (0, \frac{1}{2})$. Then equality holds in [Chebyshev's inequality](#) at t .

Proof

Note that $\mathbb{E}(X) = 0$ and $\text{var}(X) = \mathbb{E}(X^2) = 2pt^2$. So $\mathbb{P}(|X| \geq t) = 2p$ and $\sigma^2/t^2 = 2p$.

This page titled [4.3: Variance](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.4: Skewness and Kurtosis

As usual, our starting point is a random experiment, modeled by a probability space (Ω, \mathcal{F}, P) . So to review, Ω is the set of outcomes, \mathcal{F} the collection of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . Suppose that X is a real-valued random variable for the experiment. Recall that the mean of X is a measure of the center of the distribution of X . Furthermore, the variance of X is the second moment of X about the mean, and measures the spread of the distribution of X about the mean. The third and fourth moments of X about the mean also measure interesting (but more subtle) features of the distribution. The third moment measures *skewness*, the lack of symmetry, while the fourth moment measures *kurtosis*, roughly a measure of the fatness in the tails. The actual numerical measures of these characteristics are standardized to eliminate the physical units, by dividing by an appropriate power of the standard deviation. As usual, we assume that all expected values given below exist, and we will let $\mu = \mathbb{E}(X)$ and $\sigma^2 = \text{var}(X)$. We assume that $\sigma > 0$, so that the random variable is really random.

Basic Theory

Skewness

The *skewness* of X is the third moment of the standard score of X :

$$\text{skew}(X) = \mathbb{E} \left[\left(\frac{X - \mu}{\sigma} \right)^3 \right] \quad (4.4.1)$$

The distribution of X is said to be *positively skewed*, *negatively skewed* or *unskewed* depending on whether $\text{skew}(X)$ is positive, negative, or 0.

In the unimodal case, if the distribution is positively skewed then the probability density function has a long tail to the right, and if the distribution is negatively skewed then the probability density function has a long tail to the left. A symmetric distribution is unskewed.

Suppose that the distribution of X is symmetric about a . Then

1. $\mathbb{E}(X) = a$
2. $\text{skew}(X) = 0$.

Proof

By assumption, the distribution of $a - X$ is the same as the distribution of $X - a$. We proved part (a) in the section on properties of expected value. Thus, $\text{skew}(X) = \mathbb{E}[(X - a)^3] / \sigma^3$. But by symmetry and linearity, $\mathbb{E}[(X - a)^3] = \mathbb{E}[(a - X)^3] = -\mathbb{E}[(X - a)^3]$, so it follows that $\mathbb{E}[(X - a)^3] = 0$.

The converse is not true—a non-symmetric distribution can have skewness 0. Examples are given in Exercises (30) and (31) below.

$\text{skew}(X)$ can be expressed in terms of the first three moments of X .

$$\text{skew}(X) = \frac{\mathbb{E}(X^3) - 3\mu\mathbb{E}(X^2) + 2\mu^3}{\sigma^3} = \frac{\mathbb{E}(X^3) - 3\mu\sigma^2 - \mu^3}{\sigma^3} \quad (4.4.2)$$

Proof

Note that $(X - \mu)^3 = X^3 - 3X^2\mu + 3X\mu^2 - \mu^3$. From the linearity of expected value we have

$$\mathbb{E}[(X - \mu)^3] = \mathbb{E}(X^3) - 3\mu\mathbb{E}(X^2) + 3\mu^2\mathbb{E}(X) - \mu^3 = \mathbb{E}(X^3) - 3\mu\mathbb{E}(X^2) + 2\mu^3 \quad (4.4.3)$$

The second expression follows from substituting $\mathbb{E}(X^2) = \sigma^2 + \mu^2$.

Since skewness is defined in terms of an odd power of the standard score, it's invariant under a linear transformation with positive slope (a location-scale transformation of the distribution). On the other hand, if the slope is negative, skewness changes sign.

Suppose that $a \in \mathbb{R}$ and $b \in \mathbb{R} \setminus \{0\}$. Then

1. $\text{skew}(a + bX) = \text{skew}(X)$ if $b > 0$
2. $\text{skew}(a + bX) = -\text{skew}(X)$ if $b < 0$

Proof

Let $Z = (X - \mu)/\sigma$, the standard score of X . Recall from the section on variance that the standard score of $a + bX$ is Z if $b > 0$ and is $-Z$ if $b < 0$.

Recall that location-scale transformations often arise when physical units are changed, such as inches to centimeters, or degrees Fahrenheit to degrees Celsius.

Kurtosis

The *kurtosis* of X is the fourth moment of the standard score:

$$\text{kurt}(X) = \mathbb{E} \left[\left(\frac{X - \mu}{\sigma} \right)^4 \right] \quad (4.4.4)$$

Kurtosis comes from the Greek word for *bulging*. Kurtosis is always positive, since we have assumed that $\sigma > 0$ (the random variable really is random), and therefore $\mathbb{P}(X \neq \mu) > 0$. In the unimodal case, the probability density function of a distribution with large kurtosis has fatter tails, compared with the probability density function of a distribution with smaller kurtosis.

$\text{kurt}(X)$ can be expressed in terms of the first four moments of X .

$$\text{kurt}(X) = \frac{\mathbb{E}(X^4) - 4\mu\mathbb{E}(X^3) + 6\mu^2\mathbb{E}(X^2) - 3\mu^4}{\sigma^4} = \frac{\mathbb{E}(X^4) - 4\mu\mathbb{E}(X^3) + 6\mu^2\sigma^2 + 3\mu^4}{\sigma^4} \quad (4.4.5)$$

Proof

Note that $(X - \mu)^4 = X^4 - 4X^3\mu + 6X^2\mu^2 - 4X\mu^3 + \mu^4$. From linearity of expected value, we have

$$\mathbb{E}[(X - \mu)^4] = \mathbb{E}(X^4) - 4\mu\mathbb{E}(X^3) + 6\mu^2\mathbb{E}(X^2) - 4\mu^3\mathbb{E}(X) + \mu^4 = \mathbb{E}(X^4) - 4\mu\mathbb{E}(X^3) + 6\mu^2\mathbb{E}(X^2) - 3\mu^4 \quad (4.4.6)$$

The second expression follows from the substitution $\mathbb{E}(X^2) = \sigma^2 + \mu^2$.

Since kurtosis is defined in terms of an even power of the standard score, it's invariant under linear transformations.

Suppose that $a \in \mathbb{R}$ and $b \in \mathbb{R} \setminus \{0\}$. Then $\text{kurt}(a + bX) = \text{kurt}(X)$.

Proof

As before, let $Z = (X - \mu)/\sigma$ denote the standard score of X . Then the standard score of $a + bX$ is Z if $b > 0$ and is $-Z$ if $b < 0$.

We will show in [below](#) that the kurtosis of the standard normal distribution is 3. Using the standard normal distribution as a benchmark, the *excess kurtosis* of a random variable X is defined to be $\text{kurt}(X) - 3$. Some authors use the term *kurtosis* to mean what we have defined as *excess kurtosis*.

Computational Exercises

As always, be sure to try the exercises yourself before expanding the solutions and answers in the text.

Indicator Variables

Recall that an *indicator random variable* is one that just takes the values 0 and 1. Indicator variables are the building blocks of many counting random variables. The corresponding distribution is known as the *Bernoulli distribution*, named for Jacob Bernoulli.

Suppose that X is an indicator variable with $\mathbb{P}(X = 1) = p$ where $p \in (0, 1)$. Then

1. $\mathbb{E}(X) = p$
2. $\text{var}(X) = p(1 - p)$
3. $\text{skew}(X) = \frac{1 - 2p}{\sqrt{p(1 - p)}}$
4. $\text{kurt}(X) = \frac{1 - 3p + 3p^2}{p(1 - p)}$

Proof

Parts (a) and (b) have been derived before. All four parts follow easily from the fact that $X^n = X$ and hence $\mathbb{E}(X^n) = p$ for $n \in \mathbb{N}_+$.

Open the binomial coin experiment and set $n = 1$ to get an indicator variable. Vary p and note the change in the shape of the probability density function.

Dice

Recall that a *fair* die is one in which the faces are equally likely. In addition to fair dice, there are various types of crooked dice. Here are three:

- An *ace-six flat die* is a six-sided die in which faces 1 and 6 have probability $\frac{1}{4}$ each while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each.
- A *two-five flat die* is a six-sided die in which faces 2 and 5 have probability $\frac{1}{4}$ each while faces 1, 3, 4, and 6 have probability $\frac{1}{8}$ each.
- A *three-four flat die* is a six-sided die in which faces 3 and 4 have probability $\frac{1}{4}$ each while faces 1, 2, 5, and 6 have probability $\frac{1}{8}$ each.

A flat die, as the name suggests, is a die that is not a cube, but rather is shorter in one of the three directions. The particular probabilities that we use ($\frac{1}{4}$ and $\frac{1}{8}$) are fictitious, but the essential property of a flat die is that the opposite faces on the shorter axis have slightly larger probabilities than the other four faces. Flat dice are sometimes used by gamblers to cheat.

A standard, fair die is thrown and the score X is recorded. Compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{35}{12}$
3. 0
4. $\frac{303}{175}$

An ace-six flat die is thrown and the score X is recorded. Compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{15}{4}$
3. 0
4. $\frac{37}{25}$

A two-five flat die is thrown and the score X is recorded. Compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{11}{4}$
3. 0
4. $\frac{197}{121}$

A three-four flat die is thrown and the score X is recorded. Compute each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{7}{2}$
2. $\frac{9}{4}$
3. 0
4. $\frac{59}{27}$

All four die distributions above have the same mean $\frac{7}{2}$ and are symmetric (and hence have skewness 0), but differ in variance and kurtosis.

Open the dice experiment and set $n = 1$ to get a single die. Select each of the following, and note the shape of the probability density function in comparison with the computational results above. In each case, run the experiment 1000 times and compare the empirical density function to the probability density function.

1. fair
2. ace-six flat
3. two-five flat
4. three-four flat

Uniform Distributions

Recall that the continuous uniform distribution on a bounded interval corresponds to selecting a point *at random* from the interval. Continuous uniform distributions arise in geometric probability and a variety of other applied problems.

Suppose that X has uniform distribution on the interval $[a, b]$, where $a, b \in \mathbb{R}$ and $a < b$. Then

1. $\mathbb{E}(X) = \frac{1}{2}(a + b)$
2. $\text{var}(X) = \frac{1}{12}(b - a)^2$
3. $\text{skew}(X) = 0$
4. $\text{kurt}(X) = \frac{9}{5}$

Proof

Parts (a) and (b) we have seen before. For parts (c) and (d), recall that $X = a + (b - a)U$ where U has the uniform distribution on $[0, 1]$ (the *standard uniform distribution*). Hence it follows from the formulas for [skewness](#) and [kurtosis](#) under linear transformations that $\text{skew}(X) = \text{skew}(U)$ and $\text{kurt}(X) = \text{kurt}(U)$. Since $\mathbb{E}(U^n) = 1/(n + 1)$ for $n \in \mathbb{N}_+$, it's easy to compute the skewness and kurtosis of U from the computational formulas [skewness](#) and [kurtosis](#). Of course, the fact that $\text{skew}(X) = 0$ also follows trivially from the symmetry of the distribution of X about the mean.

Open the special distribution simulator, and select the continuous uniform distribution. Vary the parameters and note the shape of the probability density function in comparison with the moment results in the last exercise. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The Exponential Distribution

Recall that the *exponential distribution* is a continuous distribution on $[0, \infty)$ with probability density function f given by

$$f(t) = re^{-rt}, \quad t \in [0, \infty) \quad (4.4.7)$$

where $r \in (0, \infty)$ is the *rate parameter*. This distribution is widely used to model failure times and other “arrival times”. The exponential distribution is studied in detail in the chapter on the Poisson Process.

Suppose that X has the exponential distribution with rate parameter $r > 0$. Then

1. $\mathbb{E}(X) = \frac{1}{r}$
2. $\text{var}(X) = \frac{1}{r^2}$
3. $\text{skew}(X) = 2$
4. $\text{kurt}(X) = 9$

Proof

These results follow from the computational formulas for [skewness](#) and [kurtosis](#) and the general moment formula $\mathbb{E}(X^n) = n!/r^n$ for $n \in \mathbb{N}$.

Note that the skewness and kurtosis do not depend on the rate parameter r . That's because $1/r$ is a scale parameter for the exponential distribution

Open the gamma experiment and set $n = 1$ to get the exponential distribution. Vary the rate parameter and note the shape of the probability density function in comparison to the moment results in the last exercise. For selected values of the parameter, run the experiment 1000 times and compare the empirical density function to the true probability density function.

Pareto Distribution

Recall that the *Pareto distribution* is a continuous distribution on $[1, \infty)$ with probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (4.4.8)$$

where $a \in (0, \infty)$ is a parameter. The Pareto distribution is named for Vilfredo Pareto. It is a heavy-tailed distribution that is widely used to model financial variables such as income. The Pareto distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with shape parameter $a > 0$. Then

1. $\mathbb{E}(X) = \frac{a}{a-1}$ if $a > 1$
2. $\text{var}(X) = \frac{a}{(a-1)^2(a-2)}$ if $a > 2$
3. $\text{skew}(X) = \frac{2(1+a)}{a-3} \sqrt{1 - \frac{2}{a}}$ if $a > 3$
4. $\text{kurt}(X) = \frac{3(a-2)(3a^2+a+2)}{a(a-3)(a-4)}$ if $a > 4$

Proof

These results follow from the standard computational formulas for [skewness](#) and [kurtosis](#) and the general moment formula $\mathbb{E}(X^n) = \frac{a}{a-n}$ if $n \in \mathbb{N}$ and $n < a$.

Open the special distribution simulator and select the Pareto distribution. Vary the shape parameter and note the shape of the probability density function in comparison to the moment results in the last exercise. For selected values of the parameter, run the experiment 1000 times and compare the empirical density function to the true probability density function.

The Normal Distribution

Recall that the *standard normal distribution* is a continuous distribution on \mathbb{R} with probability density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (4.4.9)$$

Normal distributions are widely used to model physical measurements subject to small, random errors and are studied in detail in the chapter on Special Distributions.

Suppose that Z has the standard normal distribution. Then

1. $\mathbb{E}(Z) = 0$
2. $\text{var}(Z) = 1$
3. $\text{skew}(Z) = 0$
4. $\text{kurt}(Z) = 3$

Proof

Parts (a) and (b) were derived in the previous sections on expected value and variance. Part (c) follows from symmetry. For part (d), recall that $\mathbb{E}(Z^4) = 3\mathbb{E}(Z^2) = 3$.

More generally, for $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, recall that the normal distribution with mean μ and standard deviation σ is a continuous distribution on \mathbb{R} with probability density function f given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (4.4.10)$$

However, we also know that μ and σ are location and scale parameters, respectively. That is, if Z has the standard normal distribution then $X = \mu + \sigma Z$ has the normal distribution with mean μ and standard deviation σ .

If X has the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$, then

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = 3$

Proof

The results follow immediately from the formulas for [skewness](#) and [kurtosis](#) under linear transformations and the previous result.

Open the special distribution simulator and select the normal distribution. Vary the parameters and note the shape of the probability density function in comparison to the moment results in the last exercise. For selected values of the parameters, run the experiment 1000 times and compare the empirical density function to the true probability density function.

The Beta Distribution

The distributions in this subsection belong to the family of *beta distributions*, which are continuous distributions on $[0, 1]$ widely used to model random proportions and probabilities. The beta distribution is studied in detail in the chapter on Special Distributions.

Suppose that X has probability density function f given by $f(x) = 6x(1 - x)$ for $x \in [0, 1]$. Find each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{1}{2}$
2. $\frac{1}{20}$
3. 0
4. $\frac{15}{7}$

Suppose that X has probability density function f given by $f(x) = 12x^2(1 - x)$ for $x \in [0, 1]$. Find each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{3}{5}$
2. $\frac{1}{25}$
3. $-\frac{2}{7}$
4. $\frac{33}{14}$

Suppose that X has probability density function f given by $f(x) = 12x(1 - x)^2$ for $x \in [0, 1]$. Find each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{2}{5}$
2. $\frac{1}{25}$
3. $\frac{2}{7}$
4. $\frac{33}{14}$

Open the special distribution simulator and select the beta distribution. Select the parameter values below to get the distributions in the last three exercises. In each case, note the shape of the probability density function in relation to the calculated moment results. Run the simulation 1000 times and compare the empirical density function to the probability density function.

1. $a = 2, b = 2$
2. $a = 3, b = 2$
3. $a = 2, b = 3$

Suppose that X has probability density function f given by $f(x) = \frac{1}{\pi\sqrt{x(1-x)}}$ for $x \in (0, 1)$. Find

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. $\frac{1}{2}$
2. $\frac{1}{8}$
3. 0
4. 96

The particular beta distribution in the last exercise is also known as the (standard) *arcsine distribution*. It governs the last time that the Brownian motion process hits 0 during the time interval $[0, 1]$. The arcsine distribution is studied in more generality in the chapter on Special Distributions.

Open the Brownian motion experiment and select the last zero. Note the shape of the probability density function in relation to the moment results in the last exercise. Run the simulation 1000 times and compare the empirical density function to the probability density function.

Counterexamples

The following exercise gives a simple example of a discrete distribution that is not symmetric but has skewness 0.

Suppose that X is a discrete random variable with probability density function f given by $f(-3) = \frac{1}{10}, f(-1) = \frac{1}{2}, f(2) = \frac{2}{5}$. Find each of the following and then show that the distribution of X is not symmetric.

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Answer

1. 0
2. 3
3. 0
4. $\frac{5}{3}$

The PDF f is clearly not symmetric about 0, and the mean is the only possible point of symmetry.

The following exercise gives a more complicated continuous distribution that is not symmetric but has skewness 0. It is one of a collection of distributions constructed by [Erik Meijer](#).

Suppose that U , V , and I are independent random variables, and that U is normally distributed with mean $\mu = -2$ and variance $\sigma^2 = 1$, V is normally distributed with mean $\nu = 1$ and variance $\tau^2 = 2$, and I is an indicator variable with $\mathbb{P}(I = 1) = p = \frac{1}{3}$. Let $X = IU + (1 - I)V$. Find each of the following and then show that the distribution of X is not symmetric.

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\text{skew}(X)$
4. $\text{kurt}(X)$

Solution

The distribution of X is a *mixture* of normal distributions. The PDF is $f = pg + (1 - p)h$ where g is the normal PDF of U and h is the normal PDF of V . However, it's best to work with the random variables. For $n \in \mathbb{N}_+$, note that $I^n = I$ and $(1 - I)^n = 1 - I$ and note also that the random variable $I(1 - I)$ just takes the value 0. It follows that

$$X^n = IU^n + (1 - I)V^n, \quad n \in \mathbb{N}_+ \quad (4.4.11)$$

So now, using standard results for the normal distribution,

1. $\mathbb{E}(X) = p\mu + (1 - p)\nu = 0$.
2. $\text{var}(X) = \mathbb{E}(X^2) = p(\sigma^2 + \mu^2) + (1 - p)(\tau^2 + \nu^2) = \frac{11}{3}$
3. $\mathbb{E}(X^3) = p(3\mu\sigma^2 + \mu^3) + (1 - p)(3\nu\tau^2 + \nu^3) = 0$ so $\text{skew}(X) = 0$
4. $\mathbb{E}(X^4) = p(3\sigma^4 + 6\sigma^2\mu^2 + \mu^4) + (1 - p)(3\tau^4 + 6\tau^2\nu^2 + \nu^4) = 31$ so $\text{kurt}(X) = \frac{279}{121} \approx 2.306$

The graph of the PDF f of X is given below. Note that f is not symmetric about 0. (Again, the mean is the only possible point of symmetry.)

The PDF of X


This page titled [4.4: Skewness and Kurtosis](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.5: Covariance and Correlation

Recall that by taking the expected value of various transformations of a random variable, we can measure many interesting characteristics of the distribution of the variable. In this section, we will study an expected value that measures a special type of relationship between two real-valued variables. This relationship is very important both in probability and statistics.

Basic Theory

Definitions

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Unless otherwise noted, we assume that all expected values mentioned in this section exist. Suppose now that X and Y are real-valued random variables for the experiment (that is, defined on the probability space) with means $\mathbb{E}(X)$, $\mathbb{E}(Y)$ and variances $\text{var}(X)$, $\text{var}(Y)$, respectively.

The *covariance* of (X, Y) is defined by

$$\text{cov}(X, Y) = \mathbb{E}([X - \mathbb{E}(X)][Y - \mathbb{E}(Y)]) \quad (4.5.1)$$

and, assuming the variances are positive, the *correlation* of (X, Y) is defined by

$$\text{cor}(X, Y) = \frac{\text{cov}(X, Y)}{\text{sd}(X)\text{sd}(Y)} \quad (4.5.2)$$

1. If $\text{cov}(X, Y) > 0$ then X and Y are *positively correlated*.
2. If $\text{cov}(X, Y) < 0$ then X and Y are *negatively correlated*.
3. If $\text{cov}(X, Y) = 0$ then X and Y are *uncorrelated*.

Correlation is a scaled version of covariance; note that the two parameters always have the same sign (positive, negative, or 0). Note also that correlation is dimensionless, since the numerator and denominator have the same physical units, namely the product of the units of X and Y .

As these terms suggest, covariance and correlation measure a certain kind of dependence between the variables. One of our goals is a deeper understanding of this dependence. As a start, note that $(\mathbb{E}(X), \mathbb{E}(Y))$ is the center of the joint distribution of (X, Y) , and the vertical and horizontal lines through this point separate \mathbb{R}^2 into four quadrants. The function $(x, y) \mapsto [x - \mathbb{E}(X)][y - \mathbb{E}(Y)]$ is positive on the first and third quadrants and negative on the second and fourth.




Figure 4.5.1: A joint distribution with $(\mathbb{E}(X), \mathbb{E}(Y))$ as the center of mass

Properties of Covariance

The following theorems give some basic properties of covariance. The main tool that we will need is the fact that expected value is a linear operation. Other important properties will be derived below, in the subsection on the [best linear predictor](#). As usual, be sure to try the proofs yourself before reading the ones in the text. Once again, we assume that the random variables are defined on the common sample space, are real-valued, and that the indicated expected values exist (as real numbers).

Our first result is a formula that is better than the definition for computational purposes, but gives less insight.

$$\text{cov}(X, Y) = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y).$$

Proof

Let $\mu = \mathbb{E}(X)$ and $\nu = \mathbb{E}(Y)$. Then

$$\text{cov}(X, Y) = \mathbb{E}[(X - \mu)(Y - \nu)] = \mathbb{E}(XY - \mu Y - \nu X + \mu\nu) = \mathbb{E}(XY) - \mu\mathbb{E}(Y) - \nu\mathbb{E}(X) + \mu\nu = \mathbb{E}(XY) - \mu\nu \quad (4.5.3)$$

From (2), we see that X and Y are uncorrelated if and only if $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y)$, so here is a simple but important corollary:

If X and Y are independent, then they are uncorrelated.

Proof

We showed in Section 1 that if X and Y are independent then $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y)$.

However, the converse fails with a passion: Exercise (31) gives an example of two variables that are *functionally related* (the strongest form of dependence), yet uncorrelated. The computational exercises give other examples of dependent yet uncorrelated variables also. Note also that if one of the variables has mean 0, then the covariance is simply the expected product.

Trivially, covariance is a symmetric operation.

$$\text{cov}(X, Y) = \text{cov}(Y, X).$$

As the name suggests, covariance generalizes variance.

$$\text{cov}(X, X) = \text{var}(X).$$

Proof

Let $\mu = \mathbb{E}(X)$. Then $\text{cov}(X, X) = \mathbb{E}[(X - \mu)^2] = \text{var}(X)$.

Covariance is a linear operation in the first argument, if the second argument is fixed.

If X, Y, Z are random variables, and c is a constant, then

1. $\text{cov}(X + Y, Z) = \text{cov}(X, Z) + \text{cov}(Y, Z)$
2. $\text{cov}(cX, Y) = c \text{cov}(X, Y)$

Proof

We use the computational formula in (2)

$$1. \quad \text{cov}(X + Y, Z) = \mathbb{E}[(X + Y)Z] - \mathbb{E}(X + Y)\mathbb{E}(Z) = \mathbb{E}(XZ + YZ) - [\mathbb{E}(X) + \mathbb{E}(Y)]\mathbb{E}(Z) \quad (4.5.4)$$

$$= [\mathbb{E}(XZ) - \mathbb{E}(X)\mathbb{E}(Z)] + [\mathbb{E}(YZ) - \mathbb{E}(Y)\mathbb{E}(Z)] = \text{cov}(X, Z) + \text{cov}(Y, Z) \quad (4.5.5)$$

$$2. \quad \text{cov}(cX, Y) = \mathbb{E}(cXY) - \mathbb{E}(cX)\mathbb{E}(Y) = c\mathbb{E}(XY) - c\mathbb{E}(X)\mathbb{E}(Y) = c[\mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y)] = c \text{cov}(X, Y) \quad (4.5.6)$$

By symmetry, covariance is also a linear operation in the second argument, with the first argument fixed. Thus, the covariance operator is *bi-linear*. The general version of this property is given in the following theorem.

Suppose that (X_1, X_2, \dots, X_n) and (Y_1, Y_2, \dots, Y_m) are sequences of random variables, and that (a_1, a_2, \dots, a_n) and (b_1, b_2, \dots, b_m) are constants. Then

$$\text{cov}\left(\sum_{i=1}^n a_i X_i, \sum_{j=1}^m b_j Y_j\right) = \sum_{i=1}^n \sum_{j=1}^m a_i b_j \text{cov}(X_i, Y_j) \quad (4.5.7)$$

The following result shows how covariance is changed under a linear transformation of one of the variables. This is simply a special case of the basic properties, but is worth stating.

If $a, b \in \mathbb{R}$ then $\text{cov}(a + bX, Y) = b \text{cov}(X, Y)$.

Proof

A constant is independent of any random variable. Hence $\text{cov}(a + bX, Y) = \text{cov}(a, Y) + b \text{cov}(X, Y) = b \text{cov}(X, Y)$.

Of course, by symmetry, the same property holds in the second argument. Putting the two together we have that if $a, b, c, d \in \mathbb{R}$ then $\text{cov}(a + bX, c + dY) = bd \text{cov}(X, Y)$.

Properties of Correlation

Next we will establish some basic properties of correlation. Most of these follow easily from corresponding properties of covariance above. We assume that $\text{var}(X) > 0$ and $\text{var}(Y) > 0$, so that the random variable really are random and hence the correlation is well defined.

The correlation between X and Y is the covariance of the corresponding standard scores:

$$\text{cor}(X, Y) = \text{cov}\left(\frac{X - \mathbb{E}(X)}{\text{sd}(X)}, \frac{Y - \mathbb{E}(Y)}{\text{sd}(Y)}\right) = \mathbb{E}\left(\frac{X - \mathbb{E}(X)}{\text{sd}(X)} \frac{Y - \mathbb{E}(Y)}{\text{sd}(Y)}\right) \quad (4.5.8)$$

Proof

From the definitions and the linearity of expected value,

$$\text{cor}(X, Y) = \frac{\text{cov}(X, Y)}{\text{sd}(X)\text{sd}(Y)} = \frac{\mathbb{E}([X - \mathbb{E}(X)][Y - \mathbb{E}(Y)])}{\text{sd}(X)\text{sd}(Y)} = \mathbb{E}\left(\frac{X - \mathbb{E}(X)}{\text{sd}(X)} \frac{Y - \mathbb{E}(Y)}{\text{sd}(Y)}\right) \quad (4.5.9)$$

Since the standard scores have mean 0, this is also the covariance of the standard scores.

This shows again that correlation is dimensionless, since of course, the standard scores are dimensionless. Also, correlation is symmetric:

$$\text{cor}(X, Y) = \text{cor}(Y, X).$$

Under a linear transformation of one of the variables, the correlation is unchanged if the slope is positive and changes sign if the slope is negative:

If $a, b \in \mathbb{R}$ and $b \neq 0$ then

1. $\text{cor}(a + bX, Y) = \text{cor}(X, Y)$ if $b > 0$
2. $\text{cor}(a + bX, Y) = -\text{cor}(X, Y)$ if $b < 0$

Proof

Let Z denote the standard score of X . If $b > 0$, the standard score of $a + bX$ is also Z . If $b < 0$, the standard score of $a + bX$ is $-Z$. Hence the result follows from the [result above](#) for standard scores.

This result reinforces the fact that correlation is a standardized measure of association, since multiplying the variable by a positive constant is equivalent to a change of scale, and adding a constant to a variable is equivalent to a change of location. For example, in the Challenger data, the underlying variables are temperature at the time of launch (in degrees Fahrenheit) and O-ring erosion (in millimeters). The correlation between these two variables is of fundamental importance. If we decide to measure temperature in degrees Celsius and O-ring erosion in inches, the correlation is unchanged. Of course, the same property holds in the second argument, so if $a, b, c, d \in \mathbb{R}$ with $b \neq 0$ and $d \neq 0$, then $\text{cor}(a + bX, c + dY) = \text{cor}(X, Y)$ if $bd > 0$ and $\text{cor}(a + bX, c + dY) = -\text{cor}(X, Y)$ if $bd < 0$.

The most important properties of covariance and correlation will emerge from our study of the [best linear predictor](#) below.

The Variance of a Sum

We will now show that the variance of a sum of variables is the sum of the pairwise covariances. This result is very useful since many random variables with special distributions can be written as sums of simpler random variables (see in particular the [binomial distribution](#) and [hypergeometric distribution](#) below).

If (X_1, X_2, \dots, X_n) is a sequence of real-valued random variables then

$$\text{var}\left(\sum_{i=1}^n X_i\right) = \sum_{i=1}^n \sum_{j=1}^n \text{cov}(X_i, X_j) = \sum_{i=1}^n \text{var}(X_i) + 2 \sum_{\{(i,j): i < j\}} \text{cov}(X_i, X_j) \quad (4.5.10)$$

Proof

From the variance property on (5), and the linear property (7),

$$\text{var}\left(\sum_{i=1}^n X_i\right) = \text{cov}\left(\sum_{i=1}^n X_i, \sum_{j=1}^n X_j\right) = \sum_{i=1}^n \sum_{j=1}^n \text{cov}(X_i, X_j) \quad (4.5.11)$$

The second expression follows since $\text{cov}(X_i, X_i) = \text{var}(X_i)$ for each i and $\text{cov}(X_i, X_j) = \text{cov}(X_j, X_i)$ for $i \neq j$ by the symmetry property (4)

Note that the variance of a sum can be larger, smaller, or equal to the sum of the variances, depending on the pure covariance terms. As a special case of (12), when $n = 2$, we have

$$\text{var}(X + Y) = \text{var}(X) + \text{var}(Y) + 2 \text{cov}(X, Y) \quad (4.5.12)$$

The following corollary is very important.

If (X_1, X_2, \dots, X_n) is a sequence of pairwise uncorrelated, real-valued random variables then

$$\text{var}\left(\sum_{i=1}^n X_i\right) = \sum_{i=1}^n \text{var}(X_i) \quad (4.5.13)$$

Proof

This follows immediately from (12), since $\text{cov}(X_i, X_j) = 0$ for $i \neq j$.

Note that the last result holds, in particular, if the random variables are independent. We close this discussion with a couple of minor corollaries.

If X and Y are real-valued random variables then $\text{var}(X + Y) + \text{var}(X - Y) = 2[\text{var}(X) + \text{var}(Y)]$.

Proof

From (12),

$$\text{var}(X + Y) = \text{var}(X) + \text{var}(Y) + 2\text{cov}(X, Y) \quad (4.5.14)$$

Similarly,

$$\text{var}(X - Y) = \text{var}(X) + \text{var}(-Y) + 2\text{cov}(X, -Y) = \text{var}(X) + \text{var}(Y) - 2\text{cov}(X, Y) \quad (4.5.15)$$

Adding gives the result.

If X and Y are real-valued random variables with $\text{var}(X) = \text{var}(Y)$ then $X + Y$ and $X - Y$ are uncorrelated.

Proof

From the linear property (7) and the symmetry property (4),
 $\text{cov}(X + Y, X - Y) = \text{cov}(X, X) - \text{cov}(X, Y) + \text{cov}(Y, X) - \text{cov}(Y, Y) = \text{var}(X) - \text{var}(Y)$

Random Samples

In the following exercises, suppose that (X_1, X_2, \dots) is a sequence of independent, real-valued random variables with a common distribution that has mean μ and standard deviation $\sigma > 0$. In statistical terms, the variables form a random sample from the common distribution.

For $n \in \mathbb{N}_+$, let $Y_n = \sum_{i=1}^n X_i$.

1. $\mathbb{E}(Y_n) = n\mu$
2. $\text{var}(Y_n) = n\sigma^2$

Proof

1. This follows from the additive property of expected value.
2. This follows from the additive property of variance ((13) for independent variables)

For $n \in \mathbb{N}_+$, let $M_n = Y_n/n = \frac{1}{n} \sum_{i=1}^n X_i$, so that M_n is the sample mean of (X_1, X_2, \dots, X_n) .

1. $\mathbb{E}(M_n) = \mu$
2. $\text{var}(M_n) = \sigma^2/n$
3. $\text{var}(M_n) \rightarrow 0$ as $n \rightarrow \infty$
4. $\mathbb{P}(|M_n - \mu| > \epsilon) \rightarrow 0$ as $n \rightarrow \infty$ for every $\epsilon > 0$.

Proof

1. This follows from part (a) of the (16) and the scaling property of expected value.
2. This follows from part (b) of the (16) and the scaling property of variance.
3. This is an immediate consequence of (b).
4. This follows from (c) and Chebyshev's inequality: $\mathbb{P}(|M_n - \mu| > \epsilon) \leq \text{var}(M_n)/\epsilon^2 \rightarrow 0$ as $n \rightarrow \infty$

Part (c) of (17) means that $M_n \rightarrow \mu$ as $n \rightarrow \infty$ in mean square. Part (d) means that $M_n \rightarrow \mu$ as $n \rightarrow \infty$ in probability. These are both versions of the weak law of large numbers, one of the fundamental theorems of probability.

The standard score of the sum Y_n and the standard score of the sample mean M_n are the same:

$$Z_n = \frac{Y_n - n\mu}{\sqrt{n}\sigma} = \frac{M_n - \mu}{\sigma/\sqrt{n}} \quad (4.5.16)$$

1. $\mathbb{E}(Z_n) = 0$
2. $\text{var}(Z_n) = 1$

Proof

The equality of the standard score of Y_n and of Z_n is a result of simple algebra. But recall more generally that the standard score of a variable is unchanged by a linear transformation of the variable with positive slope (a location-scale transformation of the distribution). Of course, parts (a) and (b) are true for any standard score.

The central limit theorem, the other fundamental theorem of probability, states that the distribution of Z_n converges to the standard normal distribution as $n \rightarrow \infty$.

Events

If A and B are events in our random experiment then the covariance and correlation of A and B are defined to be the covariance and correlation, respectively, of their indicator random variables.

If A and B are events, define $\text{cov}(A, B) = \text{cov}(\mathbf{1}_A, \mathbf{1}_B)$ and $\text{cor}(A, B) = \text{cor}(\mathbf{1}_A, \mathbf{1}_B)$. Equivalently,

1. $\text{cov}(A, B) = \mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)$
2. $\text{cor}(A, B) = [\mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)] / \sqrt{\mathbb{P}(A)[1 - \mathbb{P}(A)] \mathbb{P}(B)[1 - \mathbb{P}(B)]}$

Proof

Recall that if X is an indicator variable with $\mathbb{P}(X = 1) = p$, then $\mathbb{E}(X) = p$ and $\text{var}(X) = p(1 - p)$. Also, if X and Y are indicator variables then XY is an indicator variable and $\mathbb{P}(XY = 1) = \mathbb{P}(X = 1, Y = 1)$. The results then follow from the definitions.

In particular, note that A and B are positively correlated, negatively correlated, or independent, respectively (as defined in the section on conditional probability) if and only if the indicator variables of A and B are positively correlated, negatively correlated, or uncorrelated, as defined in this section.

If A and B are events then

1. $\text{cov}(A, B^c) = -\text{cov}(A, B)$
2. $\text{cov}(A^c, B^c) = \text{cov}(A, B)$

Proof

These results follow from linear property (7) and the fact that $\mathbf{1}_{A^c} = 1 - \mathbf{1}_A$.

If A and B are events with $A \subseteq B$ then

1. $\text{cov}(A, B) = \mathbb{P}(A)[1 - \mathbb{P}(B)]$
2. $\text{cor}(A, B) = \sqrt{\mathbb{P}(A)[1 - \mathbb{P}(B)] / \mathbb{P}(B)[1 - \mathbb{P}(A)]}$

Proof

These results follow from (19), since $A \cap B = A$.

In the language of the experiment, $A \subseteq B$ means that A implies B . In such a case, the events are positively correlated, not surprising.

The Best Linear Predictor

What linear function of X (that is, a function of the form $a + bX$ where $a, b \in \mathbb{R}$) is closest to Y in the sense of minimizing mean square error? The question is fundamentally important in the case where random variable X (the *predictor variable*) is observable and random variable Y (the *response variable*) is not. The linear function can be used to estimate Y from an observed value of X . Moreover, the solution will have the added benefit of showing that covariance and correlation measure the *linear* relationship between X and Y . To avoid trivial cases, let us assume that $\text{var}(X) > 0$ and $\text{var}(Y) > 0$, so that the random variables really are random. The solution to our problem turns out to be the linear function of X with the same expected value as Y , and whose covariance with X is the same as that of Y .

The random variable $L(Y | X)$ defined as follows is the only linear function of X satisfying properties (a) and (b).

$$L(Y | X) = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)}[X - \mathbb{E}(X)] \quad (4.5.17)$$

1. $\mathbb{E}[L(Y | X)] = \mathbb{E}(Y)$
2. $\text{cov}[X, L(Y | X)] = \text{cov}(X, Y)$

Proof

By the linearity of expected value,

$$\mathbb{E}[L(Y | X)] = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)}[\mathbb{E}(X) - \mathbb{E}(X)] = \mathbb{E}(Y) \quad (4.5.18)$$

Next, by the linearity of covariance and the fact that a constant is independent (and hence uncorrelated) with any random variable,

$$\text{cov}[X, L(Y | X)] = \frac{\text{cov}(X, Y)}{\text{var}(X)}\text{cov}(X, X) = \frac{\text{cov}(X, Y)}{\text{var}(X)}\text{var}(X) = \text{cov}(X, Y) \quad (4.5.19)$$

Conversely, suppose that $U = a + bX$ satisfies $\mathbb{E}(U) = \mathbb{E}(Y)$ and $\text{cov}(X, U) = \text{cov}(X, Y)$. Again using linearity of covariance and the uncorrelated property of constants, the second equation gives $b \text{cov}(X, X) = \text{cov}(X, Y)$ so $b = \text{cov}(X, Y) / \text{var}(X)$. Then the first equation gives $a = \mathbb{E}(Y) - b\mathbb{E}(X)$, so $U = L(Y | X)$.

Note that in the presence of part (a), part (b) is equivalent to $\mathbb{E}[XL(Y | X)] = \mathbb{E}(XY)$. Here is another minor variation, but one that will be very useful: $L(Y | X)$ is the only linear function of X with the same mean as Y and with the property that $Y - L(Y | X)$ is uncorrelated with every linear function of X .

$L(Y | X)$ is the only linear function of X that satisfies

1. $\mathbb{E}[L(Y | X)] = \mathbb{E}(Y)$
2. $\text{cov}[Y - L(Y | X), U] = 0$ for every linear function U of X .

Proof

Of course part (a) is the same as part (a) of (22). Suppose that $U = a + bX$ where $a, b \in \mathbb{R}$. From basic properties of covariance and the previous result,

$$\text{cov}[Y - L(Y|X), U] = b \text{cov}[Y - L(Y|X), X] = b(\text{cov}(Y, X) - \text{cov}[L(Y|X), X]) = 0 \quad (4.5.20)$$

Conversely, suppose that V is a linear function of X and that $\mathbb{E}(V) = \mathbb{E}(Y)$ and $\text{cov}(Y - V, U) = 0$ for every linear function U of X . Letting $U = X$ we have $\text{cov}(Y - V, X) = 0$ so $\text{cov}(V, X) = \text{cov}(Y, X)$. Hence $V = L(Y|X)$ by (22).

The variance of $L(Y|X)$ and its covariance with Y turn out to be the same.

Additional properties of $L(Y|X)$:

1. $\text{var}[L(Y|X)] = \text{cov}^2(X, Y) / \text{var}(X)$
2. $\text{cov}[L(Y|X), Y] = \text{cov}^2(X, Y) / \text{var}(X)$

Proof

1. From basic properties of variance,

$$\text{var}[L(Y|X)] = \left[\frac{\text{cov}(X, Y)}{\text{var}(X)} \right]^2 \text{var}(X) = \frac{\text{cov}^2(X, Y)}{\text{var}(X)} \quad (4.5.21)$$

2. From basic properties of covariance,

$$\text{cov}[L(Y|X), Y] = \frac{\text{cov}(X, Y)}{\text{var}(X)} \text{cov}(X, Y) = \frac{\text{cov}^2(X, Y)}{\text{var}(X)} \quad (4.5.22)$$

We can now prove the fundamental result that $L(Y|X)$ is the linear function of X that is closest to Y in the mean square sense. We give two proofs; the first is more straightforward, but the second is more interesting and elegant.

Suppose that U is a linear function of X . Then

1. $\mathbb{E}[(Y - L(Y|X))^2] \leq \mathbb{E}[(Y - U)^2]$
2. Equality occurs in (a) if and only if $U = L(Y|X)$ with probability 1.

Proof from calculus

Let $\text{mse}(a, b)$ denote the mean square error when $U = a + bX$ is used as an estimator of Y , as a function of the parameters $a, b \in \mathbb{R}$:

$$\text{mse}(a, b) = \mathbb{E}[(Y - (a + bX))^2] \quad (4.5.23)$$

Expanding the square and using the linearity of expected value gives

$$\text{mse}(a, b) = a^2 + b^2 \mathbb{E}(X^2) + 2ab \mathbb{E}(X) - 2a \mathbb{E}(Y) - 2b \mathbb{E}(XY) + \mathbb{E}(Y^2) \quad (4.5.24)$$

In terms of the variables a and b , the first three terms are the *second-order terms*, the next two are the *first-order terms*, and the last is the *zero-order term*. The second-order terms define a *quadratic form* whose standard symmetric matrix is

$$\begin{bmatrix} 1 & \mathbb{E}(X) \\ \mathbb{E}(X) & \mathbb{E}(X^2) \end{bmatrix} \quad (4.5.25)$$

The determinant of this matrix is $\mathbb{E}(X^2) - [\mathbb{E}(X)]^2 = \text{var}(X)$ and the diagonal terms are positive. All of this means that the graph of mse is a paraboloid opening upward, so the minimum of mse will occur at the unique critical point. Setting the first derivatives of mse to 0 we have

$$-2\mathbb{E}(Y) + 2b\mathbb{E}(X) + 2a = 0 \quad (4.5.26)$$

$$-2\mathbb{E}(XY) + 2b\mathbb{E}(X^2) + 2a\mathbb{E}(X) = 0 \quad (4.5.27)$$

Solving the first equation for a gives $a = \mathbb{E}(Y) - b\mathbb{E}(X)$. Substituting this into the second equation and solving gives $b = \text{cov}(X, Y) / \text{var}(X)$.

Proof using properties

1. We abbreviate $L(Y|X)$ by L for simplicity. Suppose that U is a linear function of X . Then

$$\mathbb{E}[(Y - U)^2] = \mathbb{E}[(Y - L) + (L - U)]^2 = \mathbb{E}[(Y - L)^2] + 2\mathbb{E}[(Y - L)(L - U)] + \mathbb{E}[(L - U)^2] \quad (4.5.28)$$

Since $Y - L$ has mean 0, the middle term is $\text{cov}(Y - L, L - U)$. But L and U are linear functions of X and hence so is $L - U$. Thus $\text{cov}(Y - L, L - U) = 0$ by (23). Hence

$$\mathbb{E}[(Y - U)^2] = \mathbb{E}[(Y - L)^2] + \mathbb{E}[(L - U)^2] \geq \mathbb{E}[(Y - L)^2] \quad (4.5.29)$$

2. Equality occurs in (a) if and only if $\mathbb{E}[(L - U)^2] = 0$, if and only if $\mathbb{P}(L = U) = 1$.

The mean square error when $L(Y | X)$ is used as a predictor of Y is

$$\mathbb{E} \left([Y - L(Y | X)]^2 \right) = \text{var}(Y) [1 - \text{cor}^2(X, Y)] \quad (4.5.30)$$

Proof

Again, let $L = L(Y | X)$ for convenience. Since $Y - L$ has mean 0,

$$\mathbb{E} [(Y - L)^2] = \text{var}(Y - L) = \text{var}(Y) - 2\text{cov}(L, Y) + \text{var}(L) \quad (4.5.31)$$

But $\text{cov}(L, Y) = \text{var}(L) = \text{cov}^2(X, Y) / \text{var}(X)$ by (24). Hence

$$\mathbb{E} [(Y - L)^2] = \text{var}(Y) - \frac{\text{cov}^2(X, Y)}{\text{var}(X)} = \text{var}(Y) \left[1 - \frac{\text{cov}^2(X, Y)}{\text{var}(X)\text{var}(Y)} \right] = \text{var}(Y) [1 - \text{cor}^2(X, Y)] \quad (4.5.32)$$

Our solution to the best linear predictor problems yields important properties of covariance and correlation.

Additional properties of covariance and correlation:

1. $-1 \leq \text{cor}(X, Y) \leq 1$
2. $-\text{sd}(X)\text{sd}(Y) \leq \text{cov}(X, Y) \leq \text{sd}(X)\text{sd}(Y)$
3. $\text{cor}(X, Y) = 1$ if and only if, with probability 1, Y is a linear function of X with positive slope.
4. $\text{cor}(X, Y) = -1$ if and only if, with probability 1, Y is a linear function of X with negative slope.

Proof

Since mean square error is nonnegative, it follows from (26) that $\text{cor}^2(X, Y) \leq 1$. This gives parts (a) and (b). For parts (c) and (d), note that if $\text{cor}^2(X, Y) = 1$ then $Y = L(Y | X)$ with probability 1, and that the slope in $L(Y | X)$ has the same sign as $\text{cor}(X, Y)$.

The last two results clearly show that $\text{cov}(X, Y)$ and $\text{cor}(X, Y)$ measure the *linear* association between X and Y . The equivalent inequalities (a) and (b) above are referred to as the *correlation inequality*. They are also versions of the *Cauchy-Schwarz inequality*, named for Augustin Cauchy and Karl Schwarz

Recall from our previous discussion of variance that the best *constant* predictor of Y , in the sense of minimizing mean square error, is $\mathbb{E}(Y)$ and the minimum value of the mean square error for this predictor is $\text{var}(Y)$. Thus, the difference between the variance of Y and the *mean square error* above for $L(Y | X)$ is the reduction in the variance of Y when the linear term in X is added to the predictor:

$$\text{var}(Y) - \mathbb{E} \left([Y - L(Y | X)]^2 \right) = \text{var}(Y) \text{cor}^2(X, Y) \quad (4.5.33)$$

Thus $\text{cor}^2(X, Y)$ is the *proportion* of reduction in $\text{var}(Y)$ when X is included as a predictor variable. This quantity is called the (distribution) *coefficient of determination*. Now let

$$L(Y | X = x) = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)} [x - \mathbb{E}(X)], \quad x \in \mathbb{R} \quad (4.5.34)$$

The function $x \mapsto L(Y | X = x)$ is known as the *distribution regression function* for Y given X , and its graph is known as the *distribution regression line*. Note that the regression line passes through $(\mathbb{E}(X), \mathbb{E}(Y))$, the center of the joint distribution.




Figure 4.5.2: The distribution regression line

However, the choice of predictor variable and response variable is crucial.

The regression line for Y given X and the regression line for X given Y are not the same line, except in the trivial case where the variables are perfectly correlated. However, the coefficient of determination is the same, regardless of which variable is the predictor and which is the response.

Proof

The two regression lines are

$$y - \mathbb{E}(Y) = \frac{\text{cov}(X, Y)}{\text{var}(X)} [x - \mathbb{E}(X)] \quad (4.5.35)$$

$$x - \mathbb{E}(X) = \frac{\text{cov}(X, Y)}{\text{var}(Y)} [y - \mathbb{E}(Y)] \quad (4.5.36)$$

The two lines are the same if and only if $\text{cov}^2(X, Y) = \text{var}(X)\text{var}(Y)$. But this is equivalent to $\text{cor}^2(X, Y) = 1$.

Suppose that A and B are events with $0 < \mathbb{P}(A) < 1$ and $0 < \mathbb{P}(B) < 1$. Then

1. $\text{cor}(A, B) = 1$ if and only if $\mathbb{P}(A \setminus B) + \mathbb{P}(B \setminus A) = 0$. (That is, A and B are equivalent events.)
2. $\text{cor}(A, B) = -1$ if and only if $\mathbb{P}(A \setminus B^c) + \mathbb{P}(B^c \setminus A) = 0$. (That is, A and B^c are equivalent events.)

Proof

Recall from (19) that $\text{cor}(A, B) = \text{cor}(\mathbf{1}_A, \mathbf{1}_B)$, so if $\text{cor}^2(A, B) = 1$ then from (27), $\mathbf{1}_B = L(\mathbf{1}_B | \mathbf{1}_A)$ with probability 1. But $\mathbf{1}_A$ and $\mathbf{1}_B$ each takes values 0 and 1 only. Hence the only possible regression lines are $y = 0$, $y = 1$, $y = x$ and $y = 1 - x$. The first two correspond to $\mathbb{P}(B) = 0$ and $\mathbb{P}(B) = 1$, respectively, which are excluded by the hypotheses.

1. In this case, the slope is positive, so the regression line is $y = x$. That is, $\mathbf{1}_B = \mathbf{1}_A$ with probability 1.
2. In this case, the slope is negative, so the regression line is $y = 1 - x$. That is, $\mathbf{1}_B = 1 - \mathbf{1}_A = \mathbf{1}_{A^c}$ with probability 1.

The concept of best linear predictor is more powerful than might first appear, because it can be applied to *transformations* of the variables. Specifically, suppose that X and Y are random variables for our experiment, taking values in general spaces S and T , respectively. Suppose also that g and h are real-valued functions defined on S and T , respectively. We can find $L[h(Y) | g(X)]$, the linear function of $g(X)$ that is closest to $h(Y)$ in the mean square sense. The results of this subsection apply, of course, with $g(X)$ replacing X and $h(Y)$ replacing Y . Of course, we must be able to compute the appropriate means, variances, and covariances.

We close this subsection with two additional properties of the best linear predictor, the *linearity properties*.

Suppose that X , Y , and Z are random variables and that c is a constant. Then

1. $L(Y + Z | X) = L(Y | X) + L(Z | X)$
2. $L(cY | X) = cL(Y | X)$

Proof from the definitions

These results follow easily from the linearity of expected value and covariance.

$$1. \quad L(Y + Z | X) = \mathbb{E}(Y + Z) + \frac{\text{cov}(X, Y + Z)}{\text{var}(X)} [X - \mathbb{E}(X)] \quad (4.5.37)$$

$$= \left(\mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)} [X - \mathbb{E}(X)] \right) + \left(\mathbb{E}(Z) + \frac{\text{cov}(X, Z)}{\text{var}(X)} [X - \mathbb{E}(X)] \right) \quad (4.5.38)$$

$$= \mathbb{E}(Y | X) + \mathbb{E}(Z | X) \quad (4.5.39)$$

$$2. \quad L(cY | X) = \mathbb{E}(cY) + \frac{\text{cov}(X, cY)}{\text{var}(X)} [X - \mathbb{E}(X)] = c\mathbb{E}(Y) + c \frac{\text{cov}(X, Y)}{\text{var}(X)} [X - \mathbb{E}(X)] = cL(Y | X) \quad (4.5.40)$$

Proof by characterizing properties

1. We show that $L(Y | X) + L(Z | X)$ satisfy the properties that characterize $L(Y + Z | X)$.

$$\mathbb{E}[L(Y | X) + L(Z | X)] = \mathbb{E}[L(Y | X)] + \mathbb{E}[L(Z | X)] = \mathbb{E}(Y) + \mathbb{E}(Z) = \mathbb{E}(Y + Z) \quad (4.5.41)$$

$$\text{cov}[X, L(Y | X) + L(Z | X)] = \text{cov}[X, L(Y | X)] + \text{cov}[X, L(Z | X)] = \text{cov}(X, Y) + \text{cov}(X, Z) = \text{cov}(X, Y + Z) \quad (4.5.42)$$

2. Similarly, we show that $cL(Y | X)$ satisfies the properties that characterize $L(cY | X)$

$$\mathbb{E}[cL(Y | X)] = c\mathbb{E}[L(Y | X)] = c\mathbb{E}(Y) = \mathbb{E}(cY) \quad (4.5.43)$$

$$\text{cov}[X, cL(Y | X)] = c \text{cov}[X, L(Y | X)] = c \text{cov}(X, Y) = \text{cov}(X, cY) \quad (4.5.44)$$

There are several extensions and generalizations of the ideas in the subsection:

- The corresponding statistical problem of estimating a and b , when these distribution parameters are unknown, is considered in the section on Sample Covariance and Correlation.
- The problem finding the function of X that is closest to Y in the mean square error sense (using *all* reasonable functions, not just linear functions) is considered in the section on Conditional Expected Value.
- The best linear prediction problem when the predictor and response variables are random vectors is considered in the section on Expected Value and Covariance Matrices.

The use of characterizing properties will play a crucial role in these extensions.

Examples and Applications

Uniform Distributions

Suppose that X is uniformly distributed on the interval $[-1, 1]$ and $Y = X^2$. Then X and Y are uncorrelated even though Y is a function of X (the strongest form of dependence).

Proof

Note that $\mathbb{E}(X) = 0$ and $\mathbb{E}(Y) = \mathbb{E}(X^2) = 1/3$ and $\mathbb{E}(XY) = \mathbb{E}(X^3) = 0$. Hence $\text{cov}(X, Y) = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y) = 0$.

Suppose that (X, Y) is uniformly distributed on the region $S \subseteq \mathbb{R}^2$. Find $\text{cov}(X, Y)$ and $\text{cor}(X, Y)$ and determine whether the variables are independent in each of the following cases:

1. $S = [a, b] \times [c, d]$ where $a < b$ and $c < d$, so S is a rectangle.
2. $S = \{(x, y) \in \mathbb{R}^2 : -a \leq y \leq x \leq a\}$ where $a > 0$, so S is a triangle
3. $S = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq r^2\}$ where $r > 0$, so S is a circle

Answer

1. $\text{cov}(X, Y) = 0$, $\text{cor}(X, Y) = 0$. X and Y are independent.
2. $\text{cov}(X, Y) = \frac{a^2}{9}$, $\text{cor}(X, Y) = \frac{1}{2}$. X and Y are dependent.
3. $\text{cov}(X, Y) = 0$, $\text{cor}(X, Y) = 0$. X and Y are dependent.

In the bivariate uniform experiment, select each of the regions below in turn. For each region, run the simulation 2000 times and note the value of the correlation and the shape of the cloud of points in the scatterplot. Compare with the results in the last exercise.

1. Square
2. Triangle
3. Circle

Suppose that X is uniformly distributed on the interval $(0, 1)$ and that given $X = x \in (0, 1)$, Y is uniformly distributed on the interval $(0, x)$. Find each of the following:

1. $\text{cov}(X, Y)$
2. $\text{cor}(X, Y)$
3. $L(Y | X)$
4. $L(X | Y)$

Answer

1. $\frac{1}{24}$
2. $\sqrt{\frac{3}{7}}$
3. $\frac{1}{2}X$
4. $\frac{2}{7} + \frac{6}{7}Y$

Dice

Recall that a *standard die* is a six-sided die. A *fair die* is one in which the faces are equally likely. An *ace-six flat die* is a standard die in which faces 1 and 6 have probability $\frac{1}{4}$ each, and faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each.

A pair of standard, fair dice are thrown and the scores (X_1, X_2) recorded. Let $Y = X_1 + X_2$ denote the sum of the scores, $U = \min\{X_1, X_2\}$ the minimum scores, and $V = \max\{X_1, X_2\}$ the maximum score. Find the covariance and correlation of each of the following pairs of variables:

1. (X_1, X_2)
2. (X_1, Y)
3. (X_1, U)
4. (U, V)
5. (U, Y)

Answer

1. 0, 0
2. $\frac{35}{12}, \frac{1}{\sqrt{2}} = 0.7071$
3. $\frac{35}{24}, 0.6082$
4. $\frac{1369}{1296}, \frac{1369}{2555} = 0.5358$
5. $\frac{35}{12}, 0.8601$

Suppose that n fair dice are thrown. Find the mean and variance of each of the following variables:

1. Y_n , the sum of the scores.
2. M_n , the average of the scores.

Answer

1. $\mathbb{E}(Y_n) = \frac{7}{2}n$, $\text{var}(Y_n) = \frac{35}{12}n$
2. $\mathbb{E}(M_n) = \frac{7}{2}$, $\text{var}(M_n) = \frac{35}{12n}$

In the dice experiment, select fair dice, and select the following random variables. In each case, increase the number of dice and observe the size and location of the probability density function and the mean \pm standard deviation bar. With $n = 20$ dice, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

1. The sum of the scores.
2. The average of the scores.

Suppose that n ace-six flat dice are thrown. Find the mean and variance of each of the following variables:

1. Y_n , the sum of the scores.
2. M_n , the average of the scores.

Answer

1. $n\frac{7}{2}$, $n\frac{15}{4}$
2. $\frac{7}{2}$, $\frac{15}{4n}$

In the dice experiment, select ace-six flat dice, and select the following random variables. In each case, increase the number of dice and observe the size and location of the probability density function and the mean \pm standard deviation bar. With $n = 20$ dice, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

1. The sum of the scores.
2. The average of the scores.

A pair of fair dice are thrown and the scores (X_1, X_2) recorded. Let $Y = X_1 + X_2$ denote the sum of the scores, $U = \min\{X_1, X_2\}$ the minimum score, and $V = \max\{X_1, X_2\}$ the maximum score. Find each of the following:

1. $L(Y | X_1)$
2. $L(U | X_1)$
3. $L(V | X_1)$

Answer

1. $\frac{7}{2} + X_1$
2. $\frac{7}{9} + \frac{1}{2}X_1$
3. $\frac{49}{19} + \frac{1}{2}X_1$

Bernoulli Trials

Recall that a *Bernoulli trials process* is a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent, identically distributed indicator random variables. In the usual language of reliability, X_i denotes the outcome of trial i , where 1 denotes success and 0 denotes failure. The probability of success $p = \mathbb{P}(X_i = 1)$ is the basic parameter of the process. The process is named for Jacob Bernoulli. A separate chapter on the Bernoulli Trials explores this process in detail.

For $n \in \mathbb{N}_+$, the number of successes in the first n trials is $Y_n = \sum_{i=1}^n X_i$. Recall that this random variable has the binomial distribution with parameters n and p , which has probability density function f given by

$$f_n(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (4.5.45)$$

The mean and variance of Y_n are

1. $\mathbb{E}(Y_n) = np$
2. $\text{var}(Y_n) = np(1-p)$

Proof

These results could be derived from the PDF of Y_n , of course, but a derivation based on the sum of IID variables is much better. Recall that $\mathbb{E}(X_i) = p$ and $\text{var}(X_i) = p(1-p)$ so the results follow immediately from theorem (16).

In the binomial coin experiment, select the number of heads. Vary n and p and note the shape of the probability density function and the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the sample mean

and standard deviation to the distribution mean and standard deviation.

For $n \in \mathbb{N}_+$, the proportion of successes in the first n trials is $M_n = Y_n/n$. This random variable is sometimes used as a statistical estimator of the parameter p , when the parameter is unknown.

The mean and variance of M_n are

1. $\mathbb{E}(M_n) = p$
2. $\text{var}(M_n) = p(1-p)/n$

Proof

Recall that $\mathbb{E}(X_i) = p$ and $\text{var}(X_i) = p(1-p)$ so the results follow immediately from theorem (17).

In the binomial coin experiment, select the proportion of heads. Vary n and p and note the shape of the probability density function and the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

As a special case of (17) note that $M_n \rightarrow p$ as $n \rightarrow \infty$ in mean square and in probability.

The Hypergeometric Distribution

Suppose that a population consists of m objects; r of the objects are type 1 and $m-r$ are type 0. A sample of n objects is chosen at random, without replacement. The parameters $m, n \in \mathbb{N}_+$ and $r \in \mathbb{N}$ with $n \leq m$ and $r \leq m$. For $i \in \{1, 2, \dots, n\}$, let X_i denote the type of the i th object selected. Recall that (X_1, X_2, \dots, X_n) is a sequence of identically distributed (but *not* independent) indicator random variables.

Let Y denote the number of type 1 objects in the sample, so that $Y = \sum_{i=1}^n X_i$. Recall that this random variable has the hypergeometric distribution, which has probability density function f_n given by

$$f(y) = \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}}, \quad y \in \{0, 1, \dots, n\} \quad (4.5.46)$$

For distinct $i, j \in \{1, 2, \dots, n\}$,

1. $\mathbb{E}(X_i) = \frac{r}{m}$
2. $\text{var}(X_i) = \frac{r}{m} \left(1 - \frac{r}{m}\right)$
3. $\text{cov}(X_i, X_j) = -\frac{r}{m} \left(1 - \frac{r}{m}\right) \frac{1}{m-1}$
4. $\text{cor}(X_i, X_j) = -\frac{1}{m-1}$

Proof

Recall that $\mathbb{E}(X_i) = \mathbb{P}(X_i = 1) = \frac{r}{m}$ for each i and $\mathbb{E}(X_i X_j) = \mathbb{P}(X_i = 1, X_j = 1) = \frac{r}{m} \frac{r-1}{m-1}$ for each $i \neq j$. Technically, the sequence of indicator variables is exchangeable. The results now follow from the definitions and simple algebra.

Note that the event of a type 1 object on draw i and the event of a type 1 object on draw j are negatively correlated, but the correlation depends only on the population size and not on the number of type 1 objects. Note also that the correlation is perfect if $m = 2$. Think about these result intuitively.

The mean and variance of Y are

1. $\mathbb{E}(Y) = n \frac{r}{m}$
2. $\text{var}(Y) = n \frac{r}{m} \left(1 - \frac{r}{m}\right) \frac{m-n}{m-1}$

Proof

Again, a derivation from the representation of Y as a sum of indicator variables is far preferable to a derivation based on the PDF of Y . These results follow immediately from (45), the additive property of expected value, and Theorem (12).

Note that if the sampling were *with* replacement, Y would have a binomial distribution, and so in particular $E(Y) = n \frac{r}{m}$ and $\text{var}(Y) = n \frac{r}{m} \left(1 - \frac{r}{m}\right)$. The additional factor $\frac{m-n}{m-1}$ that occurs in the variance of the hypergeometric distribution is sometimes called the *finite population correction factor*. Note that for fixed m , $\frac{m-n}{m-1}$ is decreasing in n , and is 0 when $n = m$. Of course, we know that we must have $\text{var}(Y) = 0$ if $n = m$, since we would be sampling the entire population, and so deterministically, $Y = r$. On the other hand, for fixed n , $\frac{m-n}{m-1} \rightarrow 1$ as $m \rightarrow \infty$. More generally, the hypergeometric distribution is well approximated by the binomial when the population size m is large compared to the sample size n . These ideas are discussed more fully in the section on the hypergeometric distribution in the chapter on Finite Sampling Models.

In the ball and urn experiment, select sampling without replacement. Vary m, r , and n and note the shape of the probability density function and the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

Exercises on Basic Properties

Suppose that X and Y are real-valued random variables with $\text{cov}(X, Y) = 3$. Find $\text{cov}(2X - 5, 4Y + 2)$.

Answer

24

Suppose X and Y are real-valued random variables with $\text{var}(X) = 5$, $\text{var}(Y) = 9$, and $\text{cov}(X, Y) = -3$. Find

1. $\text{cor}(X, Y)$
2. $\text{var}(2X + 3Y - 7)$
3. $\text{cov}(5X + 2Y - 3, 3X - 4Y + 2)$
4. $\text{cor}(5X + 2Y - 3, 3X - 4Y + 2)$

Answer

1. $-\frac{1}{\sqrt{5}} \approx -0.4472$
2. 65
3. 45
4. $\frac{15}{\sqrt{2929}} \approx 0.2772$

Suppose that X and Y are independent, real-valued random variables with $\text{var}(X) = 6$ and $\text{var}(Y) = 8$. Find $\text{var}(3X - 4Y + 5)$.

Answer

182

Suppose that A and B are events in an experiment with $\mathbb{P}(A) = \frac{1}{2}$, $\mathbb{P}(B) = \frac{1}{3}$, and $\mathbb{P}(A \cap B) = \frac{1}{8}$. Find each of the following:

1. $\text{cov}(A, B)$
2. $\text{cor}(A, B)$

Answer

1. $-\frac{1}{24}$
2. $-\sqrt{2}/8$

Suppose that X , Y , and Z are real-valued random variables for an experiment, and that $L(Y | X) = 2 - 3X$ and $L(Z | X) = 5 + 4X$. Find $L(6Y - 2Z | X)$.

Answer

$2 - 26X$

Suppose that X and Y are real-valued random variables for an experiment, and that $\mathbb{E}(X) = 3$, $\text{var}(X) = 4$, and $L(Y | X) = 5 - 2X$. Find each of the following:

1. $\mathbb{E}(Y)$
2. $\text{cov}(X, Y)$

Answer

1. -1
2. -8

Simple Continuous Distributions

Suppose that (X, Y) has probability density function f given by $f(x, y) = x + y$ for $0 \leq x \leq 1$, $0 \leq y \leq 1$. Find each of the following

1. $\text{cov}(X, Y)$
2. $\text{cor}(X, Y)$
3. $L(Y | X)$
4. $L(X | Y)$

Answer

1. $-\frac{1}{144}$
2. $-\frac{1}{11} \approx -0.0909$
3. $\frac{7}{11} - \frac{1}{11}X$
4. $\frac{7}{11} = \frac{1}{11}Y$

Suppose that (X, Y) has probability density function f given by $f(x, y) = 2(x + y)$ for $0 \leq x \leq y \leq 1$. Find each of the following:

1. $\text{cov}(X, Y)$
2. $\text{cor}(X, Y)$
3. $L(Y | X)$
4. $L(X | Y)$

Answer

1. $\frac{1}{48}$
2. $\frac{5}{\sqrt{129}} \approx 0.4402$
3. $\frac{26}{43} + \frac{15}{43}X$
4. $\frac{5}{9}Y$

Suppose again that (X, Y) has probability density function f given by $f(x, y) = 2(x + y)$ for $0 \leq x \leq y \leq 1$.

1. Find $\text{cov}(X^2, Y)$.
2. Find $\text{cor}(X^2, Y)$.
3. Find $L(Y | X^2)$.
4. Which predictor of Y is better, the one based on X or the one based on X^2 ?

Answer

1. $\frac{7}{360}$
2. 0.448
3. $\frac{1255}{1920} + \frac{245}{634}X$
4. The predictor based on X^2 is slightly better.

Suppose that (X, Y) has probability density function f given by $f(x, y) = 6x^2y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$. Find each of the following:

1. $\text{cov}(X, Y)$
2. $\text{cor}(X, Y)$
3. $L(Y | X)$
4. $L(X | Y)$

Answer

Note that X and Y are independent.

1. 0
2. 0
3. $\frac{2}{3}$
4. $\frac{3}{4}$

Suppose that (X, Y) has probability density function f given by $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$. Find each of the following:

1. $\text{cov}(X, Y)$
2. $\text{cor}(X, Y)$
3. $L(Y | X)$
4. $L(X | Y)$

Answer

1. $\frac{5}{336}$
2. 0.05423
3. $\frac{30}{51} + \frac{20}{51}X$
4. $\frac{3}{4}Y$

Suppose again that (X, Y) has probability density function f given by $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$.

1. Find $\text{cov}(\sqrt{X}, Y)$.
2. Find $\text{cor}(\sqrt{X}, Y)$.
3. Find $L(Y | \sqrt{X})$.
4. Which of the predictors of Y is better, the one based on X or the one based on \sqrt{X} ?

Answer

1. $\frac{10}{1001}$
2. $\frac{24}{169}\sqrt{14}$

3. $\frac{5225}{13\,182} + \frac{1232}{2197}X$

4. The predictor based on X is slightly better.

This page titled [4.5: Covariance and Correlation](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.6: Generating Functions

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. A *generating function* of a real-valued random variable is an expected value of a certain transformation of the random variable involving another (deterministic) variable. Most generating functions share four important properties:

1. Under mild conditions, the generating function completely determines the distribution of the random variable.
2. The generating function of a sum of independent variables is the product of the generating functions
3. The moments of the random variable can be obtained from the derivatives of the generating function.
4. Ordinary (pointwise) convergence of a sequence of generating functions corresponds to the special convergence of the corresponding distributions.

Property 1 is perhaps the most important. Often a random variable is shown to have a certain distribution by showing that the generating function has a certain form. The process of recovering the distribution from the generating function is known as *inversion*. Property 2 is frequently used to determine the distribution of a sum of independent variables. By contrast, recall that the probability density function of a sum of independent variables is the convolution of the individual density functions, a much more complicated operation. Property 3 is useful because often computing moments from the generating function is easier than computing the moments directly from the probability density function. The last property is known as the *continuity theorem*. Often it is easier to show the convergence of the generating functions than to prove convergence of the distributions directly.

The numerical value of the generating function at a particular value of the free variable is of no interest, and so generating functions can seem rather unintuitive at first. But the important point is that the generating function as a whole encodes all of the information in the probability distribution in a very useful way. Generating functions are important and valuable tools in probability, as they are in other areas of mathematics, from combinatorics to differential equations.

We will study the three generating functions in the list below, which correspond to increasing levels of generality. The first is the most restrictive, but also by far the simplest, since the theory reduces to basic facts about power series that you will remember from calculus. The third is the most general and the one for which the theory is most complete and elegant, but it also requires basic knowledge of complex analysis. The one in the middle is perhaps the one most commonly used, and suffices for most distributions in applied probability.

1. [the probability generating function](#)
2. [the moment generating function](#)
3. [the characteristic function](#)

We will also study the characteristic function for multivariate distributions, although analogous results hold for the other two types. In the basic theory below, be sure to try the proofs yourself before reading the ones in the text.

Basic Theory

The Probability Generating Function

For our first generating function, assume that N is a random variable taking values in \mathbb{N} .

The *probability generating function* P of N is defined by

$$P(t) = \mathbb{E}(t^N) \quad (4.6.1)$$

for all $t \in \mathbb{R}$ for which the expected value exists in \mathbb{R} .

That is, $P(t)$ is defined when $\mathbb{E}(|t|^N) < \infty$. The probability generating function can be written nicely in terms of the probability density function.

Suppose that N has probability density function f and probability generating function P . Then

$$P(t) = \sum_{n=0}^{\infty} f(n)t^n, \quad t \in (-r, r) \quad (4.6.2)$$

where $r \in [1, \infty]$ is the radius of convergence of the series.

Proof

The expansion follows from the discrete change of variables theorem for expected value. Note that the series is a power series in t , and hence by basic calculus, converges absolutely for $t \in (-r, r)$ where $r \in [0, \infty]$ is the radius of convergence. But since $\sum_{n=0}^{\infty} f(n) = 1$ we must have $r \geq 1$, and the series converges absolutely at least for $t \in [-1, 1]$.

In the language of combinatorics, P is the *ordinary generating function* of f . Of course, if N just takes a finite set of values in \mathbb{N} then $r = \infty$. Recall from calculus that a power series can be differentiated term by term, just like a polynomial. Each derivative series has the same radius of convergence as the original series (but may behave differently at the endpoints of the interval of convergence). We denote the derivative of order n by $P^{(n)}$. Recall also that if $n \in \mathbb{N}$ and $k \in \mathbb{N}$ with $k \leq n$, then the number of permutations of size k chosen from a population of n objects is

$$n^{(k)} = n(n-1) \cdots (n-k+1) \quad (4.6.3)$$

The following theorem is the inversion result for probability generating functions: the generating function completely determines the distribution.

Suppose again that N has probability density function f and probability generating function P . Then

$$f(k) = \frac{P^{(k)}(0)}{k!}, \quad k \in \mathbb{N} \quad (4.6.4)$$

Proof

This is a standard result from the theory of power series. Differentiating k times gives $P^{(k)}(t) = \sum_{n=k}^{\infty} n^{(k)} f(n) t^{n-k}$ for $t \in (-r, r)$. Hence $P^{(k)}(0) = k^{(k)} f(k) = k! f(k)$

Our next result is not particularly important, but has a certain curiosity.

$$\mathbb{P}(N \text{ is even}) = \frac{1}{2} [1 + P(-1)] .$$

Proof

Note that

$$P(1) + P(-1) = \sum_{n=0}^{\infty} f(n) + \sum_{n=0}^{\infty} (-1)^n f(n) = 2 \sum_{k=0}^{\infty} f(2k) = 2\mathbb{P}(N \text{ is even}) \quad (4.6.5)$$

We can combine the two sum since we know that the series converge absolutely at 1 and -1 .

Recall that the *factorial moment* of N of order $k \in \mathbb{N}$ is $\mathbb{E}[N^{(k)}]$. The factorial moments can be computed from the derivatives of the probability generating function. The factorial moments, in turn, determine the ordinary moments about 0 (sometimes referred to as *raw moments*).

Suppose that the radius of convergence $r > 1$. Then $P^{(k)}(1) = \mathbb{E}[N^{(k)}]$ for $k \in \mathbb{N}$. In particular, N has finite moments of all orders.

Proof

As before, $P^{(k)}(t) = \sum_{n=k}^{\infty} n^{(k)} f(n) t^{n-k}$ for $t \in (-r, r)$. Hence if $r > 1$ then $P^{(k)}(1) = \sum_{n=k}^{\infty} n^{(k)} f(n) = \mathbb{E}[N^{(k)}]$

Suppose again that $r > 1$. Then

1. $\mathbb{E}(N) = P'(1)$
2. $\text{var}(N) = P''(1) + P'(1)[1 - P'(1)]$

Proof

1. $\mathbb{E}(N) = \mathbb{E}[N^{(1)}] = P'(1)$.
2. $\mathbb{E}(N^2) = \mathbb{E}[N(N-1)] + \mathbb{E}(N) = \mathbb{E}[N^{(2)}] + \mathbb{E}(N) = P''(1) + P'(1)$. Hence from (a),
 $\text{var}(N) = P''(1) + P'(1) - [P'(1)]^2$.

Suppose that N_1 and N_2 are independent random variables taking values in \mathbb{N} , with probability generating functions P_1 and P_2 having radii of convergence r_1 and r_2 , respectively. Then the probability generating function P of $N_1 + N_2$ is given by $P(t) = P_1(t)P_2(t)$ for $|t| < r_1 \wedge r_2$.

Proof

Recall that the expected product of independent variables is the product of the expected values. Hence

$$P(t) = \mathbb{E}(t^{N_1+N_2}) = \mathbb{E}(t^{N_1}t^{N_2}) = \mathbb{E}(t^{N_1})\mathbb{E}(t^{N_2}) = P_1(t)P_2(t), \quad |t| < r_1 \wedge r_2 \quad (4.6.6)$$

The Moment Generating Function

Our next generating function is defined more generally, so in this discussion we assume that the random variables are real-valued.

The *moment generating function* of X is the function M defined by

$$M(t) = \mathbb{E}(e^{tX}), \quad t \in \mathbb{R} \quad (4.6.7)$$

Note that since $e^{tX} \geq 0$ with probability 1, $M(t)$ exists, as a real number or ∞ , for any $t \in \mathbb{R}$. But as we will see, our interest will be in the domain where $M(t) < \infty$.

Suppose that X has a continuous distribution on \mathbb{R} with probability density function f . Then

$$M(t) = \int_{-\infty}^{\infty} e^{tx} f(x) dx \quad (4.6.8)$$

Proof

This follows from the change of variables theorem for expected value.

Thus, the moment generating function of X is closely related to the *Laplace transform* of the probability density function f . The Laplace transform is named for Pierre Simon Laplace, and is widely used in many areas of applied mathematics, particularly differential equations. The basic inversion theorem for moment generating functions (similar to the inversion theorem for Laplace transforms) states that if $M(t) < \infty$ for t in an open interval about 0, then M completely determines the distribution of X . Thus, if two distributions on \mathbb{R} have moment generating functions that are equal (and finite) in an open interval about 0, then the distributions are the same.

Suppose that X has moment generating function M that is finite in an open interval I about 0. Then X has moments of all orders and

$$M(t) = \sum_{n=0}^{\infty} \frac{\mathbb{E}(X^n)}{n!} t^n, \quad t \in I \quad (4.6.9)$$

Proof

Under the hypotheses, the expected value operator can be interchanged with the infinite series for the exponential function:

$$M(t) = \mathbb{E}(e^{tX}) = \mathbb{E}\left(\sum_{n=0}^{\infty} \frac{X^n}{n!} t^n\right) = \sum_{n=0}^{\infty} \frac{\mathbb{E}(X^n)}{n!} t^n, \quad t \in I \quad (4.6.10)$$

The interchange is a special case of *Fubini's theorem*, named for Guido Fubini. For more details see the advanced section on properties of the integral in the chapter on Distributions.

So under the finite assumption in the [last theorem](#), the moment generating function, like the probability generating function, is a power series in t .

Suppose again that X has moment generating function M that is finite in an open interval about 0. Then $M^{(n)}(0) = \mathbb{E}(X^n)$ for $n \in \mathbb{N}$

Proof

This follows by the same argument as above for the [PGF](#): $M^{(n)}(0)/n!$ is the coefficient of order n in the [power series](#) above, namely $\mathbb{E}(X^n)/n!$. Hence $M^{(n)}(0) = \mathbb{E}(X^n)$.

Thus, the derivatives of the moment generating function at 0 determine the moments of the variable (hence the name). In the language of combinatorics, the moment generating function is the *exponential generating function* of the sequence of moments. Thus, a random variable that does not have finite moments of all orders cannot have a finite moment generating function. Even when a random variable *does* have moments of all orders, the moment generating function may not exist. A [counterexample](#) is constructed below.

For nonnegative random variables (which are very common in applications), the domain where the moment generating function is finite is easy to understand.

Suppose that X takes values in $[0, \infty)$ and has moment generating function M . If $M(t) < \infty$ for $t \in \mathbb{R}$ then $M(s) < \infty$ for $s \leq t$.

Proof

Since $X \geq 0$, if $s \leq t$ then $sX \leq tX$ and hence $e^{sX} \leq e^{tX}$. Hence $\mathbb{E}(e^{sX}) \leq \mathbb{E}(e^{tX})$.

So for a nonnegative random variable, either $M(t) < \infty$ for all $t \in \mathbb{R}$ or there exists $r \in (0, \infty)$ such that $M(t) < \infty$ for $t < r$. Of course, there are complementary results for non-positive random variables, but such variables are much less common. Next we consider what happens to the moment generating function under some simple transformations of the random variables.

Suppose that X has moment generating function M and that $a, b \in \mathbb{R}$. The moment generating function N of $Y = a + bX$ is given by $N(t) = e^{at} M(bt)$ for $t \in \mathbb{R}$.

Proof

$\mathbb{E}[e^{t(a+bX)}] = \mathbb{E}(e^{ta} e^{tbX}) = e^{ta} \mathbb{E}[e^{(tb)X}] = e^{at} M(bt)$ for $t \in \mathbb{R}$.

Recall that if $a \in \mathbb{R}$ and $b \in (0, \infty)$ then the transformation $a + bX$ is a *location-scale* transformation on the distribution of X , with *location parameter* a and *scale parameter* b . Location-scale transformations frequently arise when units are changed, such as length changed from inches to centimeters or temperature from degrees Fahrenheit to degrees Celsius.

Suppose that X_1 and X_2 are independent random variables with moment generating functions M_1 and M_2 respectively. The moment generating function M of $Y = X_1 + X_2$ is given by $M(t) = M_1(t)M_2(t)$ for $t \in \mathbb{R}$.

Proof

As with the PGF, the proof for the MGF relies on the law of exponents and the fact that the expected value of a product of independent variables is the product of the expected values:

$$\mathbb{E}[e^{t(X_1+X_2)}] = \mathbb{E}(e^{tX_1} e^{tX_2}) = \mathbb{E}(e^{tX_1}) \mathbb{E}(e^{tX_2}) = M_1(t)M_2(t), \quad t \in \mathbb{R} \quad (4.6.11)$$

The probability generating function of a variable can easily be converted into the moment generating function of the variable.

Suppose that X is a random variable taking values in \mathbb{N} with [probability generating function](#) G having radius of convergence r . The moment generating function M of X is given by $M(t) = G(e^t)$ for $t < \ln(r)$.

Proof

$M(t) = \mathbb{E}(e^{tX}) = \mathbb{E}[(e^t)^X] = G(e^t)$ for $e^t < r$.

The following theorem gives the *Chernoff bounds*, named for the mathematician Herman Chernoff. These are upper bounds on the tail events of a random variable.

If X has moment generating function M then

1. $\mathbb{P}(X \geq x) \leq e^{-tx} M(t)$ for $t > 0$
2. $\mathbb{P}(X \leq x) \leq e^{-tx} M(t)$ for $t < 0$

Proof

1. From Markov's inequality, $\mathbb{P}(X \geq x) = \mathbb{P}(e^{tX} \geq e^{tx}) \leq \mathbb{E}(e^{tX}) / e^{tx} = e^{-tx} M(t)$ if $t > 0$.
2. Similarly, $\mathbb{P}(X \leq x) = \mathbb{P}(e^{tX} \geq e^{tx}) \leq e^{-tx} M(t)$ if $t < 0$.

Naturally, the best Chernoff bound (in either (a) or (b)) is obtained by finding t that minimizes $e^{-tx} M(t)$.

The Characteristic Function

Our last generating function is the nicest from a mathematical point of view. Once again, we assume that our random variables are real-valued.

The *characteristic function* of X is the function χ defined by

$$\chi(t) = \mathbb{E}(e^{itX}) = \mathbb{E}[\cos(tX)] + i\mathbb{E}[\sin(tX)], \quad t \in \mathbb{R} \quad (4.6.12)$$

Note that χ is a complex valued function, and so this subsection requires some basic knowledge of complex analysis. The function χ is defined for all $t \in \mathbb{R}$ because the random variable in the expected value is bounded in magnitude. Indeed, $|e^{itX}| = 1$ for all $t \in \mathbb{R}$. Many of the properties of the characteristic function are more elegant than the corresponding properties of the probability or moment generating functions, because the characteristic function always exists.

If X has a continuous distribution on \mathbb{R} with probability density function f and characteristic function χ then

$$\chi(t) = \int_{-\infty}^{\infty} e^{itx} f(x) dx, \quad t \in \mathbb{R} \quad (4.6.13)$$

Proof

This follows from the change of variables theorem for expected value, albeit a complex version.

Thus, the characteristic function of X is closely related to the *Fourier transform* of the probability density function f . The Fourier transform is named for Joseph Fourier, and is widely used in many areas of applied mathematics.

As with other generating functions, the characteristic function completely determines the distribution. That is, random variables X and Y have the same distribution if and only if they have the same characteristic function. Indeed, the general *inversion formula* given next is a formula for computing certain combinations of probabilities from the characteristic function.

Suppose again that X has characteristic function χ . If $a, b \in \mathbb{R}$ and $a < b$ then

$$\int_{-n}^n \frac{e^{-iat} - e^{-ibt}}{2\pi it} \chi(t) dt \rightarrow \mathbb{P}(a < X < b) + \frac{1}{2}[\mathbb{P}(X = b) - \mathbb{P}(X = a)] \text{ as } n \rightarrow \infty \quad (4.6.14)$$

The probability combinations on the right side completely determine the distribution of X . A special inversion formula holds for continuous distributions:

Suppose that X has a continuous distribution with probability density function f and characteristic function χ . At every point $x \in \mathbb{R}$ where f is differentiable,

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} \chi(t) dt \quad (4.6.15)$$

This formula is essentially the *inverse Fourier transform*. As with the other generating functions, the characteristic function can be used to find the moments of X . Moreover, this can be done even when only some of the moments exist.

Suppose again that X has characteristic function χ . If $n \in \mathbb{N}_+$ and $\mathbb{E}(|X^n|) < \infty$. Then

$$\chi(t) = \sum_{k=0}^n \frac{\mathbb{E}(X^k)}{k!} (it)^k + o(t^n) \quad (4.6.16)$$

and therefore $\chi^{(n)}(0) = i^n \mathbb{E}(X^n)$.

Details

Recall that the last term is a generic function that satisfies $o(t^n)/t^n \rightarrow 0$ as $t \rightarrow \infty$.

Next we consider how the characteristic function is changed under some simple transformations of the variables.

Suppose that X has characteristic function χ and that $a, b \in \mathbb{R}$. The characteristic function ψ of $Y = a + bX$ is given by $\psi(t) = e^{iat} \chi(bt)$ for $t \in \mathbb{R}$.

Proof

The proof is just like the one for the MGF: $\psi(t) = \mathbb{E} [e^{it(a+bX)}] = \mathbb{E} (e^{ita} e^{itbX}) = e^{ita} \mathbb{E} [e^{i(tb)X}] = e^{iat} \chi(bt)$ for $t \in \mathbb{R}$.

Suppose that X_1 and X_2 are independent random variables with characteristic functions χ_1 and χ_2 respectively. The characteristic function χ of $Y = X_1 + X_2$ is given by $\chi(t) = \chi_1(t) \chi_2(t)$ for $t \in \mathbb{R}$.

Proof

Again, the proof is just like the one for the MGF:

$$\chi(t) = \mathbb{E} [e^{it(X_1+X_2)}] = \mathbb{E} (e^{itX_1} e^{itX_2}) = \mathbb{E} (e^{itX_1}) \mathbb{E} (e^{itX_2}) = \chi_1(t) \chi_2(t), \quad t \in \mathbb{R} \quad (4.6.17)$$

The characteristic function of a random variable can be obtained from the moment generating function, under the basic existence condition that we saw earlier.

Suppose that X has moment generating function M that satisfies $M(t) < \infty$ for t in an open interval I about 0. Then the characteristic function χ of X satisfies $\chi(t) = M(it)$ for $t \in I$.

The final important property of characteristic functions that we will discuss relates to convergence in distribution. Suppose that (X_1, X_2, \dots) is a sequence of real-valued random with characteristic functions (χ_1, χ_2, \dots) respectively. Since we are only concerned with distributions, the random variables need not be defined on the same probability space.

The Continuity Theorem

1. If the distribution of X_n converges to the distribution of a random variable X as $n \rightarrow \infty$ and X has characteristic function χ , then $\chi_n(t) \rightarrow \chi(t)$ as $n \rightarrow \infty$ for all $t \in \mathbb{R}$.
2. Conversely, if $\chi_n(t)$ converges to a function $\chi(t)$ as $n \rightarrow \infty$ for t in an open interval about 0, and if χ is continuous at 0, then χ is the characteristic function of a random variable X , and the distribution of X_n converges to the distribution of X as $n \rightarrow \infty$.

There are analogous versions of the continuity theorem for probability generating functions and moment generating functions. The continuity theorem can be used to prove the central limit theorem, one of the fundamental theorems of probability. Also, the continuity theorem has a straightforward generalization to distributions on \mathbb{R}^n .

The Joint Characteristic Function

All of the generating functions that we have discussed have multivariate extensions. However, we will discuss the extension only for the characteristic function, the most important and versatile of the generating functions. There are analogous results for the other generating functions. So in this discussion, we assume that (X, Y) is a random vector for our experiment, taking values in \mathbb{R}^2 .

The (joint) *characteristic function* χ of (X, Y) is defined by

$$\chi(s, t) = \mathbb{E} [\exp(isX + itY)], \quad (s, t) \in \mathbb{R}^2 \quad (4.6.18)$$

Once again, the most important fact is that χ completely determines the distribution: two random vectors taking values in \mathbb{R}^2 have the same characteristic function if and only if they have the same distribution.

The joint moments can be obtained from the derivatives of the characteristic function.

Suppose that (X, Y) has characteristic function χ . If $m, n \in \mathbb{N}$ and $\mathbb{E} (|X^m Y^n|) < \infty$ then

$$\chi^{(m,n)}(0, 0) = e^{i(m+n)} \mathbb{E} (X^m Y^n) \quad (4.6.19)$$

The marginal characteristic functions and the characteristic function of the sum can be easily obtained from the joint characteristic function:

Suppose again that (X, Y) has characteristic function χ , and let χ_1 , χ_2 , and χ_+ denote the characteristic functions of X , Y , and $X + Y$, respectively. For $t \in \mathbb{R}$

1. $\chi(t, 0) = \chi_1(t)$
2. $\chi(0, t) = \chi_2(t)$
3. $\chi(t, t) = \chi_+(t)$

Proof

All three results follow immediately from the definitions.

Suppose again that χ_1 , χ_2 , and χ are the characteristic functions of X , Y , and (X, Y) respectively. Then X and Y are independent if and only if $\chi(s, t) = \chi_1(s)\chi_2(t)$ for all $(s, t) \in \mathbb{R}^2$.

Naturally, the results for bivariate characteristic functions have analogies in the general multivariate case. Only the notation is more complicated.

Examples and Applications

As always, be sure to try the computational problems yourself before expanding the solutions and answers in the text.

Dice

Recall that an *ace-six flat die* is a six-sided die for which faces numbered 1 and 6 have probability $\frac{1}{4}$ each, while faces numbered 2, 3, 4, and 5 have probability $\frac{1}{8}$ each. Similarly, a *3-4 flat die* is a six-sided die for which faces numbered 3 and 4 have probability $\frac{1}{4}$ each, while faces numbered 1, 2, 5, and 6 have probability $\frac{1}{8}$ each.

Suppose that an ace-six flat die and a 3-4 flat die are rolled. Use probability generating functions to find the probability density function of the sum of the scores.

Solution

Let X and Y denote the score on the ace-six die and 3-4 flat die, respectively. Then X and Y have PGFs P and Q given by

$$\begin{aligned} P(t) &= \frac{1}{4}t + \frac{1}{8}t^2 + \frac{1}{8}t^3 + \frac{1}{8}t^4 + \frac{1}{8}t^5 + \frac{1}{4}t^6, \quad t \in \mathbb{R} \\ Q(t) &= \frac{1}{8}t + \frac{1}{8}t^2 + \frac{1}{4}t^3 + \frac{1}{4}t^4 + \frac{1}{8}t^5 + \frac{1}{8}t^6, \quad t \in \mathbb{R} \end{aligned}$$

Hence $X + Y$ has PGF PQ . Expanding (a computer algebra program helps) gives

$$P(t)Q(t) = \frac{1}{32}t^2 + \frac{3}{64}t^3 + \frac{3}{32}t^4 + \frac{1}{8}t^5 + \frac{1}{8}t^6 + \frac{5}{32}t^7 + \frac{1}{8}t^8 + \frac{1}{8}t^9 + \frac{3}{32}t^{10} + \frac{3}{64}t^{11} + \frac{1}{32}t^{12}, \quad t \in \mathbb{R} \quad (4.6.20)$$

Thus the PDF f of $X + Y$ is given by $f(2) = f(12) = \frac{1}{32}$, $f(3) = f(11) = \frac{3}{64}$, $f(4) = f(10) = \frac{3}{32}$, $f(5) = f(6) = f(8) = f(9) = \frac{1}{8}$ and $f(7) = \frac{5}{32}$.

Two fair, 6-sided dice are rolled. One has faces numbered $(0, 1, 2, 3, 4, 5)$ and the other has faces numbered $(0, 6, 12, 18, 24, 30)$. Use probability generating functions to find the probability density function of the sum of the scores, and identify the distribution.

Solution

Let X and Y denote the score on the first die and the second die described, respectively. Then X and Y have PGFs P and Q given by

$$\begin{aligned} P(t) &= \frac{1}{6} \sum_{k=0}^5 t^k \quad t \in \mathbb{R} \\ Q(t) &= \frac{1}{6} \sum_{j=0}^5 t^{6j} \quad t \in \mathbb{R} \end{aligned}$$

Hence $X + Y$ has PGF PQ . Simplifying gives

$$P(t)Q(t) = \frac{1}{36} \sum_{j=0}^5 \sum_{k=0}^5 t^{6j+k} = \frac{1}{36} \sum_{n=0}^{35} t^n, \quad t \in \mathbb{R} \quad (4.6.21)$$

Hence $X + Y$ is uniformly distributed on $\{0, 1, 2, \dots, 35\}$

Suppose that random variable Y has probability generating function P given by

$$P(t) = \left(\frac{2}{5}t + \frac{3}{10}t^2 + \frac{1}{5}t^3 + \frac{1}{10}t^4 \right)^5, \quad t \in \mathbb{R} \quad (4.6.22)$$

1. Interpret Y in terms of rolling dice.
2. Use the probability generating function to find the first two factorial moments of Y .
3. Use (b) to find the variance of Y .

Answer

1. A four-sided die has faces numbered $(1, 2, 3, 4)$ with respective probabilities $\left(\frac{2}{5}, \frac{3}{10}, \frac{1}{5}, \frac{1}{10}\right)$. Y is the sum of the scores when the die is rolled 5 times.
2. $\mathbb{E}(Y) = P'(1) = 10$, $\mathbb{E}[Y(Y-1)] = P''(1) = 95$
3. $\text{var}(Y) = 5$

Bernoulli Trials

Suppose X is an indicator random variable with $p = \mathbb{P}(X = 1)$, where $p \in [0, 1]$ is a parameter. Then X has probability generating function $P(t) = 1 - p + pt$ for $t \in \mathbb{R}$.

Proof

$$P(t) = \mathbb{E}(t^X) = t^0(1-p) + t^1p = 1 - p + pt \quad \text{for } t \in \mathbb{R}.$$

Recall that a *Bernoulli trials process* is a sequence (X_1, X_2, \dots) of independent, identically distributed indicator random variables. In the usual language of reliability, X_i denotes the outcome of trial i , where 1 denotes success and 0 denotes failure. The probability of success $p = \mathbb{P}(X_i = 1)$ is the basic parameter of the process. The process is named for Jacob Bernoulli. A separate chapter on the Bernoulli Trials explores this process in more detail.

For $n \in \mathbb{N}_+$, the number of successes in the first n trials is $Y_n = \sum_{i=1}^n X_i$. Recall that this random variable has the binomial distribution with parameters n and p , which has probability density function f_n given by

$$f_n(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (4.6.23)$$

Random variable Y_n has probability generating function P_n given by $P_n(t) = (1 - p + pt)^n$ for $t \in \mathbb{R}$.

Proof

This follows immediately from the [PGF of an indicator variable](#) and the result for [sums of independent variables](#).

Random variable Y_n has the following parameters:

1. $\mathbb{E}[Y_n^{(k)}] = n^{(k)} p^k$
2. $\mathbb{E}(Y_n) = np$
3. $\text{var}(Y_n) = np(1-p)$
4. $\mathbb{P}(Y_n \text{ is even}) = \frac{1}{2}[1 - (1-2p)^n]$

Proof

1. Repeated differentiation gives $P^{(k)}(t) = n^{(k)} p^k (1-p+pt)^{n-k}$. Hence $P^{(k)}(1) = n^{(k)} p^k$, which is $\mathbb{E}[X^{(k)}]$ by the [moment result](#) above.
2. This follows from the [formula for mean](#).
3. This follows from the [formula for variance](#).
4. This follows from the [even value formula](#).

Suppose that U has the binomial distribution with parameters $m \in \mathbb{N}_+$ and $p \in [0, 1]$, V has the binomial distribution with parameters $n \in \mathbb{N}_+$ and $q \in [0, 1]$, and that U and V are independent.

1. If $p = q$ then $U + V$ has the binomial distribution with parameters $m + n$ and p .
2. If $p \neq q$ then $U + V$ does not have a binomial distribution.

Proof

From the [result for sums of independent variables](#) and the [PGF of the binomial distribution](#), note that the probability generating function of $U + V$ is $P(t) = (1 - p + pt)^m (1 - q + qt)^n$ for $t \in \mathbb{R}$.

1. If $p = q$ then $U + V$ has PGF $P(t) = (1 - p + pt)^{m+n}$, which is the PGF of the binomial distribution with parameters $m + n$ and p .
2. On the other hand, if $p \neq q$, the PGF P does not have the functional form of a binomial PGF.

Suppose now that $p \in (0, 1]$. The trial number N of the first success in the sequence of Bernoulli trials has the *geometric distribution* on \mathbb{N}_+ with success parameter p . The probability density function h is given by

$$h(n) = p(1 - p)^{n-1}, \quad n \in \mathbb{N}_+ \quad (4.6.24)$$

The geometric distribution is studied in more detail in the chapter on Bernoulli trials.

Let Q denote the probability generating function of N . Then

1. $Q(t) = \frac{pt}{1 - (1-p)t}$ for $-\frac{1}{1-p} < t < \frac{1}{1-p}$
2. $\mathbb{E}[N^{(k)}] = k! \frac{(1-p)^{k-1}}{p^k}$ for $k \in \mathbb{N}$
3. $\mathbb{E}(N) = \frac{1}{p}$
4. $\text{var}(N) = \frac{1-p}{p^2}$
5. $\mathbb{P}(N \text{ is even}) = \frac{1-p}{2-p}$

Proof

1. Using the formula for the sum of a geometric series,

$$Q(t) = \sum_{n=1}^{\infty} (1-p)^{n-1} p t^n = p t \sum_{n=1}^{\infty} [(1-p)t]^{n-1} = \frac{pt}{1 - (1-p)t}, \quad |(1-p)t| < 1 \quad (4.6.25)$$

2. Repeated differentiation gives $H^{(k)}(t) = k! p (1-p)^{k-1} [1 - (1-p)t]^{-(k+1)}$ and then the result follows from the [inversion formula](#).
3. This follows from (b) and the [formula for mean](#).
4. This follows from (b) and the [formula for variance](#).
5. This follows from [even value formula](#).

The probability that N is even comes up in the alternating coin tossing game with two players.

The Poisson Distribution

Recall that the *Poisson distribution* has probability density function f given by

$$f(n) = e^{-a} \frac{a^n}{n!}, \quad n \in \mathbb{N} \quad (4.6.26)$$

where $a \in (0, \infty)$ is a parameter. The Poisson distribution is named after Simeon Poisson and is widely used to model the number of “random points” in a region of time or space; the parameter is proportional to the size of the region of time or space. The Poisson distribution is studied in more detail in the chapter on the Poisson Process.

Suppose that N has Poisson distribution with parameter $a \in (0, \infty)$. Let P_a denote the probability generating function of N . Then

1. $P_a(t) = e^{a(t-1)}$ for $t \in \mathbb{R}$
2. $\mathbb{E}[N^{(k)}] = a^k$
3. $\mathbb{E}(N) = a$
4. $\text{var}(N) = a$

$$5. \mathbb{P}(N \text{ is even}) = \frac{1}{2}(1 + e^{-2a})$$

Proof

1. Using the exponential series,

$$P_a(t) = \sum_{n=0}^{\infty} e^{-a} \frac{a^n}{n!} t^n = e^{-a} \sum_{n=0}^{\infty} \frac{(at)^n}{n!} = e^{-a} e^{at}, \quad t \in \mathbb{R} \quad (4.6.27)$$

2. Repeated differentiation gives $P_a^{(k)}(t) = e^{a(t-1)} a^k$, so the result follows from [inversion formula](#).

3. This follows from (b) and the [formula for mean](#).

4. This follows from (b) and the [formula for variance](#).

5. This follows from [even value formula](#).

The Poisson family of distributions is closed with respect to sums of independent variables, a very important property.

Suppose that X, Y have Poisson distributions with parameters $a, b \in (0, \infty)$, respectively, and that X and Y are independent. Then $X + Y$ has the Poisson distribution with parameter $a + b$.

Proof

In the notation of the [previous result](#), note that $P_a P_b = P_{a+b}$.

The right distribution function of the Poisson distribution does not have a simple, closed-form expression. The following exercise gives an upper bound.

Suppose that N has the Poisson distribution with parameter $a > 0$. Then

$$\mathbb{P}(N \geq n) \leq e^{n-a} \left(\frac{a}{n}\right)^n, \quad n > a \quad (4.6.28)$$

Proof

The PGF of N is $P(t) = e^{a(t-1)}$ and hence the MGF is $P(e^t) = \exp(ae^t - a)$. From the [Chernov bounds](#) we have

$$\mathbb{P}(N \geq n) \leq e^{-tn} \exp(ae^t - a) = \exp(ae^t - a - tn) \quad (4.6.29)$$

If $n > a$ the expression on the right is minimized when $t = \ln(n/a)$. Substituting gives the upper bound.

The following theorem gives an important convergence result that is explored in more detail in the chapter on the Poisson process.

Suppose that $p_n \in (0, 1)$ for $n \in \mathbb{N}_+$ and that $np_n \rightarrow a \in (0, \infty)$ as $n \rightarrow \infty$. Then the binomial distribution with parameters n and p_n converges to the Poisson distribution with parameter a as $n \rightarrow \infty$.

Proof

Let P_n denote the probability generating function of the binomial distribution with parameters n and p_n . From the [PGF of the binomial distribution](#) we have

$$P_n(t) = [1 + p_n(t-1)]^n = \left[1 + \frac{np_n(t-1)}{n}\right]^n, \quad t \in \mathbb{R} \quad (4.6.30)$$

Using a famous theorem from calculus, $P_n(t) \rightarrow e^{a(t-1)}$ as $n \rightarrow \infty$. But this is the PGF of the Poisson distribution with parameter a , so the result follows from the continuity theorem for PGFs.

The Exponential Distribution

Recall that the *exponential distribution* is a continuous distribution on $[0, \infty)$ with probability density function f given by

$$f(t) = re^{-rt}, \quad t \in (0, \infty) \quad (4.6.31)$$

where $r \in (0, \infty)$ is the *rate parameter*. This distribution is widely used to model failure times and other random times, and in particular governs the time between arrivals in the Poisson model. The exponential distribution is studied in more detail in the chapter on the Poisson Process.

Suppose that T has the exponential distribution with rate parameter $r \in (0, \infty)$ and let M denote the moment generating function of T . Then

1. $M(s) = \frac{r}{r-s}$ for $s \in (-\infty, r)$.
2. $\mathbb{E}(T^n) = n!/r^n$ for $n \in \mathbb{N}$

Proof

1. $M(s) = \mathbb{E}(e^{sT}) = \int_0^\infty e^{st} r e^{-rt} dt = \int_0^\infty r e^{(s-r)t} dt = \frac{r}{r-s}$ for $s < r$.
2. $M^{(n)}(s) = \frac{r n!}{(r-s)^{n+1}}$ for $n \in \mathbb{N}$

Suppose that (T_1, T_2, \dots) is a sequence of independent random variables, each having the exponential distribution with rate parameter $r \in (0, \infty)$. For $n \in \mathbb{N}_+$, the moment generating function M_n of $U_n = \sum_{i=1}^n T_i$ is given by

$$M_n(s) = \left(\frac{r}{r-s} \right)^n, \quad s \in (-\infty, r) \quad (4.6.32)$$

Proof

This follows from the [previous result](#) and the result for [sums of independent variables](#).

Random variable U_n has the *Erlang distribution* with shape parameter n and rate parameter r , named for Agner Erlang. This distribution governs the n th arrival time in the Poisson model. The Erlang distribution is a special case of the gamma distribution and is studied in more detail in the chapter on the Poisson Process.

Uniform Distributions

Suppose that $a, b \in \mathbb{R}$ and $a < b$. Recall that the continuous uniform distribution on the interval $[a, b]$ has probability density function f given by

$$f(x) = \frac{1}{b-a}, \quad x \in [a, b] \quad (4.6.33)$$

The distribution corresponds to selecting a point *at random* from the interval. Continuous uniform distributions arise in geometric probability and a variety of other applied problems.

Suppose that X is uniformly distributed on the interval $[a, b]$ and let M denote the moment generating function of X . Then

1. $M(t) = \frac{e^{bt} - e^{at}}{(b-a)t}$ if $t \neq 0$ and $M(0) = 1$
2. $\mathbb{E}(X^n) = \frac{b^{n+1} - a^{n+1}}{(n+1)(b-a)}$ for $n \in \mathbb{N}$

Proof

1. $M(t) = \int_a^b e^{tx} \frac{1}{b-a} dx = \frac{e^{bt} - e^{at}}{(b-a)t}$ if $t \neq 0$. Trivially $M(0) = 1$
2. This is a case where the MGF is not helpful, and it's much easier to compute the moments directly:

$$\mathbb{E}(X^n) = \int_a^b x^n \frac{1}{b-a} dx = \frac{b^{n+1} - a^{n+1}}{(n+1)(b-a)}$$

Suppose that (X, Y) is uniformly distributed on the triangle $T = \{(x, y) \in \mathbb{R}^2 : 0 \leq x \leq y \leq 1\}$. Compute each of the following:

1. The joint moment generating function of (X, Y) .
2. The moment generating function of X .
3. The moment generating function of Y .
4. The moment generating function of $X + Y$.

Answer

1. $M(s, t) = 2 \frac{e^{s+t} - 1}{s(s+t)} - 2 \frac{e^t - 1}{st}$ if $s \neq 0, t \neq 0$. $M(0, 0) = 1$
2. $M_1(s) = 2 \left(\frac{e^2}{s^2} - \frac{1}{s^2} - \frac{1}{s} \right)$ if $s \neq 0$. $M_1(0) = 1$
3. $M_2(t) = 2 \frac{te^t - e^t + 1}{t^2}$ if $t \neq 0$. $M_2(0) = 1$
4. $M_+(t) = \frac{e^{2t} - 1}{t^2} - 2 \frac{e^t - 1}{t^2}$ if $t \neq 0$. $M_+(0) = 1$

A Bivariate Distribution

Suppose that (X, Y) has probability density function f given by $f(x, y) = x + y$ for $(x, y) \in [0, 1]^2$. Compute each of the following:

1. The joint moment generating function (X, Y) .
2. The moment generating function of X .
3. The moment generating function of Y .
4. The moment generating function of $X + Y$.

Answer

1. $M(s, t) = \frac{e^{s+t}(-2st+s+t)+e^s(st-s-t)+s+t}{s^2t^2}$ if $s \neq 0, t \neq 0$. $M(0, 0) = 1$
2. $M_1(s) = \frac{3se^2-2e^2-s+2}{2s^2}$ if $s \neq 0$. $M_1(0) = 1$
3. $M_2(t) = \frac{3te^2-2e^2-t+2}{2t^2}$ if $t \neq 0$. $M_2(0) = 1$
4. $M_+(t) = \frac{[e^{2t}(1-t)+e^t(t-2)+1]}{t^3}$ if $t \neq 0$. $M_+(0) = 1$

The Normal Distribution

Recall that the *standard normal distribution* is a continuous distribution on \mathbb{R} with probability density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (4.6.34)$$

Normal distributions are widely used to model physical measurements subject to small, random errors and are studied in more detail in the chapter on Special Distributions.

Suppose that Z has the standard normal distribution and let M denote the moment generating function of Z . Then

1. $M(t) = e^{\frac{1}{2}t^2}$ for $t \in \mathbb{R}$
2. $\mathbb{E}(Z^n) = 1 \cdot 3 \cdots (n-1)$ if n is even and $\mathbb{E}(Z^n) = 0$ if n is odd.

Proof

1. First,

$$M(t) = \mathbb{E}(e^{tZ}) = \int_{-\infty}^{\infty} e^{tz} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2} + tz\right) dz \quad (4.6.35)$$

Completing the square in z gives $\exp\left(-\frac{z^2}{2} + tz\right) = \exp\left[\frac{1}{2}t^2 - \frac{1}{2}(z-t)^2\right] = e^{\frac{1}{2}t^2} \exp\left[-\frac{1}{2}(z-t)^2\right]$. hence

$$M(t) = e^{\frac{1}{2}t^2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(z-t)^2\right] dz = e^{\frac{1}{2}t^2} \quad (4.6.36)$$

because the function of z in the last integral is the probability density function for the normal distribution with mean t and variance 1.

2. Note that $M'(t) = tM(t)$. Thus, repeated differentiation gives $M^{(n)}(t) = p_n(t)M(t)$ for $n \in \mathbb{N}$, where p_n is a polynomial of degree n satisfying $p'_{n+1}(t) = tp_n(t) + p'_n(t)$. Since $p_0 = 1$, it's easy to see that p_n has only even or only odd terms, depending on whether n is even or odd, respectively. Thus, $\mathbb{E}(X^n) = p_n(0)$. This is 0 if n is odd, and is the constant term $1 \cdot 3 \cdots (n-1)$ if n is even. Of course, we can also see that the odd order moments must be 0 by symmetry.

More generally, for $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, recall that the normal distribution with mean μ and standard deviation σ is a continuous distribution on \mathbb{R} with probability density function f given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (4.6.37)$$

Moreover, if Z has the standard normal distribution, then $X = \mu + \sigma Z$ has the normal distribution with mean μ and standard deviation σ . Thus, we can easily find the moment generating function of X :

Suppose that X has the normal distribution with mean μ and standard deviation σ . The moment generating function of X is

$$M(t) = \exp\left(\mu t + \frac{1}{2}\sigma^2 t^2\right), \quad t \in \mathbb{R} \quad (4.6.38)$$

Proof

This follows easily the [previous result](#) and the result for [linear transformations](#): $X = \mu + \sigma Z$ where Z has the standard normal distribution. Hence

$$M(t) = \mathbb{E}(e^{tX}) = e^{\mu t} \mathbb{E}(e^{\sigma t Z}) = e^{\mu t} e^{\frac{1}{2}\sigma^2 t^2}, \quad t \in \mathbb{R} \quad (4.6.39)$$

So the normal family of distributions is closed under location-scale transformations. The family is also closed with respect to sums of independent variables:

If X and Y are independent, normally distributed random variables then $X + Y$ has a normal distribution.

Proof

Suppose that X has the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$, and that Y has the normal distribution with mean $\nu \in \mathbb{R}$ and standard deviation $\tau \in (0, \infty)$. By [\(14\)](#), the MGF of $X + Y$ is

$$M_{X+Y}(t) = M_X(t)M_Y(t) = \exp\left(\mu t + \frac{1}{2}\sigma^2 t^2\right) \exp\left(\nu t + \frac{1}{2}\tau^2 t^2\right) = \exp\left[(\mu + \nu)t + \frac{1}{2}(\sigma^2 + \tau^2)t^2\right] \quad (4.6.40)$$

which we recognize as the MGF of the normal distribution with mean $\mu + \nu$ and variance $\sigma^2 + \tau^2$. Of course, we already knew that $\mathbb{E}(X + Y) = \mathbb{E}(X) + \mathbb{E}(Y)$, and since X and Y are independent, $\text{var}(X + Y) = \text{var}(X) + \text{var}(Y)$, so the new information is that the distribution is also normal.

The Pareto Distribution

Recall that the *Pareto distribution* is a continuous distribution on $[1, \infty)$ with probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (4.6.41)$$

where $a \in (0, \infty)$ is the shape parameter. The Pareto distribution is named for Vilfredo Pareto. It is a heavy-tailed distribution that is widely used to model financial variables such as income. The Pareto distribution is studied in more detail in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with shape parameter a , and let M denote the moment generating function of X . Then

1. $\mathbb{E}(X^n) = \frac{a}{a-n}$ if $n < a$ and $\mathbb{E}(X^n) = \infty$ if $n \geq a$
2. $M(t) = \infty$ for $t > 0$

Proof

1. We have seen this computation before. $\mathbb{E}(X^n) = \int_1^\infty x^n \frac{a}{x^{a+1}} dx = \int_1^\infty x^{n-a-1} dx$. The integral evaluates to $\frac{a}{a-n}$ if $n < a$ and ∞ if $n \geq a$.
2. This follows from part (a). Since $X \geq 1$, $M(t)$ is increasing in t . Thus $M(t) \leq 1$ if $t < 0$. If $M(t) < \infty$ for some $t > 0$, then $M(t)$ would be finite for t in an open interval about 0, in which case X would have finite moments of all orders. Of course, it's also easy to see directly from the integral that $M(t) = \infty$ for $t > 0$.

On the other hand, like all distributions on \mathbb{R} , the Pareto distribution has a characteristic function. However, the characteristic function of the Pareto distribution does not have a simple, closed form.

The Cauchy Distribution

Recall that the (standard) *Cauchy distribution* is a continuous distribution on \mathbb{R} with probability density function f given by

$$f(x) = \frac{1}{\pi(1+x^2)}, \quad x \in \mathbb{R} \quad (4.6.42)$$

and is named for Augustin Cauchy. The Cauchy distribution is studied in more generality in the chapter on Special Distributions. The graph of f is known as the *Witch of Agnesi*, named for Maria Agnesi.

Suppose that X has the standard Cauchy distribution, and let M denote the moment generating function of X . Then

1. $\mathbb{E}(X)$ does not exist.
2. $M(t) = \infty$ for $t \neq 0$.

Proof

1. We have seen this computation before. $\int_a^\infty \frac{x}{\pi(1+x^2)} dx = \infty$ and $\int_{-\infty}^a \frac{x}{\pi(1+x^2)} dx = -\infty$ for every $a \in \mathbb{R}$, so $\int_{-\infty}^\infty \frac{x}{\pi(1+x^2)} dx$ does not exist.
2. Note that $\int_0^\infty \frac{e^{tx}}{\pi(1+x^2)} dx = \infty$ if $t \geq 0$ and $\int_{-\infty}^0 \frac{e^{tx}}{\pi(1+x^2)} dx = \infty$ if $t \leq 0$.

Once again, all distributions on \mathbb{R} have characteristic functions, and the standard Cauchy distribution has a particularly simple one.

Let χ denote the characteristic function of X . Then $\chi(t) = e^{-|t|}$ for $t \in \mathbb{R}$.

Proof

The proof of this result requires contour integrals in the complex plane, and is given in the section on the Cauchy distribution in the chapter on special distributions.

Counterexample

For the Pareto distribution, only some of the moments are finite; so course, the moment generating function cannot be finite in an interval about 0. We will now give an example of a distribution for which *all* of the moments are finite, yet still the moment generating function is not finite in any interval about 0. Furthermore, we will see two different distributions that have the same moments of all orders.

Suppose that Z has the standard normal distribution and let $X = e^Z$. The distribution of X is known as the (standard) *lognormal distribution*. The lognormal distribution is studied in more generality in the chapter on Special Distributions. This distribution has finite moments of all orders, but infinite moment generating function.

X has probability density function f given by

$$f(x) = \frac{1}{\sqrt{2\pi}x} \exp\left(-\frac{1}{2}\ln^2(x)\right), \quad x > 0 \quad (4.6.43)$$

1. $\mathbb{E}(X^n) = e^{\frac{1}{2}n^2}$ for $n \in \mathbb{N}$.
2. $\mathbb{E}(e^{tX}) = \infty$ for $t > 0$.

Proof

We use the change of variables theorem. The transformation is $x = e^z$ so the inverse transformation is $z = \ln x$ for $x \in (0, \infty)$ and $z \in \mathbb{R}$. Letting ϕ denote the PDF of Z , it follows that the PDF of X is $f(x) = \phi(z) dz/dx = \phi(\ln x)/x$ for $x > 0$.

1. We use the moment generating function of the standard normal distribution given [above](#): $\mathbb{E}(X^n) = \mathbb{E}(e^{nZ}) = e^{n^2/2}$.
2. Note that

$$\mathbb{E}(e^{tX}) = \mathbb{E}\left[\sum_{n=0}^{\infty} \frac{(tX)^n}{n!}\right] = \sum_{n=0}^{\infty} \frac{\mathbb{E}(X^n)}{n!} t^n = \sum_{n=0}^{\infty} \frac{e^{n^2/2}}{n!} t^n = \infty, \quad t > 0 \quad (4.6.44)$$

The interchange of expected value and sum is justified since X is nonnegative. See the advanced section on properties of the integral in the chapter on Distributions for more details.

Next we construct a different distribution with the same moments as X .

Let h be the function defined by $h(x) = \sin(2\pi \ln x)$ for $x > 0$ and let g be the function defined by $g(x) = f(x)[1 + h(x)]$ for $x > 0$. Then

1. g is a probability density function.
2. If Y has probability density function g then $\mathbb{E}(Y^n) = e^{\frac{1}{2}n^2}$ for $n \in \mathbb{N}$

Proof




Figure 4.6.1: The graphs of f and g , probability density functions for two distributions with the same moments of all orders.

This page titled 4.6: Generating Functions is shared under a CC BY 2.0 license and was authored, remixed, and/or curated by Kyle Siegrist (Random Services) via source content that was edited to the style and standards of the LibreTexts platform.

4.7: Conditional Expected Value

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} the collection of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . Suppose next that X is a random variable taking values in a set S and that Y is a random variable taking values in $T \subseteq \mathbb{R}$. We assume that either Y has a discrete distribution, so that T is countable, or that Y has a continuous distribution so that T is an interval (or perhaps a union of intervals). In this section, we will study the conditional expected value of Y given X , a concept of fundamental importance in probability. As we will see, the expected value of Y given X is the function of X that best approximates Y in the mean square sense. Note that X is a general random variable, not necessarily real-valued, but as usual, we will assume that either X has a discrete distribution, so that S is countable or that X has a continuous distribution on $S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$. In the latter case, S is typically a region defined by inequalities involving elementary functions. We will also assume that all expected values that are mentioned exist (as real numbers).

Basic Theory

Definitions

Note that we can think of (X, Y) as a random variable that takes values in the Cartesian product set $S \times T$. We need recall some basic facts from our work with joint distributions and conditional distributions.

We assume that (X, Y) has joint probability density function f and we let g denote the (marginal) probability density function X . Recall that if Y has a discrete distribution then

$$g(x) = \sum_{y \in T} f(x, y), \quad x \in S \quad (4.7.1)$$

and if Y has a continuous distribution then

$$g(x) = \int_T f(x, y) dy, \quad x \in S \quad (4.7.2)$$

In either case, for $x \in S$, the conditional probability density function of Y given $X = x$ is defined by

$$h(y | x) = \frac{f(x, y)}{g(x)}, \quad y \in T \quad (4.7.3)$$

We are now ready for the basic definitions:

For $x \in S$, the *conditional expected value* of Y given $X = x \in S$ is simply the mean computed relative to the conditional distribution. So if Y has a discrete distribution then

$$E(Y | X = x) = \sum_{y \in T} yh(y | x), \quad x \in S \quad (4.7.4)$$

and if Y has a continuous distribution then

$$E(Y | X = x) = \int_T yh(y | x) dy, \quad x \in S \quad (4.7.5)$$

1. The function $v : S \rightarrow \mathbb{R}$ defined by $v(x) = E(Y | X = x)$ for $x \in S$ is the *regression function* of Y based on X .
2. The random variable $v(X)$ is called the *conditional expected value* of Y given X and is denoted $E(Y | X)$.

Intuitively, we treat X as known, and therefore not random, and we then average Y with respect to the probability distribution that remains. The advanced section on conditional expected value gives a much more general definition that unifies the definitions given here for the various distribution types.

Properties

The most important property of the random variable $E(Y | X)$ is given in the following theorem. In a sense, this result states that $E(Y | X)$ behaves just like Y in terms of other functions of X , and is essentially the only function of X with this property.

The fundamental property

1. $E[r(X)E(Y | X)] = E[r(X)Y]$ for every function $r : S \rightarrow \mathbb{R}$.
2. If $u : S \rightarrow \mathbb{R}$ satisfies $E[r(X)u(X)] = E[r(X)Y]$ for every $r : S \rightarrow \mathbb{R}$ then $\mathbb{P}[u(X) = E(Y | X)] = 1$.

Proof

We give the proof in the continuous case. The discrete case is analogous, with sums replacing integrals.

1. From the change of variables theorem for expected value,

$$E[r(X)E(Y | X)] = \int_S r(x)E(Y | X = x)g(x) dx = \int_S r(x) \left(\int_T yh(y | x) dy \right) g(x) dx \quad (4.7.6)$$

$$= \int_S \int_T r(x)yh(y | x)g(x) dy dx = \int_{S \times T} r(x)yf(x, y) d(x, y) = E[r(X)Y] \quad (4.7.7)$$

2. Suppose that $u_1 : S \rightarrow \mathbb{R}$ and $u_2 : S \rightarrow \mathbb{R}$ satisfy the condition in (b). Define $r : S \rightarrow \mathbb{R}$ by $r(x) = \mathbf{1}_{[u_1(x) > u_2(x)]}$. Then by assumption, $E[r(X)u_1(X)] = E[r(X)Y] = E[r(X)u_2(X)]$. But if $\mathbb{P}[u_1(X) > u_2(X)] > 0$ then $E[r(X)u_1(X)] > E[r(X)u_2(X)]$, a contradiction. Hence we must have $\mathbb{P}[u_1(X) > u_2(X)] = 0$ and by a symmetric argument, $\mathbb{P}[u_1(X) < u_2(X)] = 0$.

Two random variables that are equal with probability 1 are said to be equivalent. We often think of equivalent random variables as being essentially the same object, so the [fundamental property](#) above essentially characterizes $\mathbb{E}(Y | X)$. That is, we can think of $\mathbb{E}(Y | X)$ as any random variable that is a function of X and satisfies this property. Moreover the fundamental property can be used as a definition of conditional expected value, regardless of the type of the distribution of (X, Y) . If you are interested, read the more advanced treatment of conditional expected value.

Suppose that X is also real-valued. Recall that the best linear predictor of Y based on X was characterized by property (a), but with just two functions: $r(x) = 1$ and $r(x) = x$. Thus the characterization in the fundamental property is certainly reasonable, since (as we show below) $\mathbb{E}(Y | X)$ is the best predictor of Y among *all* functions of X , not just linear functions.

The [basic property](#) is also very useful for establishing other properties of conditional expected value. Our first consequence is the fact that Y and $\mathbb{E}(Y | X)$ have the same mean.

$$\mathbb{E}[\mathbb{E}(Y | X)] = \mathbb{E}(Y).$$

Proof

Let r be the constant function 1 in the [basic property](#).

Aside from the theoretical interest, [this theorem](#) is often a good way to compute $\mathbb{E}(Y)$ when we know the conditional distribution of Y given X . We say that we are computing the expected value of Y by *conditioning* on X .

For many basic properties of ordinary expected value, there are analogous results for conditional expected value. We start with two of the most important: every type of expected value must satisfy two critical properties: linearity and monotonicity. In the following two theorems, the random variables Y and Z are real-valued, and as before, X is a general random variable.

Linear Properties

1. $\mathbb{E}(Y + Z | X) = \mathbb{E}(Y | X) + \mathbb{E}(Z | X)$.
2. $\mathbb{E}(cY | X) = c\mathbb{E}(Y | X)$

Proof

1. Note that $\mathbb{E}(Y | X) + \mathbb{E}(Z | X)$ is a function of X . If $r : S \rightarrow \mathbb{R}$ then

$$\mathbb{E}(r(x) [\mathbb{E}(Y | X) + \mathbb{E}(Z | X)]) = \mathbb{E}[r(X)\mathbb{E}(Y | X)] + \mathbb{E}[r(X)\mathbb{E}(Z | X)] = \mathbb{E}[r(X)Y] + \mathbb{E}[r(X)Z] = \mathbb{E}[r(X)(Y + Z)] \quad (4.7.8)$$

Hence the result follows from the [basic property](#).

2. Note that $c\mathbb{E}(Y | X)$ is a function of X . If $r : S \rightarrow \mathbb{R}$ then

$$\mathbb{E}[r(X)c\mathbb{E}(Y | X)] = c\mathbb{E}[r(X)\mathbb{E}(Y | X)] = c\mathbb{E}[r(X)Y] = \mathbb{E}[r(X)(cY)] \quad (4.7.9)$$

Hence the result follows from the [basic property](#)

Part (a) is the additive property and part (b) is the scaling property. The scaling property will be significantly generalized below in (8).

Positive and Increasing Properties

1. If $Y \geq 0$ then $\mathbb{E}(Y | X) \geq 0$.
2. If $Y \leq Z$ then $\mathbb{E}(Y | X) \leq \mathbb{E}(Z | X)$.
3. $|\mathbb{E}(Y | X)| \leq \mathbb{E}(|Y| | X)$

Proof

1. This follows directly from the definition.
2. Note that if $Y \leq Z$ then $Y - Z \leq 0$ so by (a) and linearity,

$$\mathbb{E}(Y - Z | X) = \mathbb{E}(Y | X) - \mathbb{E}(Z | X) \geq 0 \quad (4.7.10)$$

3. Note that $-|Y| \leq Y \leq |Y|$ and hence by (b) and linearity, $-\mathbb{E}(|Y| | X) \leq \mathbb{E}(Y | X) \leq \mathbb{E}(|Y| | X)$.

Our next few properties relate to the idea that $\mathbb{E}(Y | X)$ is the expected value of Y given X . The first property is essentially a restatement of the [fundamental property](#).

If $r : S \rightarrow \mathbb{R}$, then $Y - \mathbb{E}(Y | X)$ and $r(X)$ are uncorrelated.

Proof

Note that $Y - \mathbb{E}(Y | X)$ has mean 0 by the [mean property](#). Hence, by the [basic property](#),

$$\text{cov}[Y - \mathbb{E}(Y | X), r(X)] = \mathbb{E}\{[Y - \mathbb{E}(Y | X)] r(X)\} = \mathbb{E}[Y r(X)] - \mathbb{E}[\mathbb{E}(Y | X) r(X)] = 0 \quad (4.7.11)$$

The next result states that any (deterministic) function of X acts like a constant in terms of the conditional expected value with respect to X .

If $s : S \rightarrow \mathbb{R}$ then

$$\mathbb{E}[s(X)Y | X] = s(X)\mathbb{E}(Y | X) \quad (4.7.12)$$

Proof

Note that $s(X)\mathbb{E}(Y | X)$ is a function of X . If $r : S \rightarrow \mathbb{R}$ then

$$\mathbb{E}[r(X)s(X)\mathbb{E}(Y | X)] = \mathbb{E}[r(X)s(X)Y] \quad (4.7.13)$$

So the result now follow from the [basic property](#).

The following rule generalizes theorem (8) and is sometimes referred to as the *substitution rule* for conditional expected value.

If $s : S \times T \rightarrow \mathbb{R}$ then

$$\mathbb{E}[s(X, Y) | X = x] = \mathbb{E}[s(x, Y) | X = x] \quad (4.7.14)$$

In particular, it follows from (8) that $\mathbb{E}[s(X) | X] = s(X)$. At the opposite extreme, we have the next result: If X and Y are independent, then knowledge of X gives no information about Y and so the conditional expected value with respect to X reduces to the ordinary (unconditional) expected value of Y .

If X and Y are independent then

$$\mathbb{E}(Y | X) = \mathbb{E}(Y) \quad (4.7.15)$$

Proof

Trivially, $\mathbb{E}(Y)$ is a (constant) function of X . If $r : S \rightarrow \mathbb{R}$ then $\mathbb{E}[\mathbb{E}(Y)r(X)] = \mathbb{E}(Y)\mathbb{E}[r(X)] = \mathbb{E}[Yr(X)]$, the last equality by independence. Hence the result follows from the [basic property](#).

Suppose now that Z is real-valued and that X and Y are random variables (all defined on the same probability space, of course). The following theorem gives a consistency condition of sorts. Iterated conditional expected values reduce to a single conditional expected value with respect to the minimum amount of information. For simplicity, we write $\mathbb{E}[Z | X, Y]$ rather than $\mathbb{E}[Z | (X, Y)]$.

Consistency

1. $\mathbb{E}[\mathbb{E}(Z | X, Y) | X] = \mathbb{E}(Z | X)$
2. $\mathbb{E}[\mathbb{E}(Z | X) | X, Y] = \mathbb{E}(Z | X)$

Proof

1. Suppose that X takes values in S and Y takes values in T , so that (X, Y) takes values in $S \times T$. By definition, $\mathbb{E}(Z | X)$ is a function of X . If $r : S \rightarrow \mathbb{R}$ then trivially r can be thought of as a function on $S \times T$ as well. Hence

$$\mathbb{E}[r(X)\mathbb{E}(Z | X)] = \mathbb{E}[r(X)Z] = \mathbb{E}[r(X)\mathbb{E}(Z | X, Y)] \quad (4.7.16)$$

It follows from the [basic property](#) that $\mathbb{E}[\mathbb{E}(Z | X, Y) | X] = \mathbb{E}(Z | X)$.

2. Note that since $\mathbb{E}(Z | X)$ is a function of X , it is trivially a function of (X, Y) . Hence from (8), $\mathbb{E}[\mathbb{E}(Z | X) | X, Y] = \mathbb{E}(Z | X)$.

Finally we show that $\mathbb{E}(Y | X)$ has the same covariance with X as does Y , not surprising since again, $\mathbb{E}(Y | X)$ behaves just like Y in its relations with X .

$$\text{cov}[X, \mathbb{E}(Y | X)] = \text{cov}(X, Y).$$

Proof

$\text{cov}[X, \mathbb{E}(Y | X)] = \mathbb{E}[X\mathbb{E}(Y | X)] - \mathbb{E}(X)\mathbb{E}[\mathbb{E}(Y | X)]$. But $\mathbb{E}[X\mathbb{E}(Y | X)] = \mathbb{E}(XY)$ by [basic property](#), and $\mathbb{E}[\mathbb{E}(Y | X)] = \mathbb{E}(Y)$ by the [mean property](#). Hence $\text{cov}[X, \mathbb{E}(Y | X)] = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y) = \text{cov}(X, Y)$.

Conditional Probability

The conditional probability of an event A , given random variable X (as above), can be defined as a special case of the conditional expected value. As usual, let $\mathbf{1}_A$ denote the indicator random variable of A .

If A is an event, defined

$$\mathbb{P}(A | X) = \mathbb{E}(\mathbf{1}_A | X) \quad (4.7.17)$$

Here is the fundamental property for conditional probability:

The fundamental property

1. $\mathbb{E}[r(X)\mathbb{P}(A | X)] = \mathbb{E}[r(X)\mathbf{1}_A]$ for every function $r : S \rightarrow \mathbb{R}$.
2. If $u : S \rightarrow \mathbb{R}$ and $u(X)$ satisfies $\mathbb{E}[r(X)u(X)] = \mathbb{E}[r(X)\mathbf{1}_A]$ for every function $r : S \rightarrow \mathbb{R}$, then $\mathbb{P}[u(X) = \mathbb{P}(A | X)] = 1$.

For example, suppose that X has a discrete distribution on a countable set S with probability density function g . Then (a) becomes

$$\sum_{x \in S} r(x)\mathbb{P}(A | X = x)g(x) = \sum_{x \in S} r(x)\mathbb{P}(A, X = x) \quad (4.7.18)$$

But this is obvious since $\mathbb{P}(A | X = x) = \mathbb{P}(A, X = x) / \mathbb{P}(X = x)$ and $g(x) = \mathbb{P}(X = x)$. Similarly, if X has a continuous distribution on $S \subseteq \mathbb{R}^n$ then (a) states that

$$\mathbb{E}[r(X)\mathbf{1}_A] = \int_S r(x)\mathbb{P}(A | X = x)g(x) dx \quad (4.7.19)$$

The properties above for conditional expected value, of course, have special cases for conditional probability.

$$\mathbb{P}(A) = \mathbb{E}[\mathbb{P}(A | X)].$$

Proof

This is a direct result of the [mean property](#), since $\mathbb{E}(\mathbf{1}_A) = \mathbb{P}(A)$.

Again, the result in the previous exercise is often a good way to compute $\mathbb{P}(A)$ when we know the conditional probability of A given X . We say that we are computing the probability of A by *conditioning* on X . This is a very compact and elegant version of the conditioning result given first in the section on Conditional Probability in the

chapter on Probability Spaces and later in the section on Discrete Distributions in the Chapter on Distributions.

The following result gives the conditional version of the axioms of probability.

Axioms of probability

1. $\mathbb{P}(A | X) \geq 0$ for every event A .
2. $\mathbb{P}(\Omega | X) = 1$
3. If $\{A_i : i \in I\}$ is a countable collection of disjoint events then $\mathbb{P}(\bigcup_{i \in I} A_i | X) = \sum_{i \in I} \mathbb{P}(A_i | X)$.

Details

There are some technical issues involving the countable additivity property (c). The conditional probabilities are random variables, and so for a given collection $\{A_i : i \in I\}$, the left and right sides are the same with probability 1. We will return to this point in the more advanced section on conditional expected value

From the last result, it follows that other standard probability rules hold for conditional probability given X . These results include

- the complement rule
- the increasing property
- Boole's inequality
- Bonferroni's inequality
- the inclusion-exclusion laws

The Best Predictor

The next result shows that, of all functions of X , $\mathbb{E}(Y | X)$ is closest to Y , in the sense of mean square error. This is fundamentally important in statistical problems where the *predictor vector* X can be observed but not the *response variable* Y . In this subsection and the next, we assume that the real-valued random variables have finite variance.

If $u : S \rightarrow \mathbb{R}$, then

1. $\mathbb{E}([\mathbb{E}(Y | X) - Y]^2) \leq \mathbb{E}([u(X) - Y]^2)$
2. Equality holds in (a) if and only if $u(X) = \mathbb{E}(Y | X)$ with probability 1.

Proof

1. Note that

$$\mathbb{E}([Y - u(X)]^2) = \mathbb{E}([Y - \mathbb{E}(Y | X) + \mathbb{E}(Y | X) - u(X)]^2) \quad (4.7.20)$$

$$= \mathbb{E}([Y - \mathbb{E}(Y | X)]^2) + 2\mathbb{E}([Y - \mathbb{E}(Y | X)][\mathbb{E}(Y | X) - u(X)]) + \mathbb{E}([\mathbb{E}(Y | X) - u(X)]^2) \quad (4.7.21)$$

But $Y - \mathbb{E}(Y | X)$ has mean 0, so the middle term on the right is $2\text{cov}[Y - \mathbb{E}(Y | X), \mathbb{E}(Y | X) - u(X)]$. Moreover, $\mathbb{E}(Y | X) - u(X)$ is a function of X and hence is uncorrelated with $Y - \mathbb{E}(Y | X)$ by the general [uncorrelated property](#). Hence the middle term is 0, so

$$\mathbb{E}([Y - u(X)]^2) = \mathbb{E}([Y - \mathbb{E}(Y | X)]^2) + \mathbb{E}([\mathbb{E}(Y | X) - u(X)]^2) \quad (4.7.22)$$

and therefore $\mathbb{E}([Y - \mathbb{E}(Y | X)]^2) \leq \mathbb{E}([Y - u(X)]^2)$.

2. Equality holds if and only if $\mathbb{E}([\mathbb{E}(Y | X) - u(X)]^2) = 0$, if and only if $\mathbb{P}[u(X) = \mathbb{E}(Y | X)] = 1$.

Suppose now that X is real-valued. In the section on covariance and correlation, we found that the best *linear* predictor of Y given X is

$$L(Y | X) = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)}[X - \mathbb{E}(X)] \quad (4.7.23)$$

On the other hand, $\mathbb{E}(Y | X)$ is the best predictor of Y among *all* functions of X . It follows that if $\mathbb{E}(Y | X)$ happens to be a linear function of X then it must be the case that $\mathbb{E}(Y | X) = L(Y | X)$. However, we will give a direct proof also:

If $\mathbb{E}(Y | X) = a + bX$ for constants a and b then $\mathbb{E}(Y | X) = L(Y | X)$; that is,

1. $b = \text{cov}(X, Y) / \text{var}(X)$
2. $a = \mathbb{E}(Y) - \mathbb{E}(X)\text{cov}(X, Y) / \text{var}(X)$

Proof

First, $\mathbb{E}(Y) = \mathbb{E}[\mathbb{E}(Y | X)] = a + b\mathbb{E}(X)$, so $a = \mathbb{E}(Y) - b\mathbb{E}(X)$. Next, $\text{cov}(X, Y) = \text{cov}[X\mathbb{E}(Y | X)] = \text{cov}(X, a + bX) = b\text{var}(X)$ and therefore $b = \text{cov}(X, Y) / \text{var}(X)$.

Conditional Variance

The conditional variance of Y given X is defined like the ordinary variance, but with all expected values conditioned on X .

The *conditional variance* of Y given X is defined as

$$\text{var}(Y | X) = \mathbb{E}([Y - \mathbb{E}(Y | X)]^2 | X) \quad (4.7.24)$$

Thus, $\text{var}(Y | X)$ is a function of X , and in particular, is a random variable. Our first result is a computational formula that is analogous to the one for standard variance—the variance is the mean of the square minus the square of the mean, but now with all expected values conditioned on X :

$$\text{var}(Y | X) = \mathbb{E}(Y^2 | X) - [\mathbb{E}(Y | X)]^2.$$

Proof

Expanding the square in the definition and using basic properties of conditional expectation, we have

$$\text{var}(Y | X) = \mathbb{E} \left(Y^2 - 2Y\mathbb{E}(Y | X) + [\mathbb{E}(Y | X)]^2 \mid X \right) = \mathbb{E}(Y^2 | X) - 2\mathbb{E}[Y\mathbb{E}(Y | X) | X] + \mathbb{E}([\mathbb{E}(Y | X)]^2 | X) \quad (4.7.25)$$

$$= \mathbb{E}(Y^2 | X) - 2\mathbb{E}(Y | X)\mathbb{E}(Y | X) + [\mathbb{E}(Y | X)]^2 = \mathbb{E}(Y^2 | X) - [\mathbb{E}(Y | X)]^2 \quad (4.7.26)$$

Our next result shows how to compute the ordinary variance of Y by *conditioning on X* .

$$\text{var}(Y) = \mathbb{E}[\text{var}(Y | X)] + \text{var}[\mathbb{E}(Y | X)].$$

Proof

From the previous theorem and properties of conditional expected value we have $\mathbb{E}[\text{var}(Y | X)] = \mathbb{E}(Y^2) - \mathbb{E}([\mathbb{E}(Y | X)]^2)$. But $\mathbb{E}(Y^2) = \text{var}(Y) + [\mathbb{E}(Y)]^2$

and similarly, $\mathbb{E}([\mathbb{E}(Y | X)]^2) = \text{var}[\mathbb{E}(Y | X)] + (\mathbb{E}[\mathbb{E}(Y | X)])^2$. But also, $\mathbb{E}[\mathbb{E}(Y | X)] = \mathbb{E}(Y)$ so substituting we get $\mathbb{E}[\text{var}(Y | X)] = \text{var}(Y) - \text{var}[\mathbb{E}(Y | X)]$.

Thus, the variance of Y is the expected conditional variance plus the variance of the conditional expected value. This result is often a good way to compute $\text{var}(Y)$ when we know the conditional distribution of Y given X . With the help of (21) we can give a formula for the mean square error when $\mathbb{E}(Y | X)$ is used a predictor of Y .

Mean square error

$$\mathbb{E}([Y - \mathbb{E}(Y | X)]^2) = \text{var}(Y) - \text{var}[\mathbb{E}(Y | X)] \quad (4.7.27)$$

Proof

From the definition of [conditional variance](#), and using [mean property](#) and [variance formula](#) we have

$$\mathbb{E}([Y - \mathbb{E}(Y | X)]^2) = \mathbb{E}[\text{var}(Y | X)] = \text{var}(Y) - \text{var}[\mathbb{E}(Y | X)] \quad (4.7.28)$$

Let us return to the study of predictors of the real-valued random variable Y , and compare the three predictors we have studied in terms of mean square error.

Suppose that Y is a real-valued random variable.

1. The best *constant predictor* of Y is $\mathbb{E}(Y)$ with mean square error $\text{var}(Y)$.
2. If X is another real-valued random variable, then the best *linear predictor* of Y given X is

$$L(Y | X) = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)}[X - \mathbb{E}(X)] \quad (4.7.29)$$

with mean square error $\text{var}(Y)[1 - \text{cor}^2(X, Y)]$.

3. If X is a general random variable, then the best *overall predictor* of Y given X is $\mathbb{E}(Y | X)$ with mean square error $\text{var}(Y) - \text{var}[\mathbb{E}(Y | X)]$.

Conditional Covariance

Suppose that Y and Z are real-valued random variables, and that X is a general random variable, all defined on our underlying probability space. Analogous to variance, the conditional covariance of Y and Z given X is defined like the ordinary covariance, but with all expected values conditioned on X .

The *conditional covariance* of Y and Z given X is defined as

$$\text{cov}(Y, Z | X) = \mathbb{E}([Y - \mathbb{E}(Y | X)][Z - \mathbb{E}(Z | X)] | X) \quad (4.7.30)$$

Thus, $\text{cov}(Y, Z | X)$ is a function of X , and in particular, is a random variable. Our first result is a computational formula that is analogous to the one for standard covariance—the covariance is the mean of the product minus the product of the means, but now with all expected values conditioned on X :

$$\text{cov}(Y, Z | X) = \mathbb{E}(YZ | X) - \mathbb{E}(Y | X)\mathbb{E}(Z | X).$$

Proof

Expanding the product in the definition and using basic properties of conditional expectation, we have

$$\begin{aligned} \text{cov}(Y, Z | X) &= \mathbb{E} \left(YZ - Y\mathbb{E}(Z | X) - Z\mathbb{E}(Y | X) + \mathbb{E}(Y | X)\mathbb{E}(Z | X) \mid X \right) = \mathbb{E}(YZ | X) - \mathbb{E}[Y\mathbb{E}(Z | X) | X] - \mathbb{E}[Z\mathbb{E}(Y | X) | X] \\ &\quad + \mathbb{E}[\mathbb{E}(Y | X)\mathbb{E}(Z | X) | X] \end{aligned} \quad (4.7.31)$$

$$= \mathbb{E}(YZ | X) - \mathbb{E}(Y | X)\mathbb{E}(Z | X) - \mathbb{E}(Y | X)\mathbb{E}(Z | X) + \mathbb{E}(Y | X)\mathbb{E}(Z | X) = \mathbb{E}(YZ | X) - \mathbb{E}(Y | X)\mathbb{E}(Z | X) \quad (4.7.32)$$

Our next result shows how to compute the ordinary covariance of Y and Z by *conditioning on X* .

$$\text{cov}(Y, Z) = \mathbb{E}[\text{cov}(Y, Z | X)] + \text{cov}[\mathbb{E}(Y | X), \mathbb{E}(Z | X)].$$

Proof

From (25) and properties of conditional expected value we have

$$\mathbb{E}[\text{cov}(Y, Z | X)] = \mathbb{E}(YZ) - \mathbb{E}[\mathbb{E}(Y | X)\mathbb{E}(Z | X)] \quad (4.7.33)$$

But $\mathbb{E}(YZ) = \text{cov}(Y, Z) + \mathbb{E}(Y)\mathbb{E}(Z)$ and similarly,

$$\mathbb{E}[\mathbb{E}(Y | X)\mathbb{E}(Z | X)] = \text{cov}[\mathbb{E}(Y | X), \mathbb{E}(Z | X)] + \mathbb{E}[\mathbb{E}(Y | X)]\mathbb{E}[\mathbb{E}(Z | X)] \quad (4.7.34)$$

But also, $\mathbb{E}[\mathbb{E}(Y | X)] = \mathbb{E}(Y)$ and $\mathbb{E}[\mathbb{E}(Z | X)] = \mathbb{E}(Z)$ so substituting we get

$$\mathbb{E}[\text{cov}(Y, Z | X)] = \text{cov}(Y, Z) - \text{cov}[\mathbb{E}(Y | X), \mathbb{E}(Z | X)] \quad (4.7.35)$$

Thus, the covariance of Y and Z is the expected conditional covariance plus the covariance of the conditional expected values. This result is often a good way to compute $\text{cov}(Y, Z)$ when we know the conditional distribution of (Y, Z) given X .

Examples and Applications

As always, be sure to try the proofs and computations yourself before reading the ones in the text.

Simple Continuous Distributions

Suppose that (X, Y) has probability density function f defined by $f(x, y) = x + y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

1. Find $L(Y | X)$.
2. Find $\mathbb{E}(Y | X)$.
3. Graph $L(Y | X = x)$ and $\mathbb{E}(Y | X = x)$ as functions of x , on the same axes.
4. Find $\text{var}(Y)$.
5. Find $\text{var}(Y) [1 - \text{cor}^2(X, Y)]$.
6. Find $\text{var}(Y) - \text{var}[\mathbb{E}(Y | X)]$.

Answer

1. $\frac{7}{11} - \frac{1}{11}X$
2. $\frac{3X+2}{6X+3}$
4. $\frac{11}{144} = 0.0764$
5. $\frac{5}{66} = 0.0758$
6. $\frac{1}{12} - \frac{1}{144} \ln 3 = 0.0757$

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 2(x + y)$ for $0 \leq x \leq y \leq 1$.

1. Find $L(Y | X)$.
2. Find $\mathbb{E}(Y | X)$.
3. Graph $L(Y | X = x)$ and $\mathbb{E}(Y | X = x)$ as functions of x , on the same axes.
4. Find $\text{var}(Y)$.
5. Find $\text{var}(Y) [1 - \text{cor}^2(X, Y)]$.
6. Find $\text{var}(Y) - \text{var}[\mathbb{E}(Y | X)]$.

Answer

1. $\frac{26}{43} + \frac{15}{43}X$
2. $\frac{5X^2+5X+2}{9X+3}$
4. $\frac{3}{80} = 0.0375$
5. $\frac{13}{430} = 0.0302$
6. $\frac{1837}{21870} - \frac{512}{6561} \ln(2) = 0.0299$

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 6x^2y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$.

1. Find $L(Y | X)$.
2. Find $\mathbb{E}(Y | X)$.
3. Graph $L(Y | X = x)$ and $\mathbb{E}(Y | X = x)$ as functions of x , on the same axes.
4. Find $\text{var}(Y)$.
5. Find $\text{var}(Y) [1 - \text{cor}^2(X, Y)]$.
6. Find $\text{var}(Y) - \text{var}[\mathbb{E}(Y | X)]$.

Answer

Note that X and Y are independent.

1. $\frac{2}{3}$
2. $\frac{2}{3}$
4. $\frac{1}{18}$
5. $\frac{1}{18}$
6. $\frac{1}{18}$

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$.

1. Find $L(Y | X)$.
2. Find $\mathbb{E}(Y | X)$.

3. Graph $L(Y | X = x)$ and $\mathbb{E}(Y | X = x)$ as functions of x , on the same axes.
4. Find $\text{var}(Y)$.
5. Find $\text{var}(Y) [1 - \text{cor}^2(X, Y)]$.
6. Find $\text{var}(Y) - \text{var}[\mathbb{E}(Y | X)]$.

Answer

1. $\frac{30}{51} + \frac{20}{51}X$
2. $\frac{2(X^2+X+1)}{3(X+1)}$
4. $\frac{5}{252} = 0.0198$
5. $\frac{5}{357} = 0.0140$
6. $\frac{292}{63} - \frac{20}{3}\ln(2) = 0.0139$

Exercises on Basic Properties

Suppose that X , Y , and Z are real-valued random variables with $\mathbb{E}(Y | X) = X^3$ and $\mathbb{E}(Z | X) = \frac{1}{1+X^2}$. Find $\mathbb{E}(Y e^X - Z \sin X | X)$.

Answer

$$X^3 e^X - \frac{\sin X}{1+X^2}$$

Uniform Distributions

As usual, continuous uniform distributions can give us some geometric insight.

Recall first that for $n \in \mathbb{N}_+$, the standard measure on \mathbb{R}^n is

$$\lambda_n(A) = \int_A 1 dx, \quad A \subseteq \mathbb{R}^n \quad (4.7.36)$$

In particular, $\lambda_1(A)$ is the length of $A \subseteq \mathbb{R}$, $\lambda_2(A)$ is the area of $A \subseteq \mathbb{R}^2$, and $\lambda_3(A)$ is the volume of $A \subseteq \mathbb{R}^3$.

Details

Technically λ_n is Lebesgue measure on the measurable subsets of \mathbb{R}^n . The integral representation is valid for the types of sets that occur in applications. In the discussion below, all subsets are assumed to be measurable.

With our usual setup, suppose that X takes values in $S \subseteq \mathbb{R}^n$, Y takes values in $T \subseteq \mathbb{R}$, and that (X, Y) is uniformly distributed on $R \subseteq S \times T \subseteq \mathbb{R}^{n+1}$. So $0 < \lambda_{n+1}(R) < \infty$, and the joint probability density function f of (X, Y) is given by $f(x, y) = 1/\lambda_{n+1}(R)$ for $(x, y) \in R$. Recall that uniform distributions, whether discrete or continuous, always have constant densities. Finally, recall that the cross section of R at $x \in S$ is $T_x = \{y \in T : (x, y) \in R\}$.

In the setting above, suppose that T_x is a bounded interval with midpoint $m(x)$ and length $l(x)$ for each $x \in S$. Then

1. $\mathbb{E}(Y | X) = m(X)$
2. $\text{var}(Y | X) = \frac{1}{12}l^2(X)$

Proof

This follows immediately from the fact that the conditional distribution of Y given $X = x$ is uniformly distributed on T_x for each $x \in S$.

So in particular, the regression curve $x \mapsto \mathbb{E}(Y | X = x)$ follows the midpoints of the cross-sectional intervals.

In each case below, suppose that (X, Y) is uniformly distributed on the give region. Find $\mathbb{E}(Y | X)$ and $\text{var}(Y | X)$

1. The rectangular region $R = [a, b] \times [c, d]$ where $a < b$ and $c < d$.
2. The triangular region $T = \{(x, y) \in \mathbb{R}^2 : -a \leq x \leq y \leq a\}$ where $a > 0$.
3. The circular region $C = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq r\}$ where $r > 0$.

Answer

1. $\mathbb{E}(Y | X) = \frac{1}{2}(c+d)$, $\text{var}(Y | X) = \frac{1}{12}(d-c)^2$. Note that X and Y are independent.
2. $\mathbb{E}(Y | X) = \frac{1}{2}(a+X)$, $\text{var}(Y | X) = \frac{1}{12}(a-X)^2$
3. $\mathbb{E}(Y | X) = 0$, $\text{var}(Y | X) = 4(r^2 - X^2)$

In the bivariate uniform experiment, select each of the following regions. In each case, run the simulation 2000 times and note the relationship between the cloud of points and the graph of the regression function.

1. square
2. triangle
3. circle

Suppose that X is uniformly distributed on the interval $(0, 1)$, and that given X , random variable Y is uniformly distributed on $(0, X)$. Find each of the following:

1. $\mathbb{E}(Y | X)$
2. $\mathbb{E}(Y)$
3. $\text{var}(Y | X)$
4. $\text{var}(Y)$

Answer

1. $\frac{1}{2}X$
2. $\frac{1}{4}$
3. $\frac{1}{12}X^2$
4. $\frac{7}{144}$

The Hypergeometric Distribution

Suppose that a population consists of m objects, and that each object is one of three types. There are a objects of type 1, b objects of type 2, and $m - a - b$ objects of type 0. The parameters a and b are positive integers with $a + b < m$. We sample n objects from the population at random, and without replacement, where $n \in \{0, 1, \dots, m\}$. Denote the number of type 1 and 2 objects in the sample by X and Y , so that the number of type 0 objects in the sample is $n - X - Y$. In the chapter on Distributions, we showed that the joint, marginal, and conditional distributions of X and Y are all hypergeometric—only the parameters change. Here is the relevant result for this section:

In the setting above,

1. $\mathbb{E}(Y | X) = \frac{b}{m-a}(n - X)$
2. $\text{var}(Y | X) = \frac{b(m-a-b)}{(m-a)^2(m-a-1)}(n - X)(m - a - n + X)$
3. $\mathbb{E}([Y - \mathbb{E}(Y | X)]^2) = \frac{n(m-n)b(m-a-b)}{m(m-1)(m-a)}$

Proof

Recall that (X, Y) has the (multivariate) hypergeometric distribution with parameters m, a, b , and n . Marginally, X has the hypergeometric distribution with parameters m, a , and n , and Y has the hypergeometric distribution with parameters m, b , and n . Given $X = x \in \{0, 1, \dots, n\}$, the remaining $n - x$ objects are chosen at random from a population of $m - a$ objects, of which b are type 2 and $m - a - b$ are type 0. Hence, the conditional distribution of Y given $X = x$ is hypergeometric with parameters $m - a, b$, and $n - x$. Parts (a) and (b) then follow from the standard formulas for the mean and variance of the hypergeometric distribution, as functions of the parameters. Part (c) is the mean square error, and in this case can be computed most easily as

$$\text{var}(Y) - \text{var}[\mathbb{E}(Y | X)] = \text{var}(Y) - \left(\frac{b}{m-a}\right)^2 \text{var}(X) = n \frac{b}{m} \frac{m-b}{m} \frac{m-n}{m-1} - \left(\frac{b}{m-a}\right)^2 n \frac{a}{m} \frac{m-a}{m} \frac{m-n}{m-1} \quad (4.7.37)$$

Simplifying gives the result.

Note that $\mathbb{E}(Y | X)$ is a linear function of X and hence $\mathbb{E}(Y | X) = L(Y | X)$.

In a collection of 120 objects, 50 are classified as *good*, 40 as *fair* and 30 as *poor*. A sample of 20 objects is selected at random and without replacement. Let X denote the number of good objects in the sample and Y the number of poor objects in the sample. Find each of the following:

1. $\mathbb{E}(Y | X)$
2. $\text{var}(Y | X)$
3. The predicted value of Y given $X = 8$

Answer

1. $\mathbb{E}(Y | X) = \frac{80}{7} - \frac{4}{7}X$
2. $\text{var}(Y | X) = \frac{4}{1127}(20 - X)(50 + X)$
3. $\frac{48}{7}$

The Multinomial Trials Model

Suppose that we have a sequence of n independent trials, and that each trial results in one of three outcomes, denoted 0, 1, and 2. On each trial, the probability of outcome 1 is p , the probability of outcome 2 is q , so that the probability of outcome 0 is $1 - p - q$. The parameters $p, q \in (0, 1)$ with $p + q < 1$, and of course $n \in \mathbb{N}_+$. Let X denote the number of trials that resulted in outcome 1, Y the number of trials that resulted in outcome 2, so that $n - X - Y$ is the number of trials that resulted in outcome 0. In the chapter on Distributions, we showed that the joint, marginal, and conditional distributions of X and Y are all multinomial—only the parameters change. Here is the relevant result for this section:

In the setting above,

1. $\mathbb{E}(Y | X) = \frac{q}{1-p}(n - X)$
2. $\text{var}(Y | X) = \frac{q(1-p-q)}{(1-p)^2}(n - X)$
3. $\mathbb{E}([Y - \mathbb{E}(Y | X)]^2) = \frac{q(1-p-q)}{1-p}n$

Proof

Recall that (X, Y) has the multinomial distribution with parameters n, p , and q . Marginally, X has the binomial distribution with parameters n and p , and Y has the binomial distribution with parameters n and q . Given $X = x \in \{0, 1, \dots, n\}$, the remaining $n - x$ trials are independent, but with just two outcomes: outcome 2 occurs with probability $q/(1 - p)$ and outcome 0 occurs with probability $1 - q/(1 - p)$. (These are the conditional probabilities of outcomes 2 and 0, respectively, given that outcome 1 did not occur.) Hence the conditional distribution of Y given $X = x$ is binomial with parameters $n - x$ and $q/(1 - p)$. Parts (a) and (b) then follow from the standard formulas for the mean and variance of the binomial distribution, as functions of the parameters. Part (c) is the mean square error and in this case can be computed most easily from

$$\mathbb{E}[\text{var}(Y | X)] = \frac{q(1-p-q)}{(1-p)^2}[n - \mathbb{E}(X)] = \frac{q(1-p-q)}{(1-p)^2}(n - np) = \frac{q(1-p-q)}{1-p}n \quad (4.7.38)$$

Note again that $\mathbb{E}(Y | X)$ is a linear function of X and hence $\mathbb{E}(Y | X) = L(Y | X)$.

Suppose that a fair, 12-sided die is thrown 50 times. Let X denote the number of throws that resulted in a number from 1 to 5, and Y the number of throws that resulted in a number from 6 to 9. Find each of the following:

1. $\mathbb{E}(Y | X)$
2. $\text{var}(Y | X)$
3. The predicted value of Y given $X = 20$

Answer

1. $\mathbb{E}(Y | X) = \frac{4}{7}(50 - X)$
2. $\text{var}(Y | X) = \frac{12}{49}(50 - X)$
3. $\frac{120}{7}$

The Poisson Distribution

Recall that the *Poisson distribution*, named for Simeon Poisson, is widely used to model the number of “random points” in a region of time or space, under certain ideal conditions. The Poisson distribution is studied in more detail in the chapter on the Poisson Process. The Poisson distribution with parameter $r \in (0, \infty)$ has probability density function f defined by

$$f(x) = e^{-r} \frac{r^x}{x!}, \quad x \in \mathbb{N} \quad (4.7.39)$$

The parameter r is the mean and variance of the distribution.

Suppose that X and Y are independent random variables, and that X has the Poisson distribution with parameter $a \in (0, \infty)$ and Y has the Poisson distribution with parameter $b \in (0, \infty)$. Let $N = X + Y$. Then

1. $\mathbb{E}(X | N) = \frac{a}{a+b} N$
2. $\text{var}(X | N) = \frac{ab}{(a+b)^2} N$
3. $\mathbb{E}([X - \mathbb{E}(X | N)]^2) = \frac{ab}{a+b}$

Proof

We have shown before that the distribution of N is also Poisson, with parameter $a + b$, and that the conditional distribution of X given $N = n \in \mathbb{N}$ is binomial with parameters n and $a/(a+b)$. Hence parts (a) and (b) follow from the standard formulas for the mean and variance of the binomial distribution, as functions of the parameters. Part (c) is the mean square error, and in this case can be computed most easily as

$$\mathbb{E}[\text{var}(X | N)] = \frac{ab}{(a+b)^2} \mathbb{E}(N) = \frac{ab}{(a+b)^2} (a+b) = \frac{ab}{a+b} \quad (4.7.40)$$

Once again, $\mathbb{E}(X | N)$ is a linear function of N and so $\mathbb{E}(X | N) = L(X | N)$. If we reverse the roles of the variables, the conditional expected value is trivial from our basic properties:

$$\mathbb{E}(N | X) = \mathbb{E}(X + Y | X) = X + b \quad (4.7.41)$$

Coins and Dice

A pair of fair dice are thrown, and the scores (X_1, X_2) recorded. Let $Y = X_1 + X_2$ denote the sum of the scores and $U = \min\{X_1, X_2\}$ the minimum score. Find each of the following:

1. $\mathbb{E}(Y | X_1)$
2. $\mathbb{E}(U | X_1)$
3. $\mathbb{E}(Y | U)$
4. $\mathbb{E}(X_2 | X_1)$

Answer

1. $\frac{7}{2} + X_1$

2. x	1	2	3	4	5	6
$\mathbb{E}(U X_1 = x)$	1	$\frac{11}{6}$	$\frac{5}{2}$	3	$\frac{10}{3}$	$\frac{7}{2}$

3. u	1	2	3	4	5	6
$\mathbb{E}(Y U = u)$	$\frac{52}{11}$	$\frac{56}{9}$	$\frac{54}{7}$	$\frac{46}{5}$	$\frac{32}{3}$	12

4. $\frac{7}{2}$

A box contains 10 coins, labeled 0 to 9. The probability of heads for coin i is $\frac{i}{9}$. A coin is chosen at random from the box and tossed. Find the probability of heads.

Answer

$$\frac{1}{2}$$

This problem is an example of Laplace's rule of succession, named for Pierre Simon Laplace.

Random Sums of Random Variables

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent and identically distributed real-valued random variables. We will denote the common mean, variance, and moment generating function, respectively, by $\mu = \mathbb{E}(X_i)$, $\sigma^2 = \text{var}(X_i)$, and $G(t) = \mathbb{E}(e^{tX_i})$. Let

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (4.7.42)$$

so that (Y_0, Y_1, \dots) is the partial sum process associated with \mathbf{X} . Suppose now that N is a random variable taking values in \mathbb{N} , independent of \mathbf{X} . Then

$$Y_N = \sum_{i=1}^N X_i \quad (4.7.43)$$

is a random sum of random variables; the terms in the sum are random, and the number of terms is random. This type of variable occurs in many different contexts. For example, N might represent the number of customers who enter a store in a given period of time, and X_i the amount spent by the customer i , so that Y_N is the total revenue of the store during the period.

The conditional and ordinary expected value of Y_N are

1. $\mathbb{E}(Y_N | N) = N\mu$
2. $\mathbb{E}(Y_N) = \mathbb{E}(N)\mu$

Proof

1. Using the substitution rule and the independence of N and \mathbf{X} we have

$$\mathbb{E}(Y_N | N = n) = \mathbb{E}(Y_n | N = n) = \mathbb{E}(Y_n) = \sum_{i=1}^n \mathbb{E}(X_i) = n\mu \quad (4.7.44)$$

so $\mathbb{E}(Y_N | N) = N\mu$.

2. From (a) and conditioning, $\mathbb{E}(Y_N) = \mathbb{E}[\mathbb{E}(Y_N | N)] = \mathbb{E}(N\mu) = \mathbb{E}(N)\mu$.

Wald's equation, named for Abraham Wald, is a generalization of the previous result to the case where N is not necessarily independent of \mathbf{X} , but rather is a *stopping time* for \mathbf{X} . Roughly, this means that the event $N = n$ depends only (X_1, X_2, \dots, X_n) . Wald's equation is discussed in the chapter on Random Samples. An elegant proof of and Wald's equation is given in the chapter on Martingales. The advanced section on stopping times is in the chapter on Probability Measures.

The conditional and ordinary variance of Y_N are

1. $\text{var}(Y_N | N) = N\sigma^2$
2. $\text{var}(Y_N) = \mathbb{E}(N)\sigma^2 + \text{var}(N)\mu^2$

Proof

1. Using the substitution rule, the independence of N and \mathbf{X} , and the fact that \mathbf{X} is an IID sequence, we have

$$\text{var}(Y_N | N = n) = \text{var}(Y_n | N = n) = \text{var}(Y_n) = \sum_{i=1}^n \text{var}(X_i) = n\sigma^2 \quad (4.7.45)$$

so $\text{var}(Y_N | N) = N\sigma^2$.

2. From (a) and the previous result,

$$\text{var}(Y_N) = \mathbb{E}[\text{var}(Y_N | N)] + \text{var}[\mathbb{E}(Y_N | N)] = \mathbb{E}(\sigma^2 N) + \text{var}(\mu N) = \mathbb{E}(N)\sigma^2 + \mu^2 \text{var}(N) \quad (4.7.46)$$

Let H denote the probability generating function of N . The conditional and ordinary moment generating function of Y_N are

1. $\mathbb{E}(e^{tY_N} | N) = [G(t)]^N$
2. $\mathbb{E}(e^{tY_N}) = H(G(t))$

Proof

1. Using the substitution rule, the independence of N and \mathbf{X} , and the fact that \mathbf{X} is an IID sequence, we have

$$\mathbb{E}(e^{tY_N} | N = n) = \mathbb{E}(e^{tY_n} | N = n) = \mathbb{E}(e^{tY_n}) = [G(t)]^n \quad (4.7.47)$$

(Recall that the MGF of the sum of independent variables is the product of the individual MGFs.)

2. From (a) and conditioning, $\mathbb{E}(e^{tY_N}) = \mathbb{E}[\mathbb{E}(e^{tY_N} | N)] = \mathbb{E}(G(t)^N) = H(G(t))$.

Thus the moment generating function of Y_N is $H \circ G$, the composition of the probability generating function of N with the common moment generating function of \mathbf{X} , a simple and elegant result.

In the die-coin experiment, a fair die is rolled and then a fair coin is tossed the number of times showing on the die. Let N denote the die score and Y the number of heads. Find each of the following:

1. The conditional distribution of Y given N .
2. $\mathbb{E}(Y | N)$
3. $\text{var}(Y | N)$
4. $\mathbb{E}(Y_i)$
5. $\text{var}(Y)$

Answer

1. Binomial with parameters N and $p = \frac{1}{2}$
2. $\frac{1}{2}N$
3. $\frac{1}{4}N$
4. $\frac{7}{4}$
5. $\frac{7}{3}$

Run the die-coin experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The number of customers entering a store in a given hour is a random variable with mean 20 and standard deviation 3. Each customer, independently of the others, spends a random amount of money with mean \$50 and standard deviation \$5. Find the mean and standard deviation of the amount of money spent during the hour.

Answer

1. \$1000
2. \$30.82

A coin has a random probability of heads V and is tossed a random number of times N . Suppose that V is uniformly distributed on $[0, 1]$; N has the Poisson distribution with parameter $a > 0$; and V and N are independent. Let Y denote the number of heads. Compute the following:

1. $\mathbb{E}(Y \mid N, V)$
2. $\mathbb{E}(Y \mid N)$
3. $\mathbb{E}(Y \mid V)$
4. $\mathbb{E}(Y)$
5. $\text{var}(Y \mid N, V)$
6. $\text{var}(Y)$

Answer

1. NV
2. $\frac{1}{2}N$
3. aV
4. $\frac{1}{2}a$
5. $NV(1 - V)$
6. $\frac{1}{12}a^2 + \frac{1}{2}a$

Mixtures of Distributions

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of real-valued random variables. Denote the mean, variance, and moment generating function of X_i by $\mu_i = \mathbb{E}(X_i)$, $\sigma_i^2 = \text{var}(X_i)$, and $M_i(t) = \mathbb{E}(e^{tX_i})$, for $i \in \mathbb{N}_+$. Suppose also that N is a random variable taking values in \mathbb{N}_+ , independent of \mathbf{X} . Denote the probability density function of N by $p_n = \mathbb{P}(N = n)$ for $n \in \mathbb{N}_+$. The distribution of the random variable X_N is a mixture of the distributions of $\mathbf{X} = (X_1, X_2, \dots)$, with the distribution of N as the mixing distribution.

The conditional and ordinary expected value of X_N are

1. $\mathbb{E}(X_N \mid N) = \mu_N$
2. $\mathbb{E}(X_N) = \sum_{n=1}^{\infty} p_n \mu_n$

Proof

1. Using the substitution rule and the independence of N and \mathbf{X} , we have $\mathbb{E}(X_N \mid N = n) = \mathbb{E}(X_n \mid N = n) = \mathbb{E}(X_n) = \mu_n$
2. From (a) and the conditioning rule,

$$\mathbb{E}(X_N) = \mathbb{E}[\mathbb{E}(X_N)] = \mathbb{E}(\mu_N) = \sum_{n=1}^{\infty} p_n \mu_n \quad (4.7.48)$$

The conditional and ordinary variance of X_N are

1. $\text{var}(X_N \mid N) = \sigma_N^2$
2. $\text{var}(X_N) = \sum_{n=1}^{\infty} p_n (\sigma_n^2 + \mu_n^2) - (\sum_{n=1}^{\infty} p_n \mu_n)^2$.

Proof

1. Using the substitution rule and the independence of N and \mathbf{X} , we have $\text{var}(X_N \mid N = n) = \text{var}(X_n \mid N = n) = \text{var}(X_n) = \sigma_n^2$
2. From (a) we have

$$\text{var}(X_N) = \mathbb{E}[\text{var}(X_N \mid N)] + \text{var}[\mathbb{E}(X_N \mid N)] = \mathbb{E}(\sigma_N^2) + \text{var}(\mu_N) = \mathbb{E}(\sigma_N^2) + \mathbb{E}(\mu_N^2) - [\mathbb{E}(\mu_N)]^2 \quad (4.7.49)$$

$$= \sum_{n=1}^{\infty} p_n \sigma_n^2 + \sum_{n=1}^{\infty} p_n \mu_n^2 - \left(\sum_{n=1}^{\infty} p_n \mu_n \right)^2 \quad (4.7.50)$$

The conditional and ordinary moment generating function of X_N are

1. $\mathbb{E}(e^{tX_N} \mid N) = M_N(t)$
2. $\mathbb{E}(e^{tX_N}) = \sum_{i=1}^{\infty} p_i M_i(t)$.

Proof

1. Using the substitution rule and the independence of N and \mathbf{X} , we have $\mathbb{E}(e^{tX_N} \mid N = n) = \mathbb{E}(e^{tX_n} \mid N = n) = \mathbb{E}(e^{tX_n}) = M_n(t)$

2. From (a) and the conditioning rule, $\mathbb{E}(e^{tX_N}) = \mathbb{E}[\mathbb{E}(e^{tX_N} | N)] = \mathbb{E}[M_N(t)] = \sum_{n=1}^{\infty} p_n M_n(t)$

In the coin-die experiment, a biased coin is tossed with probability of heads $\frac{1}{3}$. If the coin lands tails, a fair die is rolled; if the coin lands heads, an ace-six flat die is rolled (faces 1 and 6 have probability $\frac{1}{4}$ each, and faces 2, 3, 4, 5 have probability $\frac{1}{8}$ each). Find the mean and standard deviation of the die score.

Answer

1. $\frac{7}{2}$
2. 1.8634

Run the coin-die experiment 1000 times and note the apparent convergence of the empirical mean and standard deviation to the distribution mean and standard deviation.

This page titled [4.7: Conditional Expected Value](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.8: Expected Value and Covariance Matrices

The main purpose of this section is a discussion of expected value and covariance for random matrices and vectors. These topics are somewhat specialized, but are particularly important in multivariate statistical models and for the multivariate normal distribution. This section requires some prerequisite knowledge of linear algebra.

We assume that the various indices m, n, p, k that occur in this section are positive integers. Also we assume that expected values of real-valued random variables that we reference exist as real numbers, although extensions to cases where expected values are ∞ or $-\infty$ are straightforward, as long as we avoid the dreaded indeterminate form $\infty - \infty$.

Basic Theory

Linear Algebra

We will follow our usual convention of denoting random variables by upper case letters and nonrandom variables and constants by lower case letters. In this section, that convention leads to notation that is a bit nonstandard, since the objects that we will be dealing with are vectors and matrices. On the other hand, the notation we will use works well for illustrating the similarities between results for random matrices and the corresponding results in the one-dimensional case. Also, we will try to be careful to explicitly point out the underlying spaces where various objects live.

Let $\mathbb{R}^{m \times n}$ denote the space of all $m \times n$ matrices of real numbers. The (i, j) entry of $\mathbf{a} \in \mathbb{R}^{m \times n}$ is denoted a_{ij} for $i \in \{1, 2, \dots, m\}$ and $j \in \{1, 2, \dots, n\}$. We will identify \mathbb{R}^n with $\mathbb{R}^{n \times 1}$, so that an ordered n -tuple can also be thought of as an $n \times 1$ column vector. The transpose of a matrix $\mathbf{a} \in \mathbb{R}^{m \times n}$ is denoted \mathbf{a}^T —the $n \times m$ matrix whose (i, j) entry is the (j, i) entry of \mathbf{a} . Recall the definitions of matrix addition, scalar multiplication, and matrix multiplication. Recall also the standard inner product (or dot product) of $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$:

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x} \cdot \mathbf{y} = \mathbf{x}^T \mathbf{y} = \sum_{i=1}^n x_i y_i \quad (4.8.1)$$

The outer product of \mathbf{x} and \mathbf{y} is $\mathbf{x}\mathbf{y}^T$, the $n \times n$ matrix whose (i, j) entry is $x_i y_j$. Note that the inner product is the trace (sum of the diagonal entries) of the outer product. Finally recall the standard norm on \mathbb{R}^n , given by

$$\|\mathbf{x}\| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle} = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2} \quad (4.8.2)$$

Recall that inner product is *bilinear*, that is, linear (preserving addition and scalar multiplication) in each argument separately. As a consequence, for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,

$$\|\mathbf{x} + \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 + 2\langle \mathbf{x}, \mathbf{y} \rangle \quad (4.8.3)$$

Expected Value of a Random Matrix

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So to review, Ω is the set of outcomes, \mathcal{F} the collection of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . It's natural to define the expected value of a random matrix in a component-wise manner.

Suppose that \mathbf{X} is an $m \times n$ matrix of real-valued random variables, whose (i, j) entry is denoted X_{ij} . Equivalently, \mathbf{X} is as a random $m \times n$ matrix, that is, a random variable with values in $\mathbb{R}^{m \times n}$. The *expected value* $\mathbb{E}(\mathbf{X})$ is defined to be the $m \times n$ matrix whose (i, j) entry is $\mathbb{E}(X_{ij})$, the expected value of X_{ij} .

Many of the basic properties of expected value of random variables have analogous results for expected value of random matrices, with matrix operation replacing the ordinary ones. Our first two properties are the critically important *linearity properties*. The first part is the *additive property*—the expected value of a sum is the sum of the expected values.

$\mathbb{E}(\mathbf{X} + \mathbf{Y}) = \mathbb{E}(\mathbf{X}) + \mathbb{E}(\mathbf{Y})$ if \mathbf{X} and \mathbf{Y} are random $m \times n$ matrices.

Proof

This is true by definition of the matrix expected value and the ordinary additive property. Note that $\mathbb{E}(X_{ij} + Y_{ij}) = \mathbb{E}(X_{ij}) + \mathbb{E}(Y_{ij})$. The left side is the (i, j) entry of $\mathbb{E}(\mathbf{X} + \mathbf{Y})$ and the right side is the (i, j) entry of $\mathbb{E}(\mathbf{X}) + \mathbb{E}(\mathbf{Y})$.

The next part of the linearity properties is the *scaling property*—a nonrandom matrix factor can be pulled out of the expected value.

Suppose that \mathbf{X} is a random $n \times p$ matrix.

1. $\mathbb{E}(\mathbf{a}\mathbf{X}) = \mathbf{a}\mathbb{E}(\mathbf{X})$ if $\mathbf{a} \in \mathbb{R}^{m \times n}$.
2. $\mathbb{E}(\mathbf{X}\mathbf{a}) = \mathbb{E}(\mathbf{X})\mathbf{a}$ if $\mathbf{a} \in \mathbb{R}^{p \times n}$.

Proof

1. By the ordinary linearity and scaling properties, $\mathbb{E}\left(\sum_{j=1}^n a_{ij} X_{jk}\right) = \sum_{j=1}^n a_{ij} \mathbb{E}(X_{jk})$. The left side is the (i, k) entry of $\mathbb{E}(\mathbf{a}\mathbf{X})$ and the right side is the (i, k) entry of $\mathbf{a}\mathbb{E}(\mathbf{X})$.
2. The proof is similar to (a).

Recall that for independent, real-valued variables, the expected value of the product is the product of the expected values. Here is the analogous result for random matrices.

$\mathbb{E}(\mathbf{X}\mathbf{Y}) = \mathbb{E}(\mathbf{X})\mathbb{E}(\mathbf{Y})$ if \mathbf{X} is a random $m \times n$ matrix, \mathbf{Y} is a random $n \times p$ matrix, and \mathbf{X} and \mathbf{Y} are independent.

Proof

By the ordinary linearity properties and by the independence assumption,

$$\mathbb{E} \left(\sum_{j=1}^n X_{ij} Y_{jk} \right) = \sum_{j=1}^n \mathbb{E} (X_{ij} Y_{jk}) = \sum_{j=1}^n \mathbb{E} (X_{ij}) \mathbb{E} (Y_{jk}) \quad (4.8.4)$$

The left side is the (i, k) entry of $\mathbb{E}(\mathbf{XY})$ and the right side is the (i, k) entry of $\mathbb{E}(\mathbf{X})\mathbb{E}(\mathbf{Y})$.

Actually the previous result holds if \mathbf{X} and \mathbf{Y} are simply uncorrelated in the sense that X_{ij} and Y_{jk} are uncorrelated for each $i \in \{1, \dots, m\}$, $j \in \{1, 2, \dots, n\}$ and $k \in \{1, 2, \dots, p\}$. We will study covariance of random vectors in the next subsection.

Covariance Matrices

Our next goal is to define and study the covariance of two random vectors.

Suppose that \mathbf{X} is a random vector in \mathbb{R}^m and \mathbf{Y} is a random vector in \mathbb{R}^n .

1. The *covariance matrix* of \mathbf{X} and \mathbf{Y} is the $m \times n$ matrix $\text{cov}(\mathbf{X}, \mathbf{Y})$ whose (i, j) entry is $\text{cov}(X_i, Y_j)$ the ordinary covariance of X_i and Y_j .
2. Assuming that the coordinates of \mathbf{X} and \mathbf{Y} have positive variance, the *correlation matrix* of \mathbf{X} and \mathbf{Y} is the $m \times n$ matrix $\text{cor}(\mathbf{X}, \mathbf{Y})$ whose (i, j) entry is $\text{cor}(X_i, Y_j)$, the ordinary correlation of X_i and Y_j

Many of the standard properties of covariance and correlation for real-valued random variables have extensions to random vectors. For the following three results, \mathbf{X} is a random vector in \mathbb{R}^m and \mathbf{Y} is a random vector in \mathbb{R}^n .

$$\text{cov}(\mathbf{X}, \mathbf{Y}) = \mathbb{E} \left([\mathbf{X} - \mathbb{E}(\mathbf{X})] [\mathbf{Y} - \mathbb{E}(\mathbf{Y})]^T \right)$$

Proof

By the definition of the expected value of a random vector and by the definition of matrix multiplication, the (i, j) entry of $[\mathbf{X} - \mathbb{E}(\mathbf{X})] [\mathbf{Y} - \mathbb{E}(\mathbf{Y})]^T$ is simply $[X_i - \mathbb{E}(X_i)] [Y_j - \mathbb{E}(Y_j)]$. The expected value of this entry is $\text{cov}(X_i, Y_j)$, which in turn, is the (i, j) entry of $\text{cov}(\mathbf{X}, \mathbf{Y})$.

Thus, the covariance of \mathbf{X} and \mathbf{Y} is the expected value of the outer product of $\mathbf{X} - \mathbb{E}(\mathbf{X})$ and $\mathbf{Y} - \mathbb{E}(\mathbf{Y})$. Our next result is the computational formula for covariance: the expected value of the outer product of \mathbf{X} and \mathbf{Y} minus the outer product of the expected values.

$$\text{cov}(\mathbf{X}, \mathbf{Y}) = \mathbb{E}(\mathbf{XY}^T) - \mathbb{E}(\mathbf{X})[\mathbb{E}(\mathbf{Y})]^T.$$

Proof

The (i, j) entry of $\mathbb{E}(\mathbf{XY}^T) - \mathbb{E}(\mathbf{X})[\mathbb{E}(\mathbf{Y})]^T$ is $\mathbb{E}(X_i Y_j) - \mathbb{E}(X_i) \mathbb{E}(Y_j)$, which by the standard computational formula, is $\text{cov}(X_i, Y_j)$, which in turn is the (i, j) entry of $\text{cov}(\mathbf{X}, \mathbf{Y})$.

The next result is the matrix version of the symmetry property.

$$\text{cov}(\mathbf{Y}, \mathbf{X}) = [\text{cov}(\mathbf{X}, \mathbf{Y})]^T.$$

Proof

The (i, j) entry of $\text{cov}(\mathbf{X}, \mathbf{Y})$ is $\text{cov}(X_i, Y_j)$, which is the (j, i) entry of $\text{cov}(\mathbf{Y}, \mathbf{X})$.

In the following result, $\mathbf{0}$ denotes the $m \times n$ zero matrix.

$\text{cov}(\mathbf{X}, \mathbf{Y}) = \mathbf{0}$ if and only if $\text{cov}(X_i, Y_j) = 0$ for each i and j , so that each coordinate of \mathbf{X} is uncorrelated with each coordinate of \mathbf{Y} .

Proof

This follows immediately from the definition of $\text{cov}(\mathbf{X}, \mathbf{Y})$.

Naturally, when $\text{cov}(\mathbf{X}, \mathbf{Y}) = \mathbf{0}$, we say that the random vectors \mathbf{X} and \mathbf{Y} are *uncorrelated*. In particular, if the random vectors are independent, then they are uncorrelated. The following results establish the *bi-linear* properties of covariance.

The additive properties.

1. $\text{cov}(\mathbf{X} + \mathbf{Y}, \mathbf{Z}) = \text{cov}(\mathbf{X}, \mathbf{Z}) + \text{cov}(\mathbf{Y}, \mathbf{Z})$ if \mathbf{X} and \mathbf{Y} are random vectors in \mathbb{R}^m and \mathbf{Z} is a random vector in \mathbb{R}^n .
2. $\text{cov}(\mathbf{X}, \mathbf{Y} + \mathbf{Z}) = \text{cov}(\mathbf{X}, \mathbf{Y}) + \text{cov}(\mathbf{X}, \mathbf{Z})$ if \mathbf{X} is a random vector in \mathbb{R}^m , and \mathbf{Y} and \mathbf{Z} are random vectors in \mathbb{R}^n .

Proof

1. From the ordinary additive property of covariance, $\text{cov}(X_i + Y_i, Z_j) = \text{cov}(X_i, Z_j) + \text{cov}(Y_i, Z_j)$. The left side is the (i, j) entry of $\text{cov}(\mathbf{X} + \mathbf{Y}, \mathbf{Z})$ and the right side is the (i, j) entry of $\text{cov}(\mathbf{X}, \mathbf{Z}) + \text{cov}(\mathbf{Y}, \mathbf{Z})$.
2. The proof is similar to (a), using the additivity of covariance in the second argument.

The scaling properties

1. $\text{cov}(\mathbf{aX}, \mathbf{Y}) = \mathbf{a} \text{cov}(\mathbf{X}, \mathbf{Y})$ if \mathbf{X} is a random vector in \mathbb{R}^n , \mathbf{Y} is a random vector in \mathbb{R}^p , and $\mathbf{a} \in \mathbb{R}^{m \times n}$.
2. $\text{cov}(\mathbf{X}, \mathbf{aY}) = \text{cov}(\mathbf{X}, \mathbf{Y}) \mathbf{a}^T$ if \mathbf{X} is a random vector in \mathbb{R}^m , \mathbf{Y} is a random vector in \mathbb{R}^n , and $\mathbf{a} \in \mathbb{R}^{k \times n}$.

Proof

1. Using the ordinary linearity properties of covariance in the first argument, we have

$$\text{cov}\left(\sum_{j=1}^n a_{ij}X_j, Y_k\right) = \sum_{j=1}^n a_{ij}\text{cov}(X_j, Y_k) \quad (4.8.5)$$

The left side is the (i, k) entry of $\text{cov}(\mathbf{aX}, \mathbf{Y})$ and the right side is the (i, k) entry of $\mathbf{a}\text{cov}(\mathbf{X}, \mathbf{Y})$.
2. The proof is similar to (a), using the linearity of covariance in the second argument.

Variance-Covariance Matrices

Suppose that \mathbf{X} is a random vector in \mathbb{R}^n . The covariance matrix of \mathbf{X} with itself is called the *variance-covariance matrix* of \mathbf{X} :

$$\text{vc}(\mathbf{X}) = \text{cov}(\mathbf{X}, \mathbf{X}) = \mathbb{E}\left([\mathbf{X} - \mathbb{E}(\mathbf{X})][\mathbf{X} - \mathbb{E}(\mathbf{X})]^T\right) \quad (4.8.6)$$

Recall that for an ordinary real-valued random variable X , $\text{var}(X) = \text{cov}(X, X)$. Thus the variance-covariance matrix of a random vector in some sense plays the same role that variance does for a random variable.

$\text{vc}(\mathbf{X})$ is a symmetric $n \times n$ matrix with $(\text{var}(X_1), \text{var}(X_2), \dots, \text{var}(X_n))$ on the diagonal.

Proof

Recall that $\text{cov}(X_i, X_j) = \text{cov}(X_j, X_i)$. Also, the (i, i) entry of $\text{vc}(\mathbf{X})$ is $\text{cov}(X_i, X_i) = \text{var}(X_i)$.

The following result is the formula for the variance-covariance matrix of a sum, analogous to the formula for the variance of a sum of real-valued variables.

$\text{vc}(\mathbf{X} + \mathbf{Y}) = \text{vc}(\mathbf{X}) + \text{cov}(\mathbf{X}, \mathbf{Y}) + \text{cov}(\mathbf{Y}, \mathbf{X}) + \text{vc}(\mathbf{Y})$ if \mathbf{X} and \mathbf{Y} are random vectors in \mathbb{R}^n .

Proof

This follows from the additive property of covariance:

$$\text{vc}(\mathbf{X} + \mathbf{Y}) = \text{cov}(\mathbf{X} + \mathbf{Y}, \mathbf{X} + \mathbf{Y}) = \text{cov}(\mathbf{X}, \mathbf{X}) + \text{cov}(\mathbf{X}, \mathbf{Y}) + \text{cov}(\mathbf{Y}, \mathbf{X}) + \text{cov}(\mathbf{Y}, \mathbf{Y}) \quad (4.8.7)$$

Recall that $\text{var}(aX) = a^2\text{var}(X)$ if X is a real-valued random variable and $a \in \mathbb{R}$. Here is the analogous result for the variance-covariance matrix of a random vector.

$\text{vc}(\mathbf{aX}) = \mathbf{a}\text{vc}(\mathbf{X})\mathbf{a}^T$ if \mathbf{X} is a random vector in \mathbb{R}^n and $\mathbf{a} \in \mathbb{R}^{m \times n}$.

Proof

This follows from the scaling property of covariance:

$$\text{vc}(\mathbf{aX}) = \text{cov}(\mathbf{aX}, \mathbf{aX}) = \mathbf{a}\text{cov}(\mathbf{X}, \mathbf{X})\mathbf{a}^T \quad (4.8.8)$$

Recall that if X is a random variable, then $\text{var}(X) \geq 0$, and $\text{var}(X) = 0$ if and only if X is a constant (with probability 1). Here is the analogous result for a random vector:

Suppose that \mathbf{X} is a random vector in \mathbb{R}^n .

1. $\text{vc}(\mathbf{X})$ is either positive semi-definite or positive definite.
2. $\text{vc}(\mathbf{X})$ is positive semi-definite but not positive definite if and only if there exists $\mathbf{a} \in \mathbb{R}^n$ and $c \in \mathbb{R}$ such that, with probability 1, $\mathbf{a}^T \mathbf{X} = \sum_{i=1}^n a_i X_i = c$

Proof

1. From the previous result, $0 \leq \text{var}(\mathbf{a}^T \mathbf{X}) = \text{vc}(\mathbf{a}^T \mathbf{X}) = \mathbf{a}^T \text{vc}(\mathbf{X}) \mathbf{a}$ for every $\mathbf{a} \in \mathbb{R}^n$. Thus, by definition, $\text{vc}(\mathbf{X})$ is either positive semi-definite or positive definite.
2. In light of (a), $\text{vc}(\mathbf{X})$ is positive semi-definite but not positive definite if and only if there exists $\mathbf{a} \in \mathbb{R}^n$ such that $\mathbf{a}^T \text{vc}(\mathbf{X}) \mathbf{a} = \text{var}(\mathbf{a}^T \mathbf{X}) = 0$. But in turn, this is true if and only if $\mathbf{a}^T \mathbf{X}$ is constant with probability 1.

Recall that since $\text{vc}(\mathbf{X})$ is either positive semi-definite or positive definite, the eigenvalues and the determinant of $\text{vc}(\mathbf{X})$ are nonnegative. Moreover, if $\text{vc}(\mathbf{X})$ is positive semi-definite but not positive definite, then one of the coordinates of \mathbf{X} can be written as a linear transformation of the other coordinates (and hence can usually be eliminated in the underlying model). By contrast, if $\text{vc}(\mathbf{X})$ is positive definite, then this cannot happen; $\text{vc}(\mathbf{X})$ has positive eigenvalues and determinant and is invertible.

Best Linear Predictor

Suppose that \mathbf{X} is a random vector in \mathbb{R}^m and that \mathbf{Y} is a random vector in \mathbb{R}^n . We are interested in finding the function of \mathbf{X} of the form $\mathbf{a} + \mathbf{bX}$, where $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^{n \times m}$, that is closest to \mathbf{Y} in the mean square sense. Functions of this form are analogous to linear functions in the single variable case. However, unless $\mathbf{a} = \mathbf{0}$, such functions are not *linear transformations* in the sense of linear algebra, so the correct term is *affine function* of \mathbf{X} . This problem is of fundamental importance in statistics when random vector \mathbf{X} , the *predictor vector* is observable, but not random vector \mathbf{Y} , the *response vector*. Our discussion here generalizes the one-dimensional case, when X and Y are random variables. That problem was solved in the section on Covariance and Correlation. We will assume that $\text{vc}(\mathbf{X})$ is positive definite, so that $\text{vc}(\mathbf{X})$ is invertible, and none of the coordinates of \mathbf{X} can be written as an affine function of the other coordinates. We write $\text{vc}^{-1}(\mathbf{X})$ for the inverse instead of the clunkier $[\text{vc}(\mathbf{X})]^{-1}$.

As with the single variable case, the solution turns out to be the affine function that has the same expected value as \mathbf{Y} , and whose covariance with \mathbf{X} is the same as that of \mathbf{Y} .

Define $L(\mathbf{Y} | \mathbf{X}) = \mathbb{E}(\mathbf{Y}) + \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})[\mathbf{X} - \mathbb{E}(\mathbf{X})]$. Then $L(\mathbf{Y} | \mathbf{X})$ is the only affine function of \mathbf{X} in \mathbb{R}^n satisfying

1. $\mathbb{E}[L(\mathbf{Y} | \mathbf{X})] = \mathbb{E}(\mathbf{Y})$
2. $\text{cov}[L(\mathbf{Y} | \mathbf{X}), \mathbf{X}] = \text{cov}(\mathbf{Y}, \mathbf{X})$

Proof

From linearity,

$$\mathbb{E}[L(\mathbf{Y} | \mathbf{X})] = E(\mathbf{Y}) + \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})[\mathbb{E}(\mathbf{X}) - \mathbb{E}(\mathbf{X})] = 0 \quad (4.8.9)$$

From linearity and the fact that a constant vector is independent (and hence uncorrelated) with any random vector,

$$\text{cov}[L(\mathbf{Y} | \mathbf{X}), \mathbf{X}] = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{X}, \mathbf{X}) = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{vc}(\mathbf{X}) = \text{cov}(\mathbf{Y}, \mathbf{X}) \quad (4.8.10)$$

Conversely, suppose that $\mathbf{U} = \mathbf{a} + \mathbf{b}\mathbf{X}$ for some $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^{m \times n}$, and that $\mathbb{E}(\mathbf{U}) = \mathbb{E}(\mathbf{Y})$ and $\text{cov}(\mathbf{U}, \mathbf{X}) = \text{cov}(\mathbf{Y}, \mathbf{X})$. From the second equation, again using linearity and the uncorrelated property of constant vectors, we get $\mathbf{b}\text{cov}(\mathbf{X}, \mathbf{X}) = \text{cov}(\mathbf{Y}, \mathbf{X})$ and therefore $\mathbf{b} = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})$. Then from the first equation, $\mathbf{a} + \mathbf{b}\mathbb{E}(\mathbf{X}) = \mathbf{Y}$ so $\mathbf{a} = \mathbb{E}(\mathbf{Y}) - \mathbf{b}\mathbb{E}(\mathbf{X})$.

A simple corollary is the $\mathbf{Y} - L(\mathbf{Y} | \mathbf{X})$ is uncorrelated with any affine function of \mathbf{X} :

If \mathbf{U} is an affine function of \mathbf{X} then

1. $\text{cov}[\mathbf{Y} - L(\mathbf{Y} | \mathbf{X}), \mathbf{U}] = \mathbf{0}$
2. $\mathbb{E}(\langle \mathbf{Y} - L(\mathbf{Y} | \mathbf{X}), \mathbf{U} \rangle) = 0$

Proof

Suppose that $\mathbf{U} = \mathbf{a} + \mathbf{b}\mathbf{X}$ where $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^{m \times n}$. For simplicity, let $\mathbf{L} = L(\mathbf{Y} | \mathbf{X})$

1. From the previous result, $\text{cov}(\mathbf{Y}, \mathbf{X}) = \text{cov}(\mathbf{L}, \mathbf{X})$. Hence using linearity,

$$\text{cov}(\mathbf{Y} - \mathbf{L}, \mathbf{U}) = \text{cov}(\mathbf{Y} - \mathbf{L}, \mathbf{a}) + \text{cov}(\mathbf{Y} - \mathbf{L}, \mathbf{X})\mathbf{b}^T = \mathbf{0} + [\text{cov}(\mathbf{Y}, \mathbf{X}) - \text{cov}(\mathbf{L}, \mathbf{X})] = \mathbf{0} \quad (4.8.11)$$

2. Recall that $\langle \mathbf{Y} - \mathbf{L}, \mathbf{U} \rangle$ is the trace of $\text{cov}(\mathbf{Y} - \mathbf{L}, \mathbf{U})$ and hence has expected value 0 by part (a).

The variance-covariance matrix of $L(\mathbf{Y} | \mathbf{X})$, and its covariance matrix with \mathbf{Y} turn out to be the same, again analogous to the single variable case.

Additional properties of $L(\mathbf{Y} | \mathbf{X})$:

1. $\text{cov}[\mathbf{Y}, L(\mathbf{Y} | \mathbf{X})] = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{X}, \mathbf{Y})$
2. $\text{vc}[L(\mathbf{Y} | \mathbf{X})] = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{X}, \mathbf{Y})$

Proof

Recall that $L(\mathbf{Y} | \mathbf{X}) = \mathbb{E}(\mathbf{Y}) + \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})[\mathbf{X} - \mathbb{E}(\mathbf{X})]$

1. Using basic properties of covariance,

$$\text{cov}[\mathbf{Y}, L(\mathbf{Y} | \mathbf{X})] = \text{cov}[\mathbf{Y}, \mathbf{X} - \mathbb{E}(\mathbf{X})] [\text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})]^T = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{X}, \mathbf{Y}) \quad (4.8.12)$$

2. Using basic properties of variance-covariance,

$$\text{vc}[L(\mathbf{Y} | \mathbf{X})] = \text{vc}[\text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\mathbf{X}] = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{vc}(\mathbf{X})[\text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})]^T = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{X}, \mathbf{Y}) \quad (4.8.13)$$

Next is the fundamental result that $L(\mathbf{Y} | \mathbf{X})$ is the affine function of \mathbf{X} that is closest to \mathbf{Y} in the mean square sense.

Suppose that $\mathbf{U} \in \mathbb{R}^n$ is an affine function of \mathbf{X} . Then

1. $\mathbb{E}(\|\mathbf{Y} - L(\mathbf{Y} | \mathbf{X})\|^2) \leq \mathbb{E}(\|\mathbf{Y} - \mathbf{U}\|^2)$
2. Equality holds in (a) if and only if $\mathbf{U} = L(\mathbf{Y} | \mathbf{X})$ with probability 1.

Proof

Again, let $\mathbf{L} = L(\mathbf{Y} | \mathbf{X})$ for simplicity and let $\mathbf{U} \in \mathbb{R}^n$ be an affine function of \mathbf{X} .

1. Using the linearity of expected value, note that

$$\mathbb{E}(\|\mathbf{Y} - \mathbf{U}\|^2) = \mathbb{E}(\|(\mathbf{Y} - \mathbf{L}) + (\mathbf{L} - \mathbf{U})\|^2) = \mathbb{E}(\|\mathbf{Y} - \mathbf{L}\|^2) + 2\mathbb{E}(\langle \mathbf{Y} - \mathbf{L}, \mathbf{L} - \mathbf{U} \rangle) + \mathbb{E}(\|\mathbf{L} - \mathbf{U}\|^2) \quad (4.8.14)$$

But $\mathbf{L} - \mathbf{U}$ is an affine function of \mathbf{X} and hence the middle term is 0 by our previous corollary. Hence

$$\mathbb{E}(\|\mathbf{Y} - \mathbf{U}\|^2) = \mathbb{E}(\|\mathbf{L} - \mathbf{Y}\|^2) + \mathbb{E}(\|\mathbf{L} - \mathbf{U}\|^2) \geq \mathbb{E}(\|\mathbf{L} - \mathbf{Y}\|^2)$$

2. From (a), equality holds in the inequality if and only if $\mathbb{E}(\|\mathbf{L} - \mathbf{U}\|^2) = 0$ if and only if $\mathbb{P}(\mathbf{L} = \mathbf{U}) = 1$.

The variance-covariance matrix of the difference between \mathbf{Y} and the best affine approximation is given in the next theorem.

$$\text{vc}[\mathbf{Y} - L(\mathbf{Y} | \mathbf{X})] = \text{vc}(\mathbf{Y}) - \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{X}, \mathbf{Y})$$

Proof

Again, we abbreviate $L(\mathbf{Y} | \mathbf{X})$ by \mathbf{L} . Using basic properties of variance-covariance matrices,

$$\text{vc}(\mathbf{Y} - \mathbf{L}) = \text{vc}(\mathbf{Y}) - \text{cov}(\mathbf{Y}, \mathbf{L}) - \text{cov}(\mathbf{L}, \mathbf{Y}) + \text{vc}(\mathbf{L}) \quad (4.8.15)$$

But $\text{cov}(\mathbf{Y}, \mathbf{L}) = \text{cov}(\mathbf{L}, \mathbf{Y}) = \text{vc}(\mathbf{L}) = \text{cov}(\mathbf{Y}, \mathbf{X})\text{vc}^{-1}(\mathbf{X})\text{cov}(\mathbf{Y}, \mathbf{X})$. Substituting gives the result.

The actual mean square error when we use $L(\mathbf{Y} | \mathbf{X})$ to approximate \mathbf{Y} , namely $\mathbb{E}(\|\mathbf{Y} - L(\mathbf{Y} | \mathbf{X})\|^2)$, is the trace (sum of the diagonal entries) of the variance-covariance matrix above. The function of \mathbf{x} given by

$$L(\mathbf{Y} | \mathbf{X} = \mathbf{x}) = \mathbb{E}(\mathbf{Y}) + \text{cov}(\mathbf{Y}, \mathbf{X}) \text{vc}^{-1}(\mathbf{X}) [\mathbf{x} - \mathbb{E}(\mathbf{X})] \quad (4.8.16)$$

is known as the (distribution) linear regression function. If we observe \mathbf{x} then $L(\mathbf{Y} | \mathbf{X} = \mathbf{x})$ is our best affine prediction of \mathbf{Y} .

Multiple linear regression is more powerful than it may at first appear, because it can be applied to non-linear transformations of the random vectors. That is, if $g: \mathbb{R}^m \rightarrow \mathbb{R}^j$ and $h: \mathbb{R}^n \rightarrow \mathbb{R}^k$ then $L[h(\mathbf{Y}) | g(\mathbf{X})]$ is the affine function of $g(\mathbf{X})$ that is closest to $h(\mathbf{Y})$ in the mean square sense. Of course, we must be able to compute the appropriate means, variances, and covariances.

Moreover, *Non-linear regression* with a single, real-valued predictor variable can be thought of as a special case of multiple linear regression. Thus, suppose that X is the predictor variable, Y is the response variable, and that (g_1, g_2, \dots, g_n) is a sequence of real-valued functions. We can apply the results of this section to find the linear function of $(g_1(X), g_2(X), \dots, g_n(X))$ that is closest to Y in the mean square sense. We just replace X_i with $g_i(X)$ for each i . Again, we must be able to compute the appropriate means, variances, and covariances to do this.

Examples and Applications

Suppose that (X, Y) has probability density function f defined by $f(x, y) = x + y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$. Find each of the following:

1. $\mathbb{E}(X, Y)$
2. $\text{vc}(X, Y)$

Answer

1. $\left(\frac{7}{12}, \frac{7}{12}\right)$
2. $\begin{bmatrix} \frac{11}{144} & -\frac{1}{144} \\ -\frac{1}{144} & \frac{11}{144} \end{bmatrix}$

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 2(x + y)$ for $0 \leq x \leq y \leq 1$. Find each of the following:

1. $\mathbb{E}(X, Y)$
2. $\text{vc}(X, Y)$

Answer

1. $\left(\frac{5}{12}, \frac{3}{4}\right)$
2. $\begin{bmatrix} \frac{43}{720} & \frac{1}{48} \\ \frac{1}{48} & \frac{3}{80} \end{bmatrix}$

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 6x^2y$ for $0 \leq x \leq 1, 0 \leq y \leq 1$. Find each of the following:

1. $\mathbb{E}(X, Y)$
2. $\text{vc}(X, Y)$

Answer

Note that X and Y are independent.

1. $\left(\frac{3}{4}, \frac{2}{3}\right)$
2. $\begin{bmatrix} \frac{3}{80} & 0 \\ 0 & \frac{1}{18} \end{bmatrix}$

Suppose that (X, Y) has probability density function f defined by $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$. Find each of the following:

1. $\mathbb{E}(X, Y)$
2. $\text{vc}(X, Y)$
3. $L(Y | X)$
4. $L[Y | (X, X^2)]$
5. Sketch the regression curves on the same set of axes.

Answer

1. $\left(\frac{5}{8}, \frac{5}{6}\right)$
2. $\begin{bmatrix} \frac{17}{448} & \frac{5}{336} \\ \frac{5}{336} & \frac{5}{252} \end{bmatrix}$
3. $\frac{10}{17} + \frac{20}{51}X$
4. $\frac{49}{76} + \frac{10}{57}X + \frac{7}{38}X^2$

Suppose that (X, Y, Z) is uniformly distributed on the region $\{(x, y, z) \in \mathbb{R}^3 : 0 \leq x \leq y \leq z \leq 1\}$. Find each of the following:

1. $\mathbb{E}(X, Y, Z)$
2. $\text{vc}(X, Y, Z)$

3. $L[Z | (X, Y)]$
4. $L[Y | (X, Z)]$
5. $L[X | (Y, Z)]$
6. $L[(Y, Z) | X]$

Answer

1. $(\frac{1}{4}, \frac{1}{2}, \frac{3}{4})$
2. $\begin{bmatrix} \frac{3}{80} & \frac{1}{40} & \frac{1}{80} \\ \frac{1}{40} & \frac{1}{20} & \frac{1}{40} \\ \frac{1}{80} & \frac{1}{40} & \frac{3}{80} \end{bmatrix}$
3. $\frac{1}{2} + \frac{1}{2}Y$. Note that there is no X term.
4. $\frac{1}{2}X + \frac{1}{2}Z$. Note that this is the midpoint of the interval $[X, Z]$.
5. $\frac{1}{2}Y$. Note that there is no Z term.
6. $\begin{bmatrix} \frac{1}{3} + \frac{2}{3}X \\ \frac{2}{3} + \frac{1}{3}X \end{bmatrix}$

Suppose that X is uniformly distributed on $(0, 1)$, and that given X , random variable Y is uniformly distributed on $(0, X)$. Find each of the following:

1. $\mathbb{E}(X, Y)$
2. $\text{vc}(X, Y)$

Answer

1. $(\frac{1}{2}, \frac{1}{4})$
2. $\begin{bmatrix} \frac{1}{12} & \frac{1}{24} \\ \frac{1}{24} & \frac{7}{144} \end{bmatrix}$

This page titled [4.8: Expected Value and Covariance Matrices](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.9: Expected Value as an Integral

In the introductory section, we defined expected value separately for discrete, continuous, and mixed distributions, using density functions. In the section on additional properties, we showed how these definitions can be unified, by first defining expected value for nonnegative random variables in terms of the right-tail distribution function. However, by far the best and most elegant definition of expected value is as an integral with respect to the underlying probability measure. This definition and a review of the properties of expected value are the goals of this section. No proofs are necessary (you will be happy to know), since all of the results follow from the general theory of integration. However, to understand the exposition, you will need to review the advanced sections on the integral with respect to a positive measure and the properties of the integral. If you are a new student of probability, or are not interested in the measure-theoretic detail of the subject, you can safely skip this section.

Definitions

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So Ω is the set of outcomes, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) .

Recall that a random variable X for the experiment is simply a measurable function from (Ω, \mathcal{F}) into another measurable space (S, \mathcal{S}) . When $S \subseteq \mathbb{R}^n$, we assume that S is Lebesgue measurable, and we take \mathcal{S} to be the σ -algebra of Lebesgue measurable subsets of S . As noted above, here is the measure-theoretic definition:

If X is a real-valued random variable on the probability space, the expected value of X is defined as the integral of X with respect to \mathbb{P} , assuming that the integral exists:

$$\mathbb{E}(X) = \int_{\Omega} X d\mathbb{P} \quad (4.9.1)$$

Let's review how the integral is defined in stages, but now using the notation of probability theory.

Let S denote the support set of X , so that S is a measurable subset of \mathbb{R} .

1. If S is finite, then $\mathbb{E}(X) = \sum_{x \in S} x \mathbb{P}(X = x)$.
2. If $S \subseteq [0, \infty)$, then $\mathbb{E}(X) = \sup \{ \mathbb{E}(Y) : Y \text{ has finite range and } 0 \leq Y \leq X \}$
3. For general $S \subseteq \mathbb{R}$, $\mathbb{E}(X) = \mathbb{E}(X^+) - \mathbb{E}(X^-)$ as long as the right side is not of the form $\infty - \infty$, and where X^+ and X^- denote the positive and negative parts of X .
4. If $A \in \mathcal{F}$, then $\mathbb{E}(X; A) = \mathbb{E}(X \mathbf{1}_A)$, assuming that the expected value on the right exists.

Thus, as with integrals generally, an expected value can exist as a number in \mathbb{R} (in which case X is *integrable*), can exist as ∞ or $-\infty$, or can fail to exist. In reference to part (a), a random variable with a finite set of values in \mathbb{R} is a *simple function* in the terminology of general integration. In reference to part (b), note that the expected value of a nonnegative random variable always exists in $[0, \infty]$. In reference to part (c), $\mathbb{E}(X)$ exists if and only if either $\mathbb{E}(X^+) < \infty$ or $\mathbb{E}(X^-) < \infty$.

Our next goal is to restate the basic theorems and properties of integrals, but in the notation of probability. Unless otherwise noted, all random variables are assumed to be real-valued.

Basic Properties

The Linear Properties

Perhaps the most important and basic properties are the *linear properties*. Part (a) is the *additive property* and part (b) is the *scaling property*.

Suppose that X and Y are random variables whose expected values exist, and that $c \in \mathbb{R}$. Then

1. $\mathbb{E}(X + Y) = \mathbb{E}(X) + \mathbb{E}(Y)$ as long as the right side is not of the form $\infty - \infty$.
2. $\mathbb{E}(cX) = c\mathbb{E}(X)$

Thus, part (a) holds if at least one of the expected values on the right is finite, or if both are ∞ , or if both are $-\infty$. What is ruled out are the two cases where one expected value is ∞ and the other is $-\infty$, and this is what is meant by the indeterminate form $\infty - \infty$.

Equality and Order

Our next set of properties deal with equality and order. First, the expected value of a random variable over a null set is 0.

If X is a random variable and A is an event with $\mathbb{P}(A) = 0$. Then $\mathbb{E}(X; A) = 0$.

Random variables that are equivalent have the same expected value

If X is a random variable whose expected value exists, and Y is a random variable with $\mathbb{P}(X = Y) = 1$, then $\mathbb{E}(X) = \mathbb{E}(Y)$.

Our next result is the *positive property* of expected value.

Suppose that X is a random variable and $\mathbb{P}(X \geq 0) = 1$. Then

1. $\mathbb{E}(X) \geq 0$
2. $\mathbb{E}(X) = 0$ if and only if $\mathbb{P}(X = 0) = 1$.

So, if X is a nonnegative random variable then $\mathbb{E}(X) > 0$ if and only if $\mathbb{P}(X > 0) > 0$. The next result is the *increasing property* of expected value, perhaps the most important property after linearity.

Suppose that X, Y are random variables whose expected values exist, and that $\mathbb{P}(X \leq Y) = 1$. Then

1. $\mathbb{E}(X) \leq \mathbb{E}(Y)$
2. Except in the case that both expected values are ∞ or both $-\infty$, $\mathbb{E}(X) = \mathbb{E}(Y)$ if and only if $\mathbb{P}(X = Y) = 1$.

So if $X \leq Y$ with probability 1 then, except in the two cases mentioned, $\mathbb{E}(X) < \mathbb{E}(Y)$ if and only if $\mathbb{P}(X < Y) > 0$. The next result is the *absolute value inequality*.

Suppose that X is a random variable whose expected value exists. Then

1. $|\mathbb{E}(X)| \leq \mathbb{E}(|X|)$
2. If $\mathbb{E}(X)$ is finite, then equality holds in (a) if and only if $\mathbb{P}(X \geq 0) = 1$ or $\mathbb{P}(X \leq 0) = 1$.

Change of Variables and Density Functions

The Change of Variables Theorem

Suppose now that X is a general random variable on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, taking values in a measurable space (S, \mathcal{S}) . Recall that the *probability distribution* of X is the probability measure P on (S, \mathcal{S}) given by $P(A) = \mathbb{P}(X \in A)$ for $A \in \mathcal{S}$. This is a special case of a new positive measure induced by a given positive measure and a measurable function. If $g: S \rightarrow \mathbb{R}$ is measurable, then $g(X)$ is a real-valued random variable. The following result shows how to compute the expected value of $g(X)$ as an integral with respect to the distribution of X , and is known as the *change of variables theorem*.

If $g: S \rightarrow \mathbb{R}$ is measurable then, assuming that the expected value exists,

$$\mathbb{E}[g(X)] = \int_S g(x) dP(x) \quad (4.9.2)$$

So, using the original definition and the change of variables theorem, and giving the variables explicitly for emphasis, we have

$$\mathbb{E}[g(X)] = \int_{\Omega} g[X(\omega)] d\mathbb{P}(\omega) = \int_S g(x) dP(x) \quad (4.9.3)$$

The Radon-Nikodym Theorem

Suppose now μ is a positive measure on (S, \mathcal{S}) , and that the distribution of X is absolutely continuous with respect to μ . Recall that this means that $\mu(A) = 0$ implies $P(A) = \mathbb{P}(X \in A) = 0$ for $A \in \mathcal{S}$. By the *Radon-Nikodym theorem*, named for Johann Radon and Otto Nikodym, X has a *probability density function* f with respect to μ . That is,

$$P(A) = \mathbb{P}(X \in A) = \int_A f d\mu, \quad A \in \mathcal{S} \quad (4.9.4)$$

In this case, we can write the expected value of $g(X)$ as an integral with respect to the probability density function.

If $g: S \rightarrow \mathbb{R}$ is measurable then, assuming that the expected value exists,

$$\mathbb{E}[g(X)] = \int_S g f d\mu \quad (4.9.5)$$

Again, giving the variables explicitly for emphasis, we have the following chain of integrals:

$$\mathbb{E}[g(X)] = \int_{\Omega} g[X(\omega)] d\mathbb{P}(\omega) = \int_S g(x) dP(x) = \int_S g(x) f(x) d\mu(x) \quad (4.9.6)$$

There are two critically important special cases.

Discrete Distributions

Suppose first that $(S, \mathcal{S}, \#)$ is a *discrete measure space*, so that S is countable, $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S , and $\#$ is counting measure on (S, \mathcal{S}) . Thus, X has a discrete distribution on S , and this distribution is always absolutely continuous with respect to $\#$. Specifically, $\#(A) = 0$ if and only if $A = \emptyset$ and of course $\mathbb{P}(X \in \emptyset) = 0$. The probability density function f of X with respect to $\#$, as we know, is simply $f(x) = \mathbb{P}(X = x)$ for $x \in S$. Moreover, integrals with respect to $\#$ are sums, so

$$\mathbb{E}[g(X)] = \sum_{x \in S} g(x) f(x) \quad (4.9.7)$$

assuming that the expected value exists. Existence in this case means that either the sum of the positive terms is finite or the sum of the negative terms is finite, so that the sum makes sense (and in particular does not depend on the order in which the terms are added). Specializing further, if X itself is real-valued and $g = 1$ we have

$$\mathbb{E}(X) = \sum_{x \in S} x f(x) \quad (4.9.8)$$

which was our original definition of expected value in the discrete case.

Continuous Distributions

For the second special case, suppose that $(S, \mathcal{S}, \lambda_n)$ is a *Euclidean measure space*, so that S is a Lebesgue measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$, \mathcal{S} is the σ -algebra of Lebesgue measurable subsets of S , and λ_n is Lebesgue measure on (S, \mathcal{S}) . The distribution of X is absolutely continuous with respect to λ_n if $\lambda_n(A) = 0$ implies $\mathbb{P}(X \in A) = 0$ for $A \in \mathcal{S}$. If this is the case, then a probability density function f of X has its usual meaning. Thus,

$$\mathbb{E}[g(X)] = \int_S g(x)f(x) d\lambda_n(x) \quad (4.9.9)$$

assuming that the expected value exists. When g is a typically nice function, this integral reduces to an ordinary n -dimensional Riemann integral of calculus. Specializing further, if X is itself real-valued and $g = 1$ then

$$\mathbb{E}(X) = \int_S x f(x) dx \quad (4.9.10)$$

which was our original definition of expected value in the continuous case.

Interchange Properties

In this subsection, we review properties that allow the interchange of expected value and other operations: limits of sequences, infinite sums, and integrals. We assume again that the random variables are real-valued unless otherwise specified.

Limits

Our first set of convergence results deals with the interchange of expected value and limits. We start with the expected value version of *Fatou's lemma*, named in honor of Pierre Fatou. Its usefulness stems from the fact that no assumptions are placed on the random variables, except that they be nonnegative.

Suppose that X_n is a nonnegative random variable for $n \in \mathbb{N}_+$. Then

$$\mathbb{E}\left(\liminf_{n \rightarrow \infty} X_n\right) \leq \liminf_{n \rightarrow \infty} \mathbb{E}(X_n) \quad (4.9.11)$$

Our next set of results gives conditions for the interchange of expected value and limits.

Suppose that X_n is a random variable for each $n \in \mathbb{N}_+$. then

$$\mathbb{E}\left(\lim_{n \rightarrow \infty} X_n\right) = \lim_{n \rightarrow \infty} \mathbb{E}(X_n) \quad (4.9.12)$$

in each of the following cases:

1. X_n is nonnegative for each $n \in \mathbb{N}_+$ and X_n is increasing in n .
2. $\mathbb{E}(X_n)$ exists for each $n \in \mathbb{N}_+$, $\mathbb{E}(X_1) > -\infty$, and X_n is increasing in n .
3. $\mathbb{E}(X_n)$ exists for each $n \in \mathbb{N}_+$, $\mathbb{E}(X_1) < \infty$, and X_n is decreasing in n .
4. $\lim_{n \rightarrow \infty} X_n$ exists, and $|X_n| \leq Y$ for $n \in \mathbb{N}$ where Y is a nonnegative random variable with $\mathbb{E}(Y) < \infty$.
5. $\lim_{n \rightarrow \infty} X_n$ exists, and $|X_n| \leq c$ for $n \in \mathbb{N}$ where c is a positive constant.

Statements about the random variables in the theorem above (nonnegative, increasing, existence of limit, etc.) need only hold with probability 1. Part (a) is the *monotone convergence theorem*, one of the most important convergence results and in a sense, essential to the definition of the integral in the first place. Parts (b) and (c) are slight generalizations of the monotone convergence theorem. In parts (a), (b), and (c), note that $\lim_{n \rightarrow \infty} X_n$ exists (with probability 1), although the limit may be ∞ in parts (a) and (b) and $-\infty$ in part (c) (with positive probability). Part (d) is the *dominated convergence theorem*, another of the most important convergence results. It's sometimes also known as *Lebesgue's dominated convergence theorem* in honor of Henri Lebesgue. Part (e) is a corollary of the dominated convergence theorem, and is known as the *bounded convergence theorem*.

Infinite Series

Our next results involve the interchange of expected value and an infinite sum, so these results generalize the basic additivity property of expected value.

Suppose that X_n is a random variable for $n \in \mathbb{N}_+$. Then

$$\mathbb{E}\left(\sum_{n=1}^{\infty} X_n\right) = \sum_{n=1}^{\infty} \mathbb{E}(X_n) \quad (4.9.13)$$

in each of the following cases:

1. X_n is nonnegative for each $n \in \mathbb{N}_+$.
2. $\mathbb{E}(\sum_{n=1}^{\infty} |X_n|) < \infty$

Part (a) is a consequence of the monotone convergence theorem, and part (b) is a consequence of the dominated convergence theorem. In (b), note that $\sum_{n=1}^{\infty} |X_n| < \infty$ and hence $\sum_{n=1}^{\infty} X_n$ is absolutely convergent with probability 1. Our next result is the additivity of the expected value over a countably infinite collection of disjoint events.

Suppose that X is a random variable whose expected value exists, and that $\{A_n : n \in \mathbb{N}_+\}$ is a disjoint collection events. Let $A = \bigcup_{n=1}^{\infty} A_n$. Then

$$\mathbb{E}(X; A) = \sum_{n=1}^{\infty} \mathbb{E}(X; A_n) \quad (4.9.14)$$

Of course, the previous theorem applies in particular if X is nonnegative.

Integrals

Suppose that (T, \mathcal{T}, μ) is a σ -finite measure space, and that X_t is a real-valued random variable for each $t \in T$. Thus we can think of $\{X_t : t \in T\}$ is a stochastic process indexed by T . We assume that $(\omega, t) \mapsto X_t(\omega)$ is measurable, as a function from the product space $(\Omega \times T, \mathcal{F} \otimes \mathcal{T})$ into \mathbb{R} . Our next result involves the interchange of expected value and integral, and is a consequence of Fubini's theorem, named for Guido Fubini.

Under the assumptions above,

$$\mathbb{E} \left[\int_T X_t d\mu(t) \right] = \int_T \mathbb{E}(X_t) d\mu(t) \quad (4.9.15)$$

in each of the following cases:

1. X_t is nonnegative for each $t \in T$.
2. $\int_T \mathbb{E}(|X_t|) d\mu(t) < \infty$

Fubini's theorem actually states that the two iterated integrals above equal the joint integral

$$\int_{\Omega \times T} X_t(\omega) d(\mathbb{P} \otimes \mu)(\omega, t) \quad (4.9.16)$$

where of course, $\mathbb{P} \otimes \mu$ is the product measure on $(\Omega \times T, \mathcal{F} \otimes \mathcal{T})$. However, our interest is usually in evaluating the iterated integral above on the left in terms of the iterated integral on the right. Part (a) is the expected value version of *Tonelli's theorem*, named for Leonida Tonelli.

Examples and Exercises

You may have worked some of the computational exercises before, but try to see them in a new light, in terms of the general theory of integration.

The Cauchy Distribution

Recall that the *Cauchy distribution*, named for Augustin Cauchy, is a continuous distribution with probability density function f given by

$$f(x) = \frac{1}{\pi(1+x^2)}, \quad x \in \mathbb{R} \quad (4.9.17)$$

The Cauchy distribution is studied in more generality in the chapter on Special Distributions.

Suppose that X has the Cauchy distribution.

1. Show that $\mathbb{E}(X)$ does not exist.
2. Find $\mathbb{E}(X^2)$

Answer

1. $\mathbb{E}(X^+) = \mathbb{E}(X^-) = \infty$
2. ∞

Open the Cauchy Experiment and keep the default parameters. Run the experiment 1000 times and note the behavior of the sample mean.

The Pareto Distribution

Recall that the *Pareto distribution*, named for Vilfredo Pareto, is a continuous distribution with probability density function f given by

$$f(x) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (4.9.18)$$

where $a > 0$ is the *shape parameter*. The Pareto distribution is studied in more generality in the chapter on Special Distributions.

Suppose that X has the Pareto distribution with shape parameter a . Find $\mathbb{E}(X)$ in the following cases:

1. $0 < a \leq 1$
2. $a > 1$

Answer

1. ∞
2. $\frac{a}{a-1}$

Open the special distribution simulator and select the Pareto distribution. Vary the shape parameter and note the shape of the probability density function and the location of the mean. For various values of the parameter, run the experiment 1000 times and compare the sample mean with the distribution mean.

Suppose that X has the Pareto distribution with shape parameter a . Find $E(1/X^n)$ for $n \in \mathbb{N}_+$.

Answer

$$\frac{a}{a+n}$$

Special Results for Nonnegative Variables

For a nonnegative variable, the moments can be obtained from integrals of the right-tail distribution function.

If X is a nonnegative random variable then

$$\mathbb{E}(X^n) = \int_0^\infty nx^{n-1}\mathbb{P}(X > x) dx \quad (4.9.19)$$

Proof

By [Fubini's theorem](#) we can interchange an expected value and integral when the integrand is nonnegative. Hence

$$\int_0^\infty nx^{n-1}\mathbb{P}(X > x) dx = \int_0^\infty nx^{n-1}\mathbb{E}[\mathbf{1}(X > x)] dx = \mathbb{E}\left(\int_0^\infty nx^{n-1}\mathbf{1}(X > x) dx\right) = \mathbb{E}\left(\int_0^X nx^{n-1} dx\right) = \mathbb{E}(X^n) \quad (4.9.20)$$

When $n = 1$ we have $\mathbb{E}(X) = \int_0^\infty \mathbb{P}(X > x) dx$. We saw this result before in the section on additional properties of expected value, but now we can understand the proof in terms of Fubini's theorem.

For a random variable taking nonnegative integer values, the moments can be computed from sums involving the right-tail distribution function.

Suppose that X has a discrete distribution, taking values in \mathbb{N} . Then

$$\mathbb{E}(X^n) = \sum_{k=1}^\infty [k^n - (k-1)^n] \mathbb{P}(X \geq k) \quad (4.9.21)$$

Proof

By the [theorem above](#), we can interchange expected value and infinite series when the terms are nonnegative. Hence

$$\begin{aligned} \sum_{k=1}^\infty [k^n - (k-1)^n] \mathbb{P}(X \geq k) &= \sum_{k=1}^\infty [k^n - (k-1)^n] \mathbb{E}[\mathbf{1}(X \geq k)] = \mathbb{E}\left(\sum_{k=1}^\infty [k^n - (k-1)^n] \mathbf{1}(X \geq k)\right) \\ &= \mathbb{E}\left(\sum_{k=1}^X [k^n - (k-1)^n]\right) = \mathbb{E}(X^n) \end{aligned} \quad (4.9.22)$$

When $n = 1$ we have $\mathbb{E}(X) = \sum_{k=1}^\infty \mathbb{P}(X \geq k)$. We saw this result before in the section on additional properties of expected value, but now we can understand the proof in terms of the interchange of sum and expected value.

This page titled [4.9: Expected Value as an Integral](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.10: Conditional Expected Value Revisited

Conditional expected value is much more important than one might at first think. In fact, conditional expected value is at the core of modern probability theory because it provides the basic way of incorporating known information into a probability measure.

Basic Theory

Definition

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so that Ω is the set of outcomes, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . In our first elementary discussion, we studied the conditional expected value of a real-value random variable X given a general random variable Y . The more general approach is to condition on a sub σ -algebra \mathcal{G} of \mathcal{F} . The sections on σ -algebras and measure theory are essential prerequisites for this section.

Before we get to the definition, we need some preliminaries. First, all random variables mentioned are assumed to be real valued. Next the notion of equivalence plays a fundamental role in this section. Next recall that random variables X_1 and X_2 are *equivalent* if $\mathbb{P}(X_1 = X_2) = 1$. Equivalence really does define an equivalence relation on the collection of random variables defined on the sample space. Moreover, we often regard equivalent random variables as being essentially the same object. More precisely from this point of view, the objects of our study are not individual random variables but rather equivalence classes of random variables under this equivalence relation. Finally, for $A \in \mathcal{F}$, recall the notation for the expected value of X on the event A

$$\mathbb{E}(X; A) = \mathbb{E}(X \mathbf{1}_A) \quad (4.10.1)$$

assuming of course that the expected value exists. For the remainder of this subsection, suppose that \mathcal{G} is a sub σ -algebra of \mathcal{F} .

Suppose that X is a random variable with $\mathbb{E}(|X|) < \infty$. The *conditional expected value* of X given \mathcal{G} is the random variable $\mathbb{E}(X | \mathcal{G})$ defined by the following properties:

1. $\mathbb{E}(X | \mathcal{G})$ is measurable with respect to \mathcal{G} .
2. If $A \in \mathcal{G}$ then $\mathbb{E}[\mathbb{E}(X | \mathcal{G}); A] = \mathbb{E}(X; A)$

The basic idea is that $\mathbb{E}(X | \mathcal{G})$ is the expected value of X given the information in the σ -algebra \mathcal{G} . Hopefully this idea will become clearer during our study. The conditions above uniquely define $\mathbb{E}(X | \mathcal{G})$ up to equivalence. The proof of this fact is a simple application of the Radon-Nikodym theorem, named for Johann Radon and Otto Nikodym

Suppose again that X is a random variable with $\mathbb{E}(|X|) < \infty$.

1. There exists a random variable V satisfying the [definition](#).
2. If V_1 and V_2 satisfy the [definition](#), then $\mathbb{P}(V_1 = V_2) = 1$ so that V_1 and V_2 are equivalent.

Proof

1. Note that $\nu(A) = \mathbb{E}(X; A)$ for $A \in \mathcal{G}$ defines a (signed) measure on \mathcal{G} . Moreover, if $A \in \mathcal{G}$ and $\mathbb{P}(A) = 0$ then $\nu(A) = 0$. Hence ν is absolutely continuous with respect to the restriction of \mathbb{P} to \mathcal{G} . By the Radon-Nikodym theorem, there exists a random variable V that is measurable with respect to \mathcal{G} such that $\nu(A) = \mathbb{E}(V; A)$ for $A \in \mathcal{G}$. That is, V is the *density* or *derivative* of ν with respect to \mathbb{P} on \mathcal{G} .
2. This follows from the uniqueness of the Radon-Nikodym derivative, up to equivalence.

The following characterization might seem stronger but in fact is equivalent to the definition.

Suppose again that X is a random variable with $\mathbb{E}(|X|) < \infty$. Then $\mathbb{E}(X | \mathcal{G})$ is characterized by the following properties:

1. $\mathbb{E}(X | \mathcal{G})$ is measurable with respect to \mathcal{G}
2. If U is measurable with respect to \mathcal{G} and $\mathbb{E}(|UX|) < \infty$ then $\mathbb{E}[U\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(UX)$.

Proof

We have to show that part (b) in the [definition](#) is equivalent to part (b) here. First (b) here implies (b) in the definition since $\mathbf{1}_A$ is \mathcal{G} -measurable if $A \in \mathcal{G}$. Conversely suppose that (b) in the definition holds. We will show that (b) here holds by a classical *bootstrapping argument*. First $\mathbb{E}[U\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(UX)$ if $U = \mathbf{1}_A$ for some $A \in \mathcal{G}$. Next suppose that U is a simple random variable that is \mathcal{G} -measurable. That is, $U = \sum_{i \in I} a_i \mathbf{1}_{A_i}$ where I is a finite index set, $a_i \geq 0$ for $i \in I$, and $A_i \in \mathcal{G}$ for $i \in I$. then

$$\mathbb{E}[U\mathbb{E}(X | \mathcal{G})] = \mathbb{E}\left[\sum_{i \in I} a_i \mathbf{1}_{A_i} \mathbb{E}(X | \mathcal{G})\right] = \sum_{i \in I} a_i \mathbb{E}[\mathbf{1}_{A_i} \mathbb{E}(X | \mathcal{G})] = \sum_{i \in I} a_i \mathbb{E}(\mathbf{1}_{A_i} X) = \mathbb{E}\left(\sum_{i \in I} a_i \mathbf{1}_{A_i} X\right) = \mathbb{E}(UX) \quad (4.10.2)$$

Next suppose that U is nonnegative and \mathcal{G} -measurable. Then there exists a sequence of simple \mathcal{G} -measurable random variables (U_1, U_2, \dots) with $U_n \uparrow U$ as $n \rightarrow \infty$. Then by the previous step, $\mathbb{E}[U_n \mathbb{E}(X | \mathcal{G})] = \mathbb{E}(U_n X)$ for each n . Letting $n \rightarrow \infty$ and using the monotone convergence theorem we have $\mathbb{E}[U\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(UX)$. Finally, suppose that U is a general \mathcal{G} -measurable random variable. Then $U = U^+ - U^-$ where U^+ and U^- are the usual positive and negative parts of U . These parts are nonnegative and \mathcal{G} -measurable, so by the previous step, $\mathbb{E}[U^+ \mathbb{E}(X | \mathcal{G})] = \mathbb{E}(U^+ X)$ and $\mathbb{E}[U^- \mathbb{E}(X | \mathcal{G})] = \mathbb{E}(U^- X)$. hence

$$\mathbb{E}[U\mathbb{E}(X | \mathcal{G})] = \mathbb{E}[(U^+ - U^-)\mathbb{E}(X | \mathcal{G})] = \mathbb{E}[U^+ \mathbb{E}(X | \mathcal{G})] - \mathbb{E}[U^- \mathbb{E}(X | \mathcal{G})] = \mathbb{E}(U^+ X) - \mathbb{E}(U^- X) = \mathbb{E}(UX) \quad (4.10.3)$$

Properties

Our next discussion concerns some fundamental properties of conditional expected value. All equalities and inequalities are understood to hold *modulo equivalence*, that is, with probability 1. Note also that many of the proofs work by showing that the right hand side satisfies the properties in the [definition](#) for the conditional expected value on the left side. Once again we assume that \mathcal{G} is a sub *sigma*-algebra of \mathcal{F} .

Our first property is a simple consequence of the definition: X and $\mathbb{E}(X | \mathcal{G})$ have the same mean.

Suppose that X is a random variable with $\mathbb{E}(|X|) < \infty$. Then $\mathbb{E}[\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(X)$.

Proof

This follows immediately by letting $A = \Omega$ in the [definition](#).

The result above can often be used to compute $\mathbb{E}(X)$, by choosing the σ -algebra \mathcal{G} in a clever way. We say that we are computing $\mathbb{E}(X)$ by *conditioning on \mathcal{G}* . Our next properties are fundamental: every version of expected value must satisfy the *linearity* properties. The first part is the *additive property* and the second part is the *scaling property*.

Suppose that X and Y are random variables with $\mathbb{E}(|X|) < \infty$ and $\mathbb{E}(|Y|) < \infty$, and that $c \in \mathbb{R}$. Then

1. $\mathbb{E}(X + Y | \mathcal{G}) = \mathbb{E}(X | \mathcal{G}) + \mathbb{E}(Y | \mathcal{G})$
2. $\mathbb{E}(cX | \mathcal{G}) = c\mathbb{E}(X | \mathcal{G})$

Proof

1. Note that $\mathbb{E}(|X + Y|) \leq \mathbb{E}(|X|) + \mathbb{E}(|Y|) < \infty$ so $\mathbb{E}(X + Y | \mathcal{G})$ is defined. We show that $\mathbb{E}(X | \mathcal{G}) + \mathbb{E}(Y | \mathcal{G})$ satisfies the conditions in the [definition](#) for $\mathbb{E}(X + Y | \mathcal{G})$. Note first that $\mathbb{E}(X | \mathcal{G}) + \mathbb{E}(Y | \mathcal{G})$ is \mathcal{G} -measurable since both terms are. If $A \in \mathcal{G}$ then

$$\mathbb{E}\{\mathbb{E}(X | \mathcal{G}) + \mathbb{E}(Y | \mathcal{G}); A\} = \mathbb{E}[\mathbb{E}(X | \mathcal{G}); A] + \mathbb{E}[\mathbb{E}(Y | \mathcal{G}); A] = \mathbb{E}(X; A) + \mathbb{E}(Y; A) = \mathbb{E}(X + Y; A) \quad (4.10.4)$$

2. Note that $\mathbb{E}(|cX|) = |c|\mathbb{E}(|X|) < \infty$ so $\mathbb{E}(cX | \mathcal{G})$ is defined. We show that $c\mathbb{E}(X | \mathcal{G})$ satisfy the conditions in the [definition](#) for $\mathbb{E}(cX | \mathcal{G})$. Note first that $c\mathbb{E}(X | \mathcal{G})$ is \mathcal{G} -measurable since the second factor is. If $A \in \mathcal{G}$ then

$$\mathbb{E}[c\mathbb{E}(X | \mathcal{G}); A] = c\mathbb{E}[\mathbb{E}(X | \mathcal{G}); A] = c\mathbb{E}(X; A) = \mathbb{E}(cX; A) \quad (4.10.5)$$

The next set of properties are also fundamental to every notion of expected value. The first part is the *positive property* and the second part is the *increasing property*.

Suppose again that X and Y are random variables with $\mathbb{E}(|X|) < \infty$ and $\mathbb{E}(|Y|) < \infty$.

1. If $X \geq 0$ then $\mathbb{E}(X | \mathcal{G}) \geq 0$
2. If $X \leq Y$ then $\mathbb{E}(X | \mathcal{G}) \leq \mathbb{E}(Y | \mathcal{G})$

Proof

1. Let $A = \{\mathbb{E}(X | \mathcal{G}) < 0\}$. Note that $A \in \mathcal{G}$ and hence $\mathbb{E}(X; A) = \mathbb{E}[\mathbb{E}(X | \mathcal{G}); A]$. Since $X \geq 0$ with probability 1 we have $\mathbb{E}(X; A) \geq 0$. On the other hand, if $\mathbb{P}(A) > 0$ then $\mathbb{E}[\mathbb{E}(X | \mathcal{G}); A] < 0$ which is a contradiction. Hence we must have $\mathbb{P}(A) = 0$.
2. Note that if $X \leq Y$ then $Y - X \geq 0$. Hence by (a) and the additive property, $\mathbb{E}(Y - X | \mathcal{G}) = \mathbb{E}(Y | \mathcal{G}) - \mathbb{E}(X | \mathcal{G}) \geq 0$ so $\mathbb{E}(Y | \mathcal{G}) \geq \mathbb{E}(X | \mathcal{G})$.

The next few properties relate to the central idea that $\mathbb{E}(X | \mathcal{G})$ is the expected value of X given the information in the σ -algebra \mathcal{G} .

Suppose that X and V are random variables with $\mathbb{E}(|X|) < \infty$ and $\mathbb{E}(|XV|) < \infty$ and that V is measurable with respect to \mathcal{G} . Then $\mathbb{E}(VX | \mathcal{G}) = V\mathbb{E}(X | \mathcal{G})$.

Proof

We show that $V\mathbb{E}(X | \mathcal{G})$ satisfy the in [properties that characterize \$\mathbb{E}\(VX | \mathcal{G}\)\$](#) . First, $V\mathbb{E}(X | \mathcal{G})$ is \mathcal{G} -measurable since both factors are. If U is \mathcal{G} -measurable with $\mathbb{E}(|UVX|) < \infty$ then UV is also \mathcal{G} -measurable and hence

$$\mathbb{E}[UV\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(UVX) = \mathbb{E}[U(VX)] \quad (4.10.6)$$

Compare this result with the [scaling property](#). If V is measurable with respect to \mathcal{G} then V is like a constant in terms of the conditional expected value given \mathcal{G} . On the other hand, note that this result implies the [scaling property](#), since a constant can be viewed as a random variable, and as such, is measurable with respect to any σ -algebra. As a corollary to this result, note that if X itself is measurable with respect to \mathcal{G} then $\mathbb{E}(X | \mathcal{G}) = X$. The following result gives the other extreme.

Suppose that X is a random variable with $\mathbb{E}(|X|) < \infty$. If X and \mathcal{G} are independent then $\mathbb{E}(X | \mathcal{G}) = \mathbb{E}(X)$.

Proof

We show that $\mathbb{E}(X)$ satisfy the properties in the [definition](#) for $\mathbb{E}(X | \mathcal{G})$. First of course, $\mathbb{E}(X)$ is \mathcal{G} -measurable as a constant random variable. If $A \in \mathcal{G}$ then X and 1_A are independent and hence

$$\mathbb{E}(X; A) = \mathbb{E}(X)\mathbb{P}(A) = \mathbb{E}[\mathbb{E}(X); A] \quad (4.10.7)$$

Every random variable X is independent of the trivial σ -algebra $\{\emptyset, \Omega\}$ so it follows that $\mathbb{E}(X | \{\emptyset, \Omega\}) = \mathbb{E}(X)$.

The next properties are *consistency conditions*, also known as the *tower properties*. When conditioning twice, with respect to nested σ -algebras, the smaller one (representing the least amount of information) always prevails.

Suppose that X is a random variable with $\mathbb{E}(|X|) < \infty$ and that \mathcal{H} is a sub σ -algebra of \mathcal{G} . Then

1. $\mathbb{E}[\mathbb{E}(X | \mathcal{H}) | \mathcal{G}] = \mathbb{E}(X | \mathcal{H})$
2. $\mathbb{E}[\mathbb{E}(X | \mathcal{G}) | \mathcal{H}] = \mathbb{E}(X | \mathcal{H})$

Proof

1. Note first that $\mathbb{E}(X | \mathcal{H})$ is \mathcal{H} -measurable and hence also \mathcal{G} -measurable. Thus by (7), $\mathbb{E}[\mathbb{E}(X | \mathcal{H}) | \mathcal{G}] = \mathbb{E}(X | \mathcal{H})$.
2. We show that $\mathbb{E}(X | \mathcal{H})$ satisfies the conditions in the [definition](#) for $\mathbb{E}[\mathbb{E}(X | \mathcal{G}) | \mathcal{H}]$. Note again that $\mathbb{E}(X | \mathcal{H})$ is \mathcal{H} -measurable. If $A \in \mathcal{H}$ then $A \in \mathcal{G}$ and hence

$$\mathbb{E}[\mathbb{E}(X | \mathcal{G}); A] = \mathbb{E}(X; A) = \mathbb{E}[\mathbb{E}(X | \mathcal{H}); A] \quad (4.10.8)$$

The next result gives *Jensen's inequality* for conditional expected value, named for Johan Jensen.

Suppose that X takes values in an interval $S \subseteq \mathbb{R}$ and that $g: S \rightarrow \mathbb{R}$ is convex. If $\mathbb{E}(|X|) < \infty$ and $\mathbb{E}(|g(X)|) < \infty$ then


$$\mathbb{E}[g(X) | \mathcal{G}] \geq g[\mathbb{E}(X | \mathcal{G})] \quad (4.10.9)$$

Proof

As with Jensen's inequality for ordinary expected value, the best proof uses the characterization of convex functions in terms of *supporting lines*: For each $t \in S$ there exist numbers a and b (depending on t) such that

- $a + bt = g(t)$
- $a + bx \leq g(x)$ for $x \in S$

A convex function and several supporting lines



Random variables X and $\mathbb{E}(X | \mathcal{G})$ takes values in S . We can construct a *random* supporting line at $\mathbb{E}(X | \mathcal{G})$. That is, there exist random variables A and B , measurable with respect to \mathcal{G} , such that

1. $A + B\mathbb{E}(X | \mathcal{G}) = g[\mathbb{E}(X | \mathcal{G})]$
2. $A + BX \leq g(X)$

We take conditional expected value through the inequality in (b) and then use properties of conditional expected value and property (a):

$$\mathbb{E}[g(X) | \mathcal{G}] \geq \mathbb{E}(A + BX | \mathcal{G}) = A + B\mathbb{E}(X | \mathcal{G}) = g[\mathbb{E}(X | \mathcal{G})] \quad (4.10.10)$$

Note that the second step uses the fact that A and B are measurable with respect to \mathcal{G} .

Conditional Probability

For our next discussion, suppose as usual that \mathcal{G} is a sub σ -algebra of \mathcal{F} . The conditional probability of an event A given \mathcal{G} can be defined as a special case of conditional expected value. As usual, let $\mathbf{1}_A$ denote the indicator random variable of A .

For $A \in \mathcal{F}$ we define

$$\mathbb{P}(A | \mathcal{G}) = \mathbb{E}(\mathbf{1}_A | \mathcal{G}) \quad (4.10.11)$$

Thus, we have the following characterizations of conditional probability, which are special cases of the [definition](#) and the [alternate version](#):

If $A \in \mathcal{F}$ then $\mathbb{P}(A | \mathcal{G})$ is characterized (up to equivalence) by the following properties

1. $\mathbb{P}(A | \mathcal{G})$ is measurable with respect to \mathcal{G} .
2. If $B \in \mathcal{G}$ then $\mathbb{E}[\mathbb{P}(A | \mathcal{G}); B] = \mathbb{P}(A \cap B)$

Proof

For part (b), note that

$$\mathbb{E}[\mathbf{1}_B \mathbb{P}(A | \mathcal{G})] = \mathbb{E}[\mathbf{1}_B \mathbb{E}(\mathbf{1}_A | \mathcal{G})] = \mathbb{E}(\mathbf{1}_A \mathbf{1}_B) = \mathbb{E}(\mathbf{1}_{A \cap B}) = \mathbb{P}(A \cap B) \quad (4.10.12)$$

If $A \in \mathcal{F}$ then $\mathbb{P}(A | \mathcal{G})$ is characterized (up to equivalence) by the following properties

1. $\mathbb{P}(A | \mathcal{G})$ is measurable with respect to \mathcal{G} .
2. If U is measurable with respect to \mathcal{G} and $\mathbb{E}(|U|) < \infty$ then $\mathbb{E}[U \mathbb{P}(A | \mathcal{G})] = \mathbb{E}(U; A)$

The properties above for conditional expected value, of course, have special cases for conditional probability. In particular, we can compute the probability of an event by *conditioning* on a σ -algebra:

If $A \in \mathcal{F}$ then $\mathbb{P}(A) = \mathbb{E}[\mathbb{P}(A | \mathcal{G})]$.

Proof

This is a direct result of the [mean property](#) since $\mathbb{E}(\mathbf{1}_A) = \mathbb{P}(A)$.

Again, the last theorem is often a good way to compute $\mathbb{P}(A)$ when we know the conditional probability of A given \mathcal{G} . This is a very compact and elegant version of the law of total probability given first in the section on Conditional Probability in the chapter on Probability Spaces and later in the section on Discrete Distributions in the Chapter on Distributions. The following theorem gives the conditional version of the axioms of probability.

The following properties hold (as usual, modulo equivalence):

1. $\mathbb{P}(A | \mathcal{G}) \geq 0$ for every $A \in \mathcal{F}$
2. $\mathbb{P}(\Omega | \mathcal{G}) = 1$
3. If $\{A_i : i \in I\}$ is a countable disjoint subset of \mathcal{F} then $\mathbb{P}(\bigcup_{i \in I} A_i | \mathcal{G}) = \sum_{i \in I} \mathbb{P}(A_i | \mathcal{G})$

Proof

1. This is a direct consequence of (6).
2. This is trivial since $\mathbf{1}_\Omega = 1$.
3. We show that the right side satisfies the conditions in (11) that define the left side. Note that $\sum_{i \in I} \mathbb{P}(A_i | \mathcal{G})$ is \mathcal{G} -measurable since each term in the sum has this property. Let $B \in \mathcal{G}$. then

$$\mathbb{E}\left[\sum_{i \in I} \mathbb{P}(A_i | \mathcal{G}); B\right] = \sum_{i \in I} \mathbb{E}[\mathbb{P}(A_i | \mathcal{G}); B] = \sum_{i \in I} \mathbb{P}(A_i \cap B) = \mathbb{P}\left(B \cap \bigcup_{i \in I} A_i\right) \quad (4.10.13)$$

From the last result, it follows that other standard probability rules hold for conditional probability given \mathcal{G} (as always, modulo equivalence). These results include

- the complement rule

- the increasing property
- Boole's inequality
- Bonferroni's inequality
- the inclusion-exclusion laws

However, it is not correct to state that $A \mapsto \mathbb{P}(A | \mathcal{G})$ is a probability measure, because the conditional probabilities are only defined up to equivalence, and so the mapping does not make sense. We would have to specify a particular version of $\mathbb{E}(A | \mathcal{G})$ for each $A \in \mathcal{F}$ for the mapping to make sense. Even if we do this, the mapping may not define a probability measure. In part (c), the left and right sides are random variables and the equation is an event that has probability 1. However this event depends on the collection $\{A_i : i \in I\}$. In general, there will be uncountably many such collections in \mathcal{F} , and the intersection of all of the corresponding events may well have probability less than 1 (if it's measurable at all). It turns out that if the underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is sufficiently "nice" (and most probability spaces that arise in applications are nice), then there does in fact exist a *regular conditional probability*. That is, for each $A \in \mathcal{F}$, there exists a random variable $\mathbb{P}(A | \mathcal{G})$ satisfying the conditions in (12) and such that with probability 1, $A \mapsto \mathbb{P}(A | \mathcal{G})$ is a probability measure.

The following theorem gives a version of *Bayes' theorem*, named for the inimitable Thomas Bayes.

Suppose that $A \in \mathcal{G}$ and $B \in \mathcal{F}$. then

$$\mathbb{P}(A | B) = \frac{\mathbb{E}[\mathbb{P}(B | \mathcal{G}); A]}{\mathbb{E}[\mathbb{P}(B | \mathcal{G})]} \quad (4.10.14)$$

Proof

The proof is absolutely trivial. By definition of conditional probability given \mathcal{G} , the numerator is $\mathbb{P}(A \cap B)$ and the denominator is $\mathbb{P}(B)$. Nonetheless, Bayes' theorem is useful in settings where the expected values in the numerator and denominator can be computed directly

Basic Examples

The purpose of this discussion is to tie the general notions of conditional expected value that we are studying here to the more elementary concepts that you have seen before. Suppose that A is an event (that is, a member of \mathcal{F}) with $\mathbb{P}(A) > 0$. If B is another event, then of course, the conditional probability of B given A is

$$\mathbb{P}(B | A) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(A)} \quad (4.10.15)$$

If X is a random variable then the conditional distribution of X given A is the probability measure on \mathbb{R} given by

$$R \mapsto \mathbb{P}(X \in R | A) = \frac{\mathbb{P}(\{X \in R\} \cap A)}{\mathbb{P}(A)} \text{ for measurable } R \subseteq \mathbb{R} \quad (4.10.16)$$

If $\mathbb{E}(|X|) < \infty$ then the *conditional expected value* of X given A , denoted $\mathbb{E}(X | A)$, is simply the mean of this conditional distribution.

Suppose now that $\mathcal{A} = \{A_i : i \in I\}$ is a countable partition of the sample space Ω into events with positive probability. To review the jargon, $\mathcal{A} \subseteq \mathcal{F}$; the index set I is countable; $A_i \cap A_j = \emptyset$ for distinct $i, j \in I$; $\bigcup_{i \in I} A_i = \Omega$; and $\mathbb{P}(A_i) > 0$ for $i \in I$. Let $\mathcal{G} = \sigma(\mathcal{A})$, the σ -algebra generated by \mathcal{A} . The elements of \mathcal{G} are of the form $\bigcup_{j \in J} A_j$ for $J \subseteq I$. Moreover, the random variables that are measurable with respect to \mathcal{G} are precisely the variables that are constant on A_i for each $i \in I$. The σ -algebra \mathcal{G} is said to be *countably generated*.

If $B \in \mathcal{F}$ then $\mathbb{P}(B | \mathcal{G})$ is the random variable whose value on A_i is $\mathbb{P}(B | A_i)$ for each $i \in I$.

Proof

Let U denote the random variable that takes the value $\mathbb{P}(B | A_i)$ on A_i for each $i \in I$. First, U is measurable with respect to \mathcal{G} since U is constant on A_i for each $i \in I$. So we just need to show that $\mathbb{E}(U; A) = \mathbb{P}(A \cap B)$ for each $A \in \mathcal{G}$. Thus, let $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$. Then

$$\mathbb{E}(U; A) = \sum_{j \in J} \mathbb{E}(U; A_j) = \sum_{j \in J} \mathbb{P}(B | A_j) \mathbb{P}(A_j) = \mathbb{P}(A \cap B) \quad (4.10.17)$$

In this setting, the version of Bayes' theorem in (15) reduces to the usual elementary formulation: For $i \in I$, $\mathbb{E}[\mathbb{P}(B | \mathcal{G}); A_i] = \mathbb{P}(A_i) \mathbb{P}(B | A_i)$ and $\mathbb{E}[\mathbb{P}(B | \mathcal{G})] = \sum_{j \in I} \mathbb{P}(A_j) \mathbb{P}(B | A_j)$. Hence

$$\mathbb{P}(A_i | B) = \frac{\mathbb{P}(A_i) \mathbb{P}(B | A_i)}{\sum_{j \in I} \mathbb{P}(A_j) \mathbb{P}(B | A_j)} \quad (4.10.18)$$

If X is a random variable with $\mathbb{E}(|X|) < \infty$, then $\mathbb{E}(X | \mathcal{G})$ is the random variable whose value on A_i is $\mathbb{E}(X | A_i)$ for each $i \in I$.

Proof

Let U denote the random variable that takes the value $\mathbb{E}(X | A_i)$ on A_i for each $i \in I$. First, U is measurable with respect to \mathcal{G} since U is constant on A_i for each $i \in I$. So we just need to show that $\mathbb{E}(U; A) = \mathbb{E}(X; A)$ for each $A \in \mathcal{G}$. Thus, let $A = \bigcup_{j \in J} A_j$ where $J \subseteq I$. Then

$$\mathbb{E}(U; A) = \sum_{j \in J} \mathbb{E}(U; A_j) = \sum_{j \in J} \mathbb{E}(X | A_j) \mathbb{P}(A_j) = \mathbb{E}(X; A) \quad (4.10.19)$$

The previous examples would apply to $\mathcal{G} = \sigma(Y)$ if Y is a discrete random variable taking values in a countable set T . In this case, the partition is simply $\mathcal{A} = \{\{Y = y\} : y \in T\}$. On the other hand, suppose that Y is a random variable taking values in a general set T with σ -algebra \mathcal{T} . The real-valued random variables that are measurable with respect to $\mathcal{G} = \sigma(Y)$ are (up to equivalence) the measurable, real-valued functions of Y .

Specializing further, Suppose that X takes values in $S \subseteq \mathbb{R}$, Y takes values in $T \subseteq \mathbb{R}^n$ (where S and T are Lebesgue measurable) and that (X, Y) has a joint continuous distribution with probability density function f . Then Y has probability density function h given by

$$h(y) = \int_S f(x, y) dx, \quad y \in T \quad (4.10.20)$$

Assume that $h(y) > 0$ for $y \in T$. Then for $y \in T$, a conditional probability density function of X given $Y = y$ is defined by

$$g(x | y) = \frac{f(x, y)}{h(y)}, \quad x \in S \quad (4.10.21)$$

This is precisely the setting of our elementary discussion of conditional expected value. If $\mathbb{E}(|X|) < \infty$ then we usually write $\mathbb{E}(X | Y)$ instead of the clunkier $\mathbb{E}[X | \sigma(Y)]$.

In this setting above suppose that $\mathbb{E}(|X|) < \infty$. Then

$$\mathbb{E}(X | Y) = \int_S x g(x | Y) dx \quad (4.10.22)$$

Proof

Once again, we show that the integral on the right satisfies the properties in the [definition](#) for $\mathbb{E}(X | Y) = \mathbb{E}[X | \sigma(Y)]$. First, $y \mapsto \int_S x g(x | y) dx$ is measurable as a function from T into \mathbb{R} and hence the random variable $\int_S x g(x | Y) dx$ is a measurable function of Y and so is measurable with respect to $\sigma(Y)$. Next suppose that $B \in \sigma(Y)$. Then $B = \{Y \in A\}$ for some $A \in \mathcal{F}$. Then

$$\begin{aligned} \mathbb{E} \left[\int_S x g(x | Y) dx; B \right] &= \mathbb{E} \left[\int_S x g(x | Y) dx; Y \in A \right] \\ &= \mathbb{E} \left[\int_S x \frac{f(x, y)}{h(y)} dx; Y \in A \right] = \int_A \int_S x \frac{f(x, y)}{h(y)} h(y) dx dy \\ &= \int_{S \times A} x f(x, y) d(x, y) = \mathbb{E}(X; Y \in A) = \mathbb{E}(X; B) \end{aligned}$$

Best Predictor

In our elementary treatment of conditional expected value, we showed that the conditional expected value of a real-valued random variable X given a general random variable Y is the best predictor of X , in the least squares sense, among all real-valued functions of Y . A more careful statement is that $\mathbb{E}(X | Y)$ is the best predictor of X among all real-valued random variables that are *measurable* with respect to $\sigma(Y)$. Thus, it should come as not surprise that if \mathcal{G} is a sub σ -algebra of \mathcal{F} , then $\mathbb{E}(X | \mathcal{G})$ is the best predictor of X , in the least squares sense, among all real-valued random variables that are measurable with respect to \mathcal{G} . We will show that this is indeed the case in this subsection. The proofs are very similar to the ones given in the elementary section. For the rest of this discussion, we assume that \mathcal{G} is a sub σ -algebra of \mathcal{F} and that all random variables mentioned are real valued.

Suppose that X and U are random variables with $\mathbb{E}(|X|) < \infty$ and $\mathbb{E}(|XU|) < \infty$ and that U is measurable with respect to \mathcal{G} . Then $X - \mathbb{E}(X | \mathcal{G})$ and U are uncorrelated.

Proof

Note that $X - \mathbb{E}(X | \mathcal{G})$ has mean 0 by the [mean property](#). Using the [properties that characterize](#) $\mathbb{E}(X | \mathcal{G})$ we have

$$\text{cov}[X - \mathbb{E}(X | \mathcal{G}), U] = \mathbb{E}[U[X - \mathbb{E}(X | \mathcal{G})]] = \mathbb{E}(UX) - \mathbb{E}[U\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(UX) - \mathbb{E}(UX) = 0 \quad (4.10.23)$$

The next result is the main one: $\mathbb{E}(X | \mathcal{G})$ is closer to X in the mean square sense than any other random variable that is measurable with respect to \mathcal{G} . Thus, if \mathcal{G} represents the information that we have, then $\mathbb{E}(X | \mathcal{G})$ is the best we can do in estimating X .

Suppose that X and U are random variables with $\mathbb{E}(X^2) < \infty$ and $\mathbb{E}(U^2) < \infty$ and that U is measurable with respect to \mathcal{G} . Then

1. $\mathbb{E}([X - \mathbb{E}(X | \mathcal{G})]^2) \leq \mathbb{E}([X - U]^2)$.
2. Equality holds if and only if $\mathbb{P}[U = \mathbb{E}(X | \mathcal{G})] = 1$, so U and $\mathbb{E}(X | \mathcal{G})$ are equivalent.

Proof

1. Note that

$$\mathbb{E}([X - U]^2) = \mathbb{E}([X - \mathbb{E}(X | \mathcal{G}) + \mathbb{E}(X | \mathcal{G}) - U]^2) \quad (4.10.24)$$

$$= \mathbb{E}([X - \mathbb{E}(X | \mathcal{G})]^2) + 2\mathbb{E}([X - \mathbb{E}(X | \mathcal{G})][\mathbb{E}(X | \mathcal{G}) - U]) + \mathbb{E}([\mathbb{E}(X | \mathcal{G}) - U]^2) \quad (4.10.25)$$

By [mean property](#), $X - \mathbb{E}(X | \mathcal{G})$ has mean 0, so the middle term in the displayed equation is $2\text{cov}[X - \mathbb{E}(X | \mathcal{G}), \mathbb{E}(X | \mathcal{G}) - U]$. But $\mathbb{E}(X | \mathcal{G}) - U$ is \mathcal{G} -measurable and hence this covariance is 0 by [uncorrelated property](#). Therefore

$$\mathbb{E}([X - U]^2) = \mathbb{E}([X - \mathbb{E}(X | \mathcal{G})]^2) + \mathbb{E}([\mathbb{E}(X | \mathcal{G}) - U]^2) \geq \mathbb{E}([X - \mathbb{E}(X | \mathcal{G})]^2) \quad (4.10.26)$$

2. Equality holds if and only if $\mathbb{E}([\mathbb{E}(X | \mathcal{G}) - U]^2) = 0$ if and only if $\mathbb{P}[U = \mathbb{E}(X | \mathcal{G})] = 1$

Conditional Variance

Once again, we assume that \mathcal{G} is a sub σ -algebra of \mathcal{F} and that all random variables mentioned are real valued, unless otherwise noted. It's natural to define the conditional variance of a random variable given \mathcal{G} in the same way as ordinary variance, but with all expected values conditioned on \mathcal{G} .

Suppose that X is a random variable with $\mathbb{E}(X^2) < \infty$. The *conditional variance* of X given \mathcal{G} is

$$\text{var}(X | \mathcal{G}) = \mathbb{E} \left([X - \mathbb{E}(X | \mathcal{G})]^2 \middle| \mathcal{G} \right) \quad (4.10.27)$$

Like all conditional expected values relative to \mathcal{G} , $\text{var}(X | \mathcal{G})$ is a random variable that is measurable with respect to \mathcal{G} and is unique up to equivalence. The first property is analogous to the computational formula for ordinary variance.

Suppose again that X is a random variable with $\mathbb{E}(X^2) < \infty$. Then

$$\text{var}(X | \mathcal{G}) = \mathbb{E}(X^2 | \mathcal{G}) - [\mathbb{E}(X | \mathcal{G})]^2 \quad (4.10.28)$$

Proof

Expanding the square in the definition and using basic properties of conditional expectation, we have

$$\text{var}(X | \mathcal{G}) = \mathbb{E}(X^2 - 2X\mathbb{E}(X | \mathcal{G}) + [\mathbb{E}(X | \mathcal{G})]^2 | \mathcal{G}) = \mathbb{E}(X^2 | \mathcal{G}) - 2\mathbb{E}[X\mathbb{E}(X | \mathcal{G}) | \mathcal{G}] + \mathbb{E}([\mathbb{E}(X | \mathcal{G})]^2 | \mathcal{G}) \quad (4.10.29)$$

$$= \mathbb{E}(X^2 | \mathcal{G}) - 2\mathbb{E}(X | \mathcal{G})\mathbb{E}(X | \mathcal{G}) + [\mathbb{E}(X | \mathcal{G})]^2 = \mathbb{E}(X^2 | \mathcal{G}) - [\mathbb{E}(X | \mathcal{G})]^2 \quad (4.10.30)$$

Next is a formula for the ordinary variance in terms of conditional variance and expected value.

Suppose again that X is a random variable with $\mathbb{E}(X^2) < \infty$. Then

$$\text{var}(X) = \mathbb{E}[\text{var}(X | \mathcal{G})] + \text{var}[\mathbb{E}(X | \mathcal{G})] \quad (4.10.31)$$

Proof

From the previous theorem and properties of conditional expected value we have $\mathbb{E}[\text{var}(X | \mathcal{G})] = \mathbb{E}(X^2) - \mathbb{E}([\mathbb{E}(X | \mathcal{G})]^2)$. But $\mathbb{E}(X^2) = \text{var}(X) + [\mathbb{E}(X)]^2$ and similarly, $\mathbb{E}([\mathbb{E}(X | \mathcal{G})]^2) = \text{var}[\mathbb{E}(X | \mathcal{G})] + (\mathbb{E}[\mathbb{E}(X | \mathcal{G})])^2$. But also, $\mathbb{E}[\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(X)$ so substituting we get $\mathbb{E}[\text{var}(X | \mathcal{G})] = \text{var}(X) - \text{var}[\mathbb{E}(X | \mathcal{G})]$.

So the variance of X is the expected conditional variance plus the variance of the conditional expected value. This result is often a good way to compute $\text{var}(X)$ when we know the conditional distribution of X given \mathcal{G} . In turn, this property leads to a formula for the mean square error when $\mathbb{E}(X | \mathcal{G})$ is thought of as a predictor of X .

Suppose again that X is a random variable with $\mathbb{E}(X^2) < \infty$.

$$\mathbb{E}([X - \mathbb{E}(X | \mathcal{G})]^2) = \text{var}(X) - \text{var}[\mathbb{E}(X | \mathcal{G})] \quad (4.10.32)$$

Proof

From the [definition](#) and from the [mean property](#) and [variance formula](#),

$$\mathbb{E}([X - \mathbb{E}(X | \mathcal{G})]^2) = \mathbb{E}[\text{var}(X | \mathcal{G})] = \text{var}(X) - \text{var}[\mathbb{E}(X | \mathcal{G})] \quad (4.10.33)$$

Let us return to the study of predictors of the real-valued random variable X , and compare them in terms of mean square error.

Suppose again that X is a random variable with $\mathbb{E}(X^2) < \infty$.

1. The best constant predictor of X is $\mathbb{E}(X)$ with mean square error $\text{var}(X)$.
2. If Y is another random variable with $\mathbb{E}(Y^2) < \infty$, then the best predictor of X among linear functions of Y is

$$L(X | Y) = \mathbb{E}(X) + \frac{\text{cov}(X, Y)}{\text{var}(Y)}[Y - \mathbb{E}(Y)] \quad (4.10.34)$$

with mean square error $\text{var}(X)[1 - \text{cor}^2(X, Y)]$.

3. If Y is a (general) random variable, then the best predictor of X among all real-valued functions of Y with finite variance is $\mathbb{E}(X | Y)$ with mean square error $\text{var}(X) - \text{var}[\mathbb{E}(X | Y)]$.
4. If \mathcal{G} is a sub σ -algebra of \mathcal{F} , then the best predictor of X among random variables with finite variance that are measurable with respect to \mathcal{G} is $\mathbb{E}(X | \mathcal{G})$ with mean square error $\text{var}(X) - \text{var}[\mathbb{E}(X | \mathcal{G})]$.

Of course, (a) is a special case of (d) with $\mathcal{G} = \{\emptyset, \Omega\}$ and (c) is a special case of (d) with $\mathcal{G} = \sigma(Y)$. Only (b), the linear case, cannot be interpreted in terms of conditioning with respect to a σ -algebra.

Conditional Covariance

Suppose again that \mathcal{G} is a sub σ -algebra of \mathcal{F} . The conditional covariance of two random variables is defined like the ordinary covariance, but with all expected values conditioned on \mathcal{G} .

Suppose that X and Y are random variables with $\mathbb{E}(X^2) < \infty$ and $\mathbb{E}(Y^2) < \infty$. The *conditional covariance* of X and Y given \mathcal{G} is defined as

$$\text{cov}(X, Y | \mathcal{G}) = \mathbb{E}\left([X - \mathbb{E}(X | \mathcal{G})][Y - \mathbb{E}(Y | \mathcal{G})] | \mathcal{G}\right) \quad (4.10.35)$$

So $\text{cov}(X, Y | \mathcal{G})$ is a random variable that is measurable with respect to \mathcal{G} and is unique up to equivalence. As should be the case, conditional covariance generalizes conditional variance.

Suppose that X is a random variable with $\mathbb{E}(X^2) < \infty$. Then $\text{cov}(X, X | \mathcal{G}) = \text{var}(X | \mathcal{G})$.

Proof

This follows immediately from the two definitions.

Our next result is a computational formula that is analogous to the one for standard covariance—the covariance is the mean of the product minus the product of the means, but now with all expected values conditioned on \mathcal{G} :

Suppose again that X and Y are random variables with $\mathbb{E}(X^2) < \infty$ and $\mathbb{E}(Y^2) < \infty$. Then

$$\text{cov}(X, Y | \mathcal{G}) = \mathbb{E}(XY | \mathcal{G}) - \mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G}) \quad (4.10.36)$$

Proof

Expanding the product in the definition and using basic properties of conditional expectation, we have

$$\text{cov}(X, Y | \mathcal{G}) = \mathbb{E} \left(XY - X\mathbb{E}(Y | \mathcal{G}) - Y\mathbb{E}(X | \mathcal{G}) + \mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G}) \mid \mathcal{G} \right) = \mathbb{E}(XY | \mathcal{G}) - \mathbb{E}[X\mathbb{E}(Y | \mathcal{G}) | \mathcal{G}] - \mathbb{E}[Y\mathbb{E}(X | \mathcal{G}) | \mathcal{G}] \quad (4.10.37)$$

$$= \mathbb{E}(XY | \mathcal{G}) - \mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G}) - \mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G}) + \mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G}) = \mathbb{E}(XY | \mathcal{G}) - \mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G}) \quad (4.10.38)$$

Our next result shows how to compute the ordinary covariance of X and Y by *conditioning on X* .

Suppose again that X and Y are random variables with $\mathbb{E}(X^2) < \infty$ and $\mathbb{E}(Y^2) < \infty$. Then

$$\text{cov}(X, Y) = \mathbb{E}[\text{cov}(X, Y | \mathcal{G})] + \text{cov}[\mathbb{E}(X | \mathcal{G}), \mathbb{E}(Y | \mathcal{G})] \quad (4.10.39)$$

Proof

From (29) and properties of conditional expected value we have

$$\mathbb{E}[\text{cov}(X, Y | \mathcal{G})] = \mathbb{E}(XY) - \mathbb{E}[\mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G})] \quad (4.10.40)$$

But $\mathbb{E}(XY) = \text{cov}(X, Y) + \mathbb{E}(X)\mathbb{E}(Y)$ and similarly,

$$\mathbb{E}[\mathbb{E}(X | \mathcal{G})\mathbb{E}(Y | \mathcal{G})] = \text{cov}[\mathbb{E}(X | \mathcal{G}), \mathbb{E}(Y | \mathcal{G})] + \mathbb{E}[\mathbb{E}(X | \mathcal{G})]\mathbb{E}[\mathbb{E}(Y | \mathcal{G})] \quad (4.10.41)$$

But also, $\mathbb{E}[\mathbb{E}(X | \mathcal{G})] = \mathbb{E}(X)$ and $\mathbb{E}[\mathbb{E}(Y | \mathcal{G})] = \mathbb{E}(Y)$ so substituting we get

$$\mathbb{E}[\text{cov}(X, Y | \mathcal{G})] = \text{cov}(X, Y) - \text{cov}[\mathbb{E}(X | \mathcal{G}), \mathbb{E}(Y | \mathcal{G})] \quad (4.10.42)$$

Thus, the covariance of X and Y is the expected conditional covariance plus the covariance of the conditional expected values. This result is often a good way to compute $\text{cov}(X, Y)$ when we know the conditional distribution of (X, Y) given \mathcal{G} .

This page titled [4.10: Conditional Expected Value Revisited](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.11: Vector Spaces of Random Variables

Basic Theory

Many of the concepts in this chapter have elegant interpretations if we think of real-valued random variables as vectors in a vector space. In particular, variance and higher moments are related to the concept of norm and distance, while covariance is related to inner product. These connections can help unify and illuminate some of the ideas in the chapter from a different point of view. Of course, real-valued random variables are simply measurable, real-valued functions defined on the sample space, so much of the discussion in this section is a special case of our discussion of function spaces in the chapter on Distributions, but recast in the notation of probability.

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Thus, Ω is the set of outcomes, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . Our basic *vector space* \mathcal{V} consists of all real-valued random variables defined on $(\Omega, \mathcal{F}, \mathbb{P})$ (that is, defined for the experiment). Recall that random variables X_1 and X_2 are equivalent if $\mathbb{P}(X_1 = X_2) = 1$, in which case we write $X_1 \equiv X_2$. We consider two such random variables as the same vector, so that technically, our vector space consists of *equivalence classes* under this equivalence relation. The *addition operator* corresponds to the usual addition of two real-valued random variables, and the operation of *scalar multiplication* corresponds to the usual multiplication of a real-valued random variable by a real (non-random) number. These operations are compatible with the equivalence relation in the sense that if $X_1 \equiv X_2$ and $Y_1 \equiv Y_2$ then $X_1 + Y_1 \equiv X_2 + Y_2$ and $cX_1 \equiv cX_2$ for $c \in \mathbb{R}$. In short, the vector space \mathcal{V} is well-defined.

Norm

Suppose that $k \in [1, \infty)$. The k norm of $X \in \mathcal{V}$ is defined by

$$\|X\|_k = \left[\mathbb{E}(|X|^k) \right]^{1/k} \quad (4.11.1)$$

Thus, $\|X\|_k$ is a measure of the size of X in a certain sense, and of course it's possible that $\|X\|_k = \infty$. The following theorems establish the fundamental properties. The first is the *positive property*.

Suppose again that $k \in [1, \infty)$. For $X \in \mathcal{V}$,

1. $\|X\|_k \geq 0$
2. $\|X\|_k = 0$ if and only if $\mathbb{P}(X = 0) = 1$ (so that $X \equiv 0$).

Proof

These results follow from the basic inequality properties of expected value. First $|X|^k \geq 0$ with probability 1, so $\mathbb{E}(|X|^k) \geq 0$. In addition, $\mathbb{E}(|X|^k) = 0$ if and only if $\mathbb{P}(X = 0) = 1$.

The next result is the *scaling property*.

Suppose again that $k \in [1, \infty)$. Then $\|cX\|_k = |c| \|X\|_k$ for $X \in \mathcal{V}$ and $c \in \mathbb{R}$.

Proof

$$\|cX\|_k = \left[\mathbb{E}(|cX|^k) \right]^{1/k} = \left[\mathbb{E}(|c|^k |X|^k) \right]^{1/k} = \left[|c|^k \mathbb{E}(|X|^k) \right]^{1/k} = |c| \left[\mathbb{E}(|X|^k) \right]^{1/k} = |c| \|X\|_k \quad (4.11.2)$$

The next result is *Minkowski's inequality*, named for Hermann Minkowski, and also known as the *triangle inequality*.

Suppose again that $k \in [1, \infty)$. Then $\|X + Y\|_k \leq \|X\|_k + \|Y\|_k$ for $X, Y \in \mathcal{V}$.

Proof

The first quadrant $S = \{(x, y) \in \mathbb{R}^2 : x \geq 0, y \geq 0\}$ is a convex set and $g(x, y) = (x^{1/k} + y^{1/k})^k$ is concave on S . From Jensen's inequality, if U and V are nonnegative random variables, then

$$\mathbb{E} \left[(U^{1/k} + V^{1/k})^k \right] \leq \left([\mathbb{E}(U)]^{1/k} + [\mathbb{E}(V)]^{1/k} \right)^k \quad (4.11.3)$$

Letting $U = |X|^k$ and $V = |Y|^k$ and simplifying gives the result. To show that g really is concave on S , we can compute the second partial derivatives. Let $h(x, y) = x^{1/k} + y^{1/k}$ so that $g = h^k$. Then

$$g_{xx} = \frac{k-1}{k} h^{k-2} x^{1/k-2} (x^{1/k} - h) \quad (4.11.4)$$

$$g_{yy} = \frac{k-1}{k} h^{k-2} y^{1/k-2} (y^{1/k} - h) \quad (4.11.5)$$

$$g_{xy} = \frac{k-1}{k} h^{k-2} x^{1/k-1} y^{1/k-1} \quad (4.11.6)$$

Clearly $h(x, y) \geq x^{1/k}$ and $h(x, y) \geq y^{1/k}$ for $x \geq 0$ and $y \geq 0$, so g_{xx} and g_{yy} , the diagonal entries of the second derivative matrix, are nonpositive on S . A little algebra shows that the determinant of the second derivative matrix $g_{xx}g_{yy} - g_{xy}^2 = 0$ on S . Thus, the second derivative matrix of g is negative semi-definite.

It follows from the last three results that the set of random variables (again, modulo equivalence) with finite k norm forms a *subspace* of our parent vector space \mathcal{V} , and that the k norm really is a norm on this vector space.

For $k \in [1, \infty)$, \mathcal{L}_k denotes the vector space of $X \in \mathcal{V}$ with $\|X\|_k < \infty$, and with norm $\|\cdot\|_k$.

In analysis, p is often used as the index rather than k as we have used here, but p seems too much like a probability, so we have broken with tradition on this point. The \mathcal{L} is in honor of Henri Lebesgue, who developed much of this theory. Sometimes, when we need to indicate the dependence on the underlying σ -algebra \mathcal{F} , we write $\mathcal{L}_k(\mathcal{F})$. Our next result is *Lyapunov's inequality*, named for Aleksandr Lyapunov. This inequality shows that the k -norm of a random variable is increasing in k .

Suppose that $j, k \in [1, \infty)$ with $j \leq k$. Then $\|X\|_j \leq \|X\|_k$ for $X \in \mathcal{V}$.

Proof

Note that $S = \{x \in \mathbb{R} : x \geq 0\}$ is convex and $g(x) = x^{k/j}$ is convex on S . From Jensen's inequality, if U is a nonnegative random variable then $[\mathbb{E}(U)]^{k/j} \leq \mathbb{E}(U^{k/j})$. Letting $U = |X|^j$ and simplifying gives the result.

Lyapunov's inequality shows that if $1 \leq j \leq k$ and $\|X\|_k < \infty$ then $\|X\|_j < \infty$. Thus, \mathcal{L}_k is a subspace of \mathcal{L}_j .

Metric

The k norm, like any norm on a vector space, can be used to define a metric, or distance function; we simply compute the norm of the difference between two vectors.

For $k \in [1, \infty)$, the k distance (or k metric) between $X, Y \in \mathcal{V}$ is defined by

$$d_k(X, Y) = \|X - Y\|_k = \left[\mathbb{E}(|X - Y|^k) \right]^{1/k} \quad (4.11.7)$$

The following properties are analogous to the properties in [norm properties](#) (and thus very little additional work is required for the proofs). These properties show that the k metric really is a metric on \mathcal{L}_k (as always, modulo equivalence). The first is the *positive property*.

Suppose again that $k \in [1, \infty)$ $X, Y \in \mathcal{V}$. Then

1. $d_k(X, Y) \geq 0$
2. $d_k(X, Y) = 0$ if and only if $\mathbb{P}(X = Y) = 1$ (so that $X \equiv Y$ and Y).

Proof

These results follow directly from the [positive property](#).

Next is the obvious *symmetry property*:

$$d_k(X, Y) = d_k(Y, X) \text{ for } X, Y \in \mathcal{V}.$$

Next is the distance version of the *triangle inequality*.

$$d_k(X, Z) \leq d_k(X, Y) + d_k(Y, Z) \text{ for } X, Y, Z \in \mathcal{V}$$

Proof

From [Minkowski's inequality](#),

$$d_k(X, Z) = \|X - Z\|_k = \|(X - Y) + (Y - Z)\|_k \leq \|X - Y\|_k + \|Y - Z\|_k = d_k(X, Y) + d_k(Y, Z) \quad (4.11.8)$$

The last three properties mean that d_k is indeed a metric on \mathcal{L}_k for $k \geq 1$. In particular, note that the standard deviation is simply the 2-distance from X to its mean $\mu = \mathbb{E}(X)$:

$$\text{sd}(X) = d_2(X, \mu) = \|X - \mu\|_2 = \sqrt{\mathbb{E}[(X - \mu)^2]} \quad (4.11.9)$$

and the variance is the square of this. More generally, the k th moment of X about a is simply the k th power of the k -distance from X to a . The 2-distance is especially important for reasons that will become clear below, in the discussion of [inner product](#). This distance is also called the *root mean square distance*.

Center and Spread Revisited

Measures of center and measures of spread are best thought of together, in the context of a *measure of distance*. For a real-valued random variable X , we first try to find the constants $t \in \mathbb{R}$ that are closest to X , as measured by the given distance; any such t is a *measure of center* relative to the distance. The minimum distance itself is the corresponding *measure of spread*.

Let us apply this procedure to the 2-distance.

For $X \in \mathcal{L}_2$, define the *root mean square error* function by

$$d_2(X, t) = \|X - t\|_2 = \sqrt{\mathbb{E}[(X - t)^2]}, \quad t \in \mathbb{R} \quad (4.11.10)$$

For $X \in \mathcal{L}_2$, $d_2(X, t)$ is minimized when $t = \mathbb{E}(X)$ and the minimum value is $\text{sd}(X)$.

Proof

Note that the minimum value of $d_2(X, t)$ occurs at the same points as the minimum value of $d_2^2(X, t) = \mathbb{E}[(X - t)^2]$ (this is the *mean square error* function). Expanding and taking expected values term by term gives

$$\mathbb{E}[(X - t)^2] = \mathbb{E}(X^2) - 2t\mathbb{E}(X) + t^2 \quad (4.11.11)$$

This is a quadratic function of t and hence the graph is a parabola opening upward. The minimum occurs at $t = \mathbb{E}(X)$, and the minimum value is $\text{var}(X)$. Hence the minimum value of $t \mapsto d_2(X, t)$ also occurs at $t = \mathbb{E}(X)$ and the minimum value is $\text{sd}(X)$.

We have seen this computation several times before. The best constant predictor of X is $\mathbb{E}(X)$, with mean square error $\text{var}(X)$. The physical interpretation of this result is that the moment of inertia of the mass distribution of X about t is minimized when $t = \mu$, the center of mass. Next, let us apply our procedure to the 1-distance.

For $X \in \mathcal{L}_1$, define the *mean absolute error* function by

$$d_1(X, t) = \|X - t\|_1 = \mathbb{E}[|X - t|], \quad t \in \mathbb{R} \quad (4.11.12)$$

We will show that $d_1(X, t)$ is minimized when t is any median of X . (Recall that the set of medians of X forms a closed, bounded interval.) We start with a discrete case, because it's easier and has special interest.

Suppose that $X \in \mathcal{L}_1$ has a discrete distribution with values in a finite set $S \subseteq \mathbb{R}$. Then $d_1(X, t)$ is minimized when t is any median of X .

Proof

Note first that $\mathbb{E}(|X - t|) = \mathbb{E}(t - X, X \leq t) + \mathbb{E}(X - t, X > t)$. Hence $\mathbb{E}(|X - t|) = a_t t + b_t$, where $a_t = 2\mathbb{P}(X \leq t) - 1$ and where $b_t = \mathbb{E}(X) - 2\mathbb{E}(X, X \leq t)$. Note that $\mathbb{E}(|X - t|)$ is a continuous, piecewise linear function of t , with corners at the values in S . That is, the function is a *linear spline*. Let m be the smallest median of X . If $t < m$ and $t \notin S$, then the slope of the linear piece at t is negative. Let M be the largest median of X . If $t > M$ and $t \notin S$, then the slope of the linear piece at t is positive. If $t \in (m, M)$ then the slope of the linear piece at t is 0. Thus $\mathbb{E}(|X - t|)$ is minimized for every t in the median interval $[m, M]$.

The last result shows that mean absolute error has a couple of basic deficiencies as a measure of error:

- The function may not be smooth (differentiable).
- The function may not have a unique minimizing value of t .

Indeed, when X does not have a unique median, there is no compelling reason to choose one value in the median interval, as the measure of center, over any other value in the interval.

Suppose now that $X \in \mathcal{L}_1$ has a general distribution on \mathbb{R} . Then $d_1(X, t)$ is minimized when t is any median of X .

Proof

Let $s, t \in \mathbb{R}$. Suppose first that $s < t$. Computing the expected value over the events $X \leq s$, $s < X \leq t$, and $X > t$, and simplifying gives

$$\mathbb{E}(|X - t|) = \mathbb{E}(|X - s|) + (t - s) [2\mathbb{P}(X \leq s) - 1] + 2\mathbb{E}(t - X, s < X \leq t) \quad (4.11.13)$$

Suppose next that $t < s$. Using similar methods gives

$$\mathbb{E}(|X - t|) = \mathbb{E}(|X - s|) + (t - s) [2\mathbb{P}(X < s) - 1] + 2\mathbb{E}(X - t, t \leq X < s) \quad (4.11.14)$$

Note that the last terms on the right in these equations are nonnegative. If we take s to be a median of X , then the middle terms on the right in the equations are also nonnegative. Hence if s is a median of X and t is any other number then $\mathbb{E}(|X - t|) \geq \mathbb{E}(|X - s|)$.

Convergence

Whenever we have a measure of distance, we automatically have a criterion for convergence.

Suppose that $X_n \in \mathcal{L}_k$ for $n \in \mathbb{N}_+$ and that $X \in \mathcal{L}_k$, where $k \in [1, \infty)$. Then $X_n \rightarrow X$ as $n \rightarrow \infty$ in k th mean if $X_n \rightarrow X$ as $n \rightarrow \infty$ in the vector space \mathcal{L}_k . That is,

$$d_k(X_n, X) = \|X_n - X\|_k \rightarrow 0 \text{ as } n \rightarrow \infty \quad (4.11.15)$$

or equivalently $\mathbb{E}(|X_n - X|^k) \rightarrow 0$ as $n \rightarrow \infty$.

When $k = 1$, we simply say that $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean; when $k = 2$, we say that $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean square. These are the most important special cases.

Suppose that $1 \leq j \leq k$. If $X_n \rightarrow X$ as $n \rightarrow \infty$ in k th mean then $X_n \rightarrow X$ as $n \rightarrow \infty$ in j th mean.

Proof

This follows from [Lyapunov's inequality](#). Note that $0 \leq d_j(X_n, X) \leq d_k(X_n, X) \rightarrow 0$ as $n \rightarrow \infty$.

Convergence in k th mean implies that the k norms converge.

Suppose that $X_n \in \mathcal{L}_k$ for $n \in \mathbb{N}_+$ and that $X \in \mathcal{L}_k$, where $k \in [1, \infty)$. If $X_n \rightarrow X$ as $n \rightarrow \infty$ in k th mean then $\|X_n\|_k \rightarrow \|X\|_k$ as $n \rightarrow \infty$. Equivalently, if $\mathbb{E}(|X_n - X|^k) \rightarrow 0$ as $n \rightarrow \infty$ then $\mathbb{E}(|X_n|^k) \rightarrow \mathbb{E}(|X|^k)$ as $n \rightarrow \infty$.

Proof

This is a simple consequence of the reverse triangle inequality, which holds in any normed vector space. The general result is that if a sequence of vectors in a normed vector space converge then the norms converge. In our notation here,

$$|\|X_n\|_k - \|X\|_k| \leq \|X_n - X\|_k \quad (4.11.16)$$

so if the right side converges to 0 as $n \rightarrow \infty$, then so does the left side.

The converse is not true; a [counterexample](#) is given below. Our next result shows that convergence in mean is stronger than convergence in probability.

Suppose that $X_n \in \mathcal{L}_1$ for $n \in \mathbb{N}_+$ and that $X \in \mathcal{L}_1$. If $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean, then $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability.

Proof

This follows from Markov's inequality. For $\epsilon > 0$, $0 \leq \mathbb{P}(|X_n - X| > \epsilon) \leq \mathbb{E}(|X_n - X|) / \epsilon \rightarrow 0$ as $n \rightarrow \infty$.

The converse is not true. That is, convergence with probability 1 does not imply convergence in k th mean; a [counterexample](#) is given below. Also convergence in k th mean does not imply convergence with probability 1; a [counterexample](#) to this is given below. In summary, the implications in the various modes of convergence are shown below; no other implications hold in general.

- Convergence with probability 1 implies convergence in probability.
- Convergence in k th mean implies convergence in j th mean if $j \leq k$.
- Convergence in k th mean implies convergence in probability.
- Convergence in probability implies convergence in distribution.

However, the next section on uniformly integrable variables gives a condition under which convergence in probability implies convergence in mean.

Inner Product

The vector space \mathcal{L}_2 of real-valued random variables on $(\Omega, \mathcal{F}, \mathbb{P})$ (modulo equivalence of course) with finite second moment is special, because it's the only one in which the norm corresponds to an inner product.

The inner product of $X, Y \in \mathcal{L}_2$ is defined by

$$\langle X, Y \rangle = \mathbb{E}(XY) \quad (4.11.17)$$

The following results are analogous to the basic properties of covariance, and show that this definition really does give an inner product on the vector space

For $X, Y, Z \in \mathcal{L}_2$ and $a \in \mathbb{R}$,

1. $\langle X, Y \rangle = \langle Y, X \rangle$, the *symmetric property*.
2. $\langle X, X \rangle \geq 0$ and $\langle X, X \rangle = 0$ if and only if $\mathbb{P}(X = 0) = 1$ (so that $X \equiv 0$), the *positive property*.
3. $\langle aX, Y \rangle = a\langle X, Y \rangle$, the *scaling property*.
4. $\langle X + Y, Z \rangle = \langle X, Z \rangle + \langle Y, Z \rangle$, the *additive property*.

Proof

1. This property is trivial from the definition.
2. Note that $\mathbb{E}(X^2) \geq 0$ and $\mathbb{E}(X^2) = 0$ if and only if $\mathbb{P}(X = 0) = 1$.
3. This follows from the scaling property of expected value: $\mathbb{E}(aXY) = a\mathbb{E}(XY)$
4. This follows from the additive property of expected value: $\mathbb{E}[(X + Y)Z] = \mathbb{E}(XZ) + \mathbb{E}(YZ)$.

From parts (a), (c), and (d) it follows that inner product is *bi-linear*, that is, linear in each variable with the other fixed. Of course bi-linearity holds for any inner product on a vector space. Covariance and correlation can easily be expressed in terms of this inner product. The covariance of two random variables is the inner product of the corresponding *centered* variables. The correlation is the inner product of the corresponding standard scores.

For $X, Y \in \mathcal{L}_2$,

1. $\text{cov}(X, Y) = \langle X - \mathbb{E}(X), Y - \mathbb{E}(Y) \rangle$
2. $\text{cor}(X, Y) = \langle [X - \mathbb{E}(X)]/\text{sd}(X), [Y - \mathbb{E}(Y)]/\text{sd}(Y) \rangle$

Proof

1. This is simply a restatement of the definition of covariance.
2. This is a restatement of the fact that the correlation of two variables is the covariance of their corresponding standard scores.

Thus, real-valued random variables X and Y are uncorrelated if and only if the centered variables $X - \mathbb{E}(X)$ and $Y - \mathbb{E}(Y)$ are *perpendicular* or *orthogonal* as elements of \mathcal{L}_2 .

For $X \in \mathcal{L}_2$, $\langle X, X \rangle = \|X\|_2^2 = \mathbb{E}(X^2)$.

Thus, the norm associated with the inner product is the 2-norm studied above, and corresponds to the *root mean square* operation on a random variable. This fact is a fundamental reason why the 2-norm plays such a special, honored role; of all the k -norms, only the 2-norm corresponds to an inner product. In turn, this is one of the reasons that *root mean square difference* is of fundamental importance in probability and statistics. Technically, the vector space \mathcal{L}_2 is a *Hilbert space*, named for David Hilbert.

The next result is *Hölder's inequality*, named for Otto Hölder.

Suppose that $j, k \in [1, \infty)$ and $\frac{1}{j} + \frac{1}{k} = 1$. For $X \in \mathcal{L}_j$ and $Y \in \mathcal{L}_k$,

$$\langle |X|, |Y| \rangle \leq \|X\|_j \|Y\|_k \quad (4.11.18)$$

Proof

Note that $S = \{(x, y) \in \mathbb{R}^2 : x \geq 0, y \geq 0\}$ is a convex set and $g(x, y) = x^{1/j} y^{1/k}$ is concave on S . From Jensen's inequality, if U and V are nonnegative random variables then $\mathbb{E}(U^{1/j} V^{1/k}) \leq [\mathbb{E}(U)]^{1/j} [\mathbb{E}(V)]^{1/k}$. Substituting $U = |X|^j$ and $V = |Y|^k$ gives the result.

To show that g really is concave on S , we compute the second derivative matrix:

$$\begin{bmatrix} (1/j)(1/j-1)x^{1/j-2}y^{1/k} & (1/j)(1/k)x^{1/j-1}y^{1/k-1} \\ (1/j)(1/k)x^{1/j-1}y^{1/k-1} & (1/k)(1/k-1)x^{1/j}y^{1/k-2} \end{bmatrix} \quad (4.11.19)$$

Since $1/j < 1$ and $1/k < 1$, the diagonal entries are negative on S . The determinant simplifies to

$$(1/j)(1/k)x^{2/j-2}y^{2/k-2}[1 - (1/j + 1/k)] = 0 \quad (4.11.20)$$

In the context of the last theorem, j and k are called *conjugate exponents*. If we let $j = k = 2$ in Hölder's inequality, then we get the *Cauchy-Schwarz inequality*, named for Augustin Cauchy and Karl Schwarz: For $X, Y \in \mathcal{L}_2$,

$$\mathbb{E}(|X| |Y|) \leq \sqrt{\mathbb{E}(X^2)} \sqrt{\mathbb{E}(Y^2)} \quad (4.11.21)$$

In turn, the Cauchy-Schwarz inequality is equivalent to the basic inequalities for covariance and correlations: For $X, Y \in \mathcal{L}_2$,

$$|\text{cov}(X, Y)| \leq \text{sd}(X)\text{sd}(Y), \quad |\text{cor}(X, Y)| \leq 1 \quad (4.11.22)$$

If $j, k \in [1, \infty)$ are conjugate exponents then

1. $k = \frac{j}{j-1}$.
2. $k \downarrow 1$ as $j \uparrow \infty$.

The following result is an equivalent to the identity $\text{var}(X+Y) + \text{var}(X-Y) = 2[\text{var}(X) + \text{var}(Y)]$ that we studied in the section on covariance and correlation. In the context of vector spaces, the result is known as the *parallelogram rule*:

If $X, Y \in \mathcal{L}_2$ then

$$\|X+Y\|_2^2 + \|X-Y\|_2^2 = 2\|X\|_2^2 + 2\|Y\|_2^2 \quad (4.11.23)$$

Proof

This result follows from the bi-linearity of inner product:

$$\|X+Y\|_2^2 + \|X-Y\|_2^2 = \langle X+Y, X+Y \rangle + \langle X-Y, X-Y \rangle \quad (4.11.24)$$

$$= (\langle X, X \rangle + 2\langle X, Y \rangle + \langle Y, Y \rangle) + (\langle X, X \rangle - 2\langle X, Y \rangle + \langle Y, Y \rangle) = 2\|X\|^2 + 2\|Y\|^2 \quad (4.11.25)$$

The following result is equivalent to the statement that the variance of the sum of uncorrelated variables is the sum of the variances, which again we proved in the section on covariance and correlation. In the context of vector spaces, the result is the famous *Pythagorean theorem*, named for Pythagoras of course.

If (X_1, X_2, \dots, X_n) is a sequence of random variables in \mathcal{L}_2 with $\langle X_i, X_j \rangle = 0$ for $i \neq j$ then

$$\left\| \sum_{i=1}^n X_i \right\|_2^2 = \sum_{i=1}^n \|X_i\|_2^2 \quad (4.11.26)$$

Proof

Again, this follows from the bi-linearity of inner product:

$$\left\| \sum_{i=1}^n X_i \right\|_2^2 = \left\langle \sum_{i=1}^n X_i, \sum_{j=1}^n X_j \right\rangle = \sum_{i=1}^n \sum_{j=1}^n \langle X_i, X_j \rangle \quad (4.11.27)$$

The terms with $i \neq j$ are 0 by the orthogonality assumption, so

$$\left\| \sum_{i=1}^n X_i \right\|_2^2 = \sum_{i=1}^n \langle X_i, X_i \rangle = \sum_{i=1}^n \|X_i\|_2^2 \quad (4.11.28)$$

Projections

The best linear predictor studied in the section on covariance and correlation and conditional expected values have nice interpretation in terms of projections onto subspaces of \mathcal{L}_2 . First let's review the concepts. Recall that \mathcal{U} is a *subspace* of \mathcal{L}_2 if $\mathcal{U} \subseteq \mathcal{L}_2$ and \mathcal{U} is also a vector space (under the same operations of addition and scalar multiplication). To show that $\mathcal{U} \subseteq \mathcal{L}_2$ is a subspace, we just need to show the *closure properties* (the other axioms of a vector space are inherited).

- If $U, V \in \mathcal{U}$ then $U+V \in \mathcal{U}$.
- If $U \in \mathcal{U}$ and $c \in \mathbb{R}$ then $cU \in \mathcal{U}$.

Suppose now that \mathcal{U} is a subspace of \mathcal{L}_2 and that $X \in \mathcal{L}_2$. Then the *projection* of X onto \mathcal{U} (if it exists) is the vector $V \in \mathcal{U}$ with the property that $X-V$ is perpendicular to \mathcal{U} :

$$\langle X-V, U \rangle = 0, \quad U \in \mathcal{U} \quad (4.11.29)$$

The projection has two critical properties: It is unique (if it exists) and it is the vector in \mathcal{U} closest to X . If you look at the proofs of these results, you will see that they are essentially the same as the ones used for the best predictors of X mentioned at the beginning of this subsection. Moreover, the proofs use only vector space concepts—the fact that our vectors are random variables on a probability space plays no special role.

The projection of X onto \mathcal{U} (if it exists) is unique.

Proof

Suppose that V_1 and V_2 satisfy the definition. then

$$\|V_1 - V_2\|_2^2 = \langle V_1 - V_2, V_1 - V_2 \rangle = \langle V_1 - X + X - V_2, V_1 - V_2 \rangle = \langle V_1 - X, V_1 - V_2 \rangle + \langle X - V_2, V_1 - V_2 \rangle = 0 \quad (4.11.30)$$

Hence $V_1 \equiv V_2$. The last equality in the displayed equation holds by assumption and the fact that $V_1 - V_2 \in \mathcal{U}$

Suppose that V is the projection of X onto \mathcal{U} . Then

1. $\|X - V\|_2^2 \leq \|X - U\|_2^2$ for all $U \in \mathcal{U}$.
2. Equality holds in (a) if and only if $U \equiv V$

Proof

1. If $U \in \mathcal{U}$ then

$$\|X - U\|_2^2 = \|X - V + V - U\|_2^2 = \|X - V\|_2^2 + 2\langle X - V, V - U \rangle + \|V - U\|_2^2 \quad (4.11.31)$$

But the middle terms is 0 so

$$\|X - U\|_2^2 = \|X - V\|_2^2 + \|V - U\|_2^2 \geq \|X - V\|_2^2 \quad (4.11.32)$$

2. Equality holds if and only if $\|V - U\|_2^2 = 0$, if and only if $V \equiv U$.

Now let's return to our study of best predictors of a random variable.

If $X \in \mathcal{L}_2$ then the set $\mathcal{W}_X = \{a + bX : a \in \mathbb{R}, b \in \mathbb{R}\}$ is a subspace of \mathcal{L}_2 . In fact, it is the subspace *generated* by X and 1.

Proof

Note that \mathcal{W}_X is the set of all *linear combinations* of the vectors 1 and X . If $U, V \in \mathcal{W}_X$ then $U + V \in \mathcal{W}_X$. If $U \in \mathcal{W}_X$ and $c \in \mathbb{R}$ then $cU \in \mathcal{W}_X$.

Recall that for $X, Y \in \mathcal{L}_2$, the best linear predictor of Y based on X is

$$L(Y | X) = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)} [X - \mathbb{E}(X)] \quad (4.11.33)$$

Here is the meaning of the predictor in the context of our vector spaces.

If $X, Y \in \mathcal{L}_2$ then $L(Y | X)$ is the projection of Y onto \mathcal{W}_X .

Proof

Note first that $L(Y | X) \in \mathcal{W}_X$. Thus, we just need to show that $Y - L(Y | X)$ is perpendicular to \mathcal{W}_X . For this, it suffices to show

1. $\langle Y - L(Y | X), X \rangle = 0$
2. $\langle Y - L(Y | X), 1 \rangle = 0$

We have already done this in the earlier sections, but for completeness, we do it again. Note that $\mathbb{E}(X[X - \mathbb{E}(X)]) = \text{var}(X)$. Hence $\mathbb{E}[XL(Y | X)] = \mathbb{E}(X)\mathbb{E}(Y) + \text{cov}(X, Y) = \mathbb{E}(XY)$. This gives (a). By linearity, $\mathbb{E}[L(Y | X)] = \mathbb{E}(Y)$ so (b) holds as well.

The previous result is actually just the random variable version of the standard formula for the projection of a vector onto a space spanned by two other vectors. Note that 1 is a unit vector and that $X_0 = X - \mathbb{E}(X) = X - \langle X, 1 \rangle 1$ is perpendicular to 1. Thus, $L(Y | X)$ is just the sum of the projections of Y onto 1 and X_0 :

$$L(Y | X) = \langle Y, 1 \rangle 1 + \frac{\langle Y, X_0 \rangle}{\langle X_0, X_0 \rangle} X_0 \quad (4.11.34)$$

Suppose now that \mathcal{G} is a sub σ -algebra of \mathcal{F} . Of course if $X : \Omega \rightarrow \mathbb{R}$ is \mathcal{G} -measurable then X is \mathcal{F} -measurable, so $\mathcal{L}_2(\mathcal{G})$ is a subspace of $\mathcal{L}_2(\mathcal{F})$.

If $X \in \mathcal{L}_2(\mathcal{F})$ then $\mathbb{E}(X | \mathcal{G})$ is the projection of X onto $\mathcal{L}_2(\mathcal{G})$.

Proof

This is essentially the definition of $\mathbb{E}(X | \mathcal{G})$ as the only (up to equivalence) random variable in $\mathcal{L}_2(\mathcal{G})$ with $\mathbb{E}[\mathbb{E}(X | \mathcal{G})U] = \mathbb{E}(XU)$ for every $U \in \mathcal{L}_2(\mathcal{G})$.

But remember that $\mathbb{E}(X | \mathcal{G})$ is defined more generally for $X \in \mathcal{L}_1(\mathcal{F})$. Our final result in this discussion concerns convergence.

Suppose that $k \in [1, \infty)$ and that \mathcal{G} is a sub σ -algebra of \mathcal{F} .

1. If $X \in \mathcal{L}_k(\mathcal{F})$ then $\mathbb{E}(X | \mathcal{G}) \in \mathcal{L}_k(\mathcal{G})$
2. If $X_n \in \mathcal{L}_k(\mathcal{F})$ for $n \in \mathbb{N}_+$, $X \in \mathcal{L}_k(\mathcal{F})$, and $X_n \rightarrow X$ as $n \rightarrow \infty$ in $\mathcal{L}_k(\mathcal{F})$ then $\mathbb{E}(X_n | \mathcal{G}) \rightarrow \mathbb{E}(X | \mathcal{G})$ as $n \rightarrow \infty$ in $\mathcal{L}_k(\mathcal{G})$

Proof

1. Note that $|\mathbb{E}(X | \mathcal{G})| \leq \mathbb{E}(|X| | \mathcal{G})$. Since $t \mapsto t^k$ is increasing and convex on $[0, \infty)$ we have

$$|\mathbb{E}(X | \mathcal{G})|^k \leq [\mathbb{E}(|X| | \mathcal{G})]^k \leq \mathbb{E}(|X|^k | \mathcal{G}) \quad (4.11.35)$$

The last step uses Jensen's inequality. Taking expected values gives

$$\mathbb{E}[|\mathbb{E}(X | \mathcal{G})|^k] \leq \mathbb{E}(|X|^k) < \infty \quad (4.11.36)$$

2. Using the same ideas,

$$\mathbb{E} \left[|\mathbb{E}(X_n | \mathcal{G}) - \mathbb{E}(X | \mathcal{G})|^k \right] = \mathbb{E} \left[|\mathbb{E}(X_n - X | \mathcal{G})|^k \right] \leq \mathbb{E}[|X_n - X|^k] \quad (4.11.37)$$

By assumption, the right side converges to 0 as $n \rightarrow \infty$ and hence so does the left side.

Examples and Applications

App Exercises

In the error function app, select the root mean square error function. Click on the x -axis to generate an empirical distribution, and note the shape and location of the graph of the error function.

In the error function app, select the mean absolute error function. Click on the x -axis to generate an empirical distribution, and note the shape and location of the graph of the error function.

Computational Exercises

Suppose that X is uniformly distributed on the interval $[0, 1]$.

1. Find $\|X\|_k$ for $k \in [1, \infty)$.
2. Graph $\|X\|_k$ as a function of $k \in [1, \infty)$.
3. Find $\lim_{k \rightarrow \infty} \|X\|_k$.

Answer

1. $\frac{1}{(k+1)^{1/k}}$
3. 1

Suppose that X has probability density function $f(x) = \frac{a}{x^{a+1}}$ for $1 \leq x < \infty$, where $a > 0$ is a parameter. Thus, X has the Pareto distribution with shape parameter a .

1. Find $\|X\|_k$ for $k \in [1, \infty)$.
2. Graph $\|X\|_k$ as a function of $k \in (1, a)$.
3. Find $\lim_{k \uparrow a} \|X\|_k$.

Answer

1. $\left(\frac{a}{a-k}\right)^{1/k}$ if $k < a$, ∞ if $k \geq a$
3. ∞

Suppose that (X, Y) has probability density function $f(x, y) = x + y$ for $0 \leq x \leq 1$, $0 \leq y \leq 1$. Verify Minkowski's inequality.

Answer

1. $\|X + Y\|_k = \left(\frac{2^{k+2}-2}{(k+2)(k+3)}\right)^{1/k}$
2. $\|X\|_k + \|Y\|_k = 2\left(\frac{1}{k+2} + \frac{1}{2(k+1)}\right)^{1/k}$

Let X be an indicator random variable with $\mathbb{P}(X = 1) = p$, where $0 \leq p \leq 1$. Graph $\mathbb{E}(|X - t|)$ as a function of $t \in \mathbb{R}$ in each of the cases below. In each case, find the minimum value of the function and the values of t where the minimum occurs.

1. $p < \frac{1}{2}$
2. $p = \frac{1}{2}$
3. $p > \frac{1}{2}$

Answer

1. The minimum is p and occurs at $t = 0$.
2. The minimum is $\frac{1}{2}$ and occurs for $t \in [0, 1]$
3. The minimum is $1 - p$ and occurs at $t = 1$

Suppose that X is uniformly distributed on the interval $[0, 1]$. Find $d_1(X, t) = \mathbb{E}(|X - t|)$ as a function of t and sketch the graph. Find the minimum value of the function and the value of t where the minimum occurs.

Suppose that X is uniformly distributed on the set $[0, 1] \cup [2, 3]$. Find $d_1(X, t) = \mathbb{E}(|X - t|)$ as a function of t and sketch the graph. Find the minimum value of the function and the values of t where the minimum occurs.

Suppose that (X, Y) has probability density function $f(x, y) = x + y$ for $0 \leq x \leq 1$, $0 \leq y \leq 1$. Verify Hölder's inequality in the following cases:

1. $j = k = 2$
2. $j = 3, k = \frac{3}{2}$

Answer

1. $\|X\|_2 \|Y\|_2 = \frac{5}{12}$
2. $\|X\|_3 + \|Y\|_{3/2} \approx 0.4248$

Counterexamples

The following exercise shows that convergence with probability 1 does not imply convergence in mean.

Suppose that (X_1, X_2, \dots) is a sequence of independent random variables with

$$\mathbb{P}(X = n^3) = \frac{1}{n^2}, \quad \mathbb{P}(X_n = 0) = 1 - \frac{1}{n^2}; \quad n \in \mathbb{N}_+ \quad (4.11.38)$$

1. $X_n \rightarrow 0$ as $n \rightarrow \infty$ with probability 1.
2. $X_n \rightarrow 0$ as $n \rightarrow \infty$ in probability.
3. $\mathbb{E}(X_n) \rightarrow \infty$ as $n \rightarrow \infty$.

Proof

1. This follows from the basic characterization of convergence with probability 1: $\sum_{n=1}^{\infty} \mathbb{P}(X_n > \epsilon) = \sum_{n=1}^{\infty} 1/n^2 < \infty$ for $0 < \epsilon < 1$.
2. This follows since convergence with probability 1 implies convergence in probability.
3. Note that $\mathbb{E}(X_n) = n^3/n^2 = n$ for $n \in \mathbb{N}_+$.

The following exercise shows that convergence in mean does not imply convergence with probability 1.

Suppose that (X_1, X_2, \dots) is a sequence of independent indicator random variables with

$$\mathbb{P}(X_n = 1) = \frac{1}{n}, \quad \mathbb{P}(X_n = 0) = 1 - \frac{1}{n}; \quad n \in \mathbb{N}_+ \quad (4.11.39)$$

1. $\mathbb{P}(X_n = 0 \text{ for infinitely many } n) = 1$.
2. $\mathbb{P}(X_n = 1 \text{ for infinitely many } n) = 1$.
3. $\mathbb{P}(X_n \text{ does not converge as } n \rightarrow \infty) = 1$.
4. $X_n \rightarrow 0$ as $n \rightarrow \infty$ in k th mean for every $k \geq 1$.

Proof

1. This follows from the second Borel-Cantelli lemma since $\sum_{n=1}^{\infty} \mathbb{P}(X_n = 1) = \sum_{n=1}^{\infty} 1/n = \infty$.
2. This also follows from the second Borel-Cantelli lemma since $\sum_{n=1}^{\infty} \mathbb{P}(X_n = 0) = \sum_{n=1}^{\infty} (1 - 1/n) = \infty$.
3. This follows from parts (a) and (b).
4. Note that $\mathbb{E}(X_n) = 1/n \rightarrow 0$ as $n \rightarrow \infty$.

The following exercise show that convergence of the k th means does not imply convergence *in* k th mean.

Suppose that U has the Bernoulli distribution with parameter $\frac{1}{2}$, so that $\mathbb{P}(U = 1) = \mathbb{P}(U = 0) = \frac{1}{2}$. Let $X_n = U$ for $n \in \mathbb{N}_+$ and let $X = 1 - U$. Let $k \in [1, \infty)$. Then

1. $\mathbb{E}(X_n^k) = \mathbb{E}(X^k) = \frac{1}{2}$ for $n \in \mathbb{N}_+$, so $\mathbb{E}(X_n^k) \rightarrow \mathbb{E}(X^k)$ as $n \rightarrow \infty$
2. $\mathbb{E}(|X_n - X|^k) = 1$ for $n \in \mathbb{N}$ so X_n does not converge to X as $n \rightarrow \infty$ in \mathcal{L}_k .

Proof

1. Note that $X_n^k = U^k = U$ for $n \in \mathbb{N}_+$, since U just takes values 0 and 1. Also, U and $1 - U$ have the same distribution so $\mathbb{E}(U) = \mathbb{E}(1 - U) = \frac{1}{2}$.
2. Note that $X_n - X = U - (1 - U) = 2U - 1$ for $n \in \mathbb{N}_+$. Again, U just takes values 0 and 1, so $|2U - 1| = 1$.

This page titled [4.11: Vector Spaces of Random Variables](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.12: Uniformly Integrable Variables

Two of the most important modes of convergence in probability theory are convergence with probability 1 and convergence in mean. As we have noted several times, neither mode of convergence implies the other. However, if we impose an additional condition on the sequence of variables, convergence with probability 1 will imply convergence in mean. The purpose of this brief, but advanced section, is to explore the additional condition that is needed. This section is particularly important for the theory of martingales.

Basic Theory

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. So Ω is the set of outcomes, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is the probability measure on the sample space (Ω, \mathcal{F}) . In this section, all random variables that are mentioned are assumed to be real valued, unless otherwise noted. Next, recall from the section on vector spaces that for $k \in [1, \infty)$, \mathcal{L}_k is the vector space of random variables X with $\mathbb{E}(|X|^k) < \infty$, endowed with the norm $\|X\|_k = [\mathbb{E}(X^k)]^{1/k}$. In particular, $X \in \mathcal{L}_1$ simply means that $\mathbb{E}(|X|) < \infty$ so that $\mathbb{E}(X)$ exists as a real number. From the section on expected value as an integral, recall the following notation, assuming of course that the expected value makes sense:

$$\mathbb{E}(X; A) = \mathbb{E}(X\mathbf{1}_A) = \int_A X d\mathbb{P} \quad (4.12.1)$$

Definition

The following result is motivation for the main definition in this section.

If X is a random variable then $\mathbb{E}(|X|) < \infty$ if and only if $\mathbb{E}(|X|; |X| \geq x) \rightarrow 0$ as $x \rightarrow \infty$.

Proof

Note that that $|X|\mathbf{1}(|X| \leq x)$ is nonnegative, increasing in $x \in [0, \infty)$ and $|X|\mathbf{1}(|X| \leq x) \rightarrow |X|$ as $x \rightarrow \infty$. From the monotone convergence theorem, $\mathbb{E}(|X|; |X| \leq x) \rightarrow \mathbb{E}(|X|)$ as $x \rightarrow \infty$. On the other hand,

$$\mathbb{E}(|X|) = \mathbb{E}(|X|; |X| \leq x) + \mathbb{E}(|X|; |X| > x) \quad (4.12.2)$$

If $\mathbb{E}(|X|) < \infty$ then taking limits in the displayed equation shows that $\mathbb{E}(|X|; |X| > x) \rightarrow 0$ as $x \rightarrow \infty$. On the other hand, $\mathbb{E}(|X|; |X| \leq x) \leq x$. So if $\mathbb{E}(|X|) = \infty$ then $\mathbb{E}(|X|; |X| > x) = \infty$ for every $x \in [0, \infty)$.

Suppose now that X_i is a random variable for each i in a nonempty index set I (not necessarily countable). The critical definition for this section is to require the convergence in the previous theorem to hold *uniformly* for the collection of random variables $\mathbf{X} = \{X_i : i \in I\}$.

The collection $\mathbf{X} = \{X_i : i \in I\}$ is *uniformly integrable* if for each $\epsilon > 0$ there exists $x > 0$ such that for all $i \in I$,

$$\mathbb{E}(|X_i|; |X_i| > x) < \epsilon \quad (4.12.3)$$

Equivalently $\mathbb{E}(|X_i|; |X_i| > x) \rightarrow 0$ as $x \rightarrow \infty$ uniformly in $i \in I$.

Properties

Our next discussion centers on conditions that ensure that the collection of random variables $\mathbf{X} = \{X_i : i \in I\}$ is uniformly integrable. Here is an equivalent characterization:

The collection $\mathbf{X} = \{X_i : i \in I\}$ is uniformly integrable if and only if the following conditions hold:

1. $\{\mathbb{E}(|X_i|) : i \in I\}$ is bounded.
2. For each $\epsilon > 0$ there exists $\delta > 0$ such that if $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta$ then $\mathbb{E}(|X_i|; A) < \epsilon$ for all $i \in I$.

Proof

Suppose that \mathbf{X} is uniformly integrable. With $\epsilon = 1$ there exists $x > 0$ such that $\mathbb{E}(|X_i|; |X_i| > x) < 1$ for all $i \in I$. Hence

$$\mathbb{E}(|X_i|) = \mathbb{E}(|X_i|; |X_i| \leq x) + \mathbb{E}(|X_i|; |X_i| > x) \leq x + 1, \quad i \in I \quad (4.12.4)$$

so (a) holds. For (b), let $\epsilon > 0$. There exists $x > 0$ such that $\mathbb{E}(|X_i|; |X_i| > x) < \epsilon/2$ for all $i \in I$. Let $\delta = \epsilon/2x$. If $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta$ then

$$\mathbb{E}(|X_i|; A) = \mathbb{E}(|X_i|; A \cap \{|X| \leq x\}) + \mathbb{E}(|X_i|; A \cap \{|X| > x\}) \leq x\mathbb{P}(A) + \mathbb{E}(|X_i|; |X| > x) < \epsilon/2 + \epsilon/2 = \epsilon \quad (4.12.5)$$

Conversely, suppose that (a) and (b) hold. By (a), there exists $c > 0$ such that $\mathbb{E}(|X_i|) \leq c$ for all $i \in I$. Let $\epsilon > 0$. By (b) there exists $\delta > 0$ such that if $A \in \mathcal{F}$ with $\mathbb{P}(A) < \delta$ then $\mathbb{E}(|X_i|; A) < \epsilon$ for all $i \in I$. Next, by Markov's inequality,

$$\mathbb{P}(|X_i| > x) \leq \frac{\mathbb{E}(|X_i|)}{x} \leq \frac{c}{x}, \quad i \in I \quad (4.12.6)$$

Pick $x > 0$ such that $c/x < \delta$, so that $\mathbb{P}(|X_i| > x) < \delta$ for each $i \in I$. Then for each $j \in I$, $\mathbb{E}(|X_i|; |X_j| > x) < \epsilon$ for all $i \in I$ and so in particular, $\mathbb{E}(|X_i|; |X_i| > x) < \epsilon$ for all $i \in I$. Hence \mathbf{X} is uniformly integrable.

Condition (a) means that \mathbf{X} is bounded (in norm) as a subset of the vector space \mathcal{L}_1 . Trivially, a *finite* collection of integrable random variables is uniformly integrable.

Suppose that I is finite and that $\mathbb{E}(|X_i|) < \infty$ for each $i \in I$. Then $\mathbf{X} = \{X_i : i \in I\}$ is uniformly integrable.

A subset of a uniformly integrable set of variables is also uniformly integrable.

If $\{X_i : i \in I\}$ is uniformly integrable and J is a nonempty subset of I , then $\{X_j : j \in J\}$ is uniformly integrable.

If the random variables in the collection are dominated in absolute value by a random variable with finite mean, then the collection is uniformly integrable.

Suppose that Y is a nonnegative random variable with $\mathbb{E}(Y) < \infty$ and that $|X_i| \leq Y$ for each $i \in I$. Then $\mathbf{X} = \{X_i : i \in I\}$ is uniformly integrable.

Proof

Clearly $\mathbb{E}(|X_i|; |X_i| > x) \leq \mathbb{E}(Y; Y > x)$ for $x \in [0, \infty)$ and for all $i \in I$. The right side is independent of $i \in I$, and by the [theorem above](#), converges to 0 as $x \rightarrow \infty$. Hence \mathbf{X} is uniformly integrable.

The following result is more general, but essentially the same proof works.

Suppose that $\mathbf{Y} = \{Y_j : j \in J\}$ is uniformly integrable, and $\mathbf{X} = \{X_i : i \in I\}$ is a set of variables with the property that for each $i \in I$ there exists $j \in J$ such that $|X_i| \leq |Y_j|$. Then \mathbf{X} is uniformly integrable.

As a simple corollary, if the variables are bounded in absolute value then the collection is uniformly integrable.

If there exists $c > 0$ such that $|X_i| \leq c$ for all $i \in I$ then $\mathbf{X} = \{X_i : i \in I\}$ is uniformly integrable.

Just having $\mathbb{E}(|X_i|)$ bounded in $i \in I$ (condition (a) in the [characterization above](#)) is not sufficient for $\mathbf{X} = \{X_i : i \in I\}$ to be uniformly integrable; a [counterexample](#) is given below. However, if $\mathbb{E}(|X_i|^k)$ is bounded in $i \in I$ for some $k > 1$, then \mathbf{X} is uniformly integrable. This condition means that \mathbf{X} is bounded (in norm) as a subset of the vector space \mathcal{L}_k .

If $\{\mathbb{E}(|X_i|^k) : i \in I\}$ is bounded for some $k > 1$, then $\{X_i : i \in I\}$ is uniformly integrable.

Proof

Suppose that for some $k > 1$ and $c > 0$, $\mathbb{E}(|X_i|^k) \leq c$ for all $i \in I$. Then $k - 1 > 0$ and so $t \mapsto t^{k-1}$ is increasing on $(0, \infty)$. So if $|X_i| > x$ for $x > 0$ then

$$|X_i|^k = |X_i| |X_i|^{k-1} \geq |X_i| x^{k-1} \quad (4.12.7)$$

Hence $|X_i| \leq |X_i|^k / x^{k-1}$ on the event $|X_i| > x$. Therefore

$$\mathbb{E}(|X_i|; |X_i| > x) \leq \mathbb{E}\left(\frac{|X_i|^k}{x^{k-1}}; |X_i| > x\right) \leq \frac{\mathbb{E}(|X_i|^k)}{x^{k-1}} \leq \frac{c}{x^{k-1}} \quad (4.12.8)$$

The last expression is independent of $i \in I$ and converges to 0 as $x \rightarrow \infty$. Hence \mathbf{X} is uniformly integrable.

Uniformly integrability is closed under the operations of addition and scalar multiplication.

Suppose that $\mathbf{X} = \{X_i : i \in I\}$ and $\mathbf{Y} = \{Y_i : i \in I\}$ are uniformly integrable and that $c \in \mathbb{R}$. Then each of the following collections is also uniformly integrable.

1. $\mathbf{X} + \mathbf{Y} = \{X_i + Y_i : i \in I\}$
2. $c\mathbf{X} = \{cX_i : i \in I\}$

Proof

We use the [characterization above](#). The proofs use standard techniques, so try them yourself.

1. There exists $a, b \in (0, \infty)$ such that $\mathbb{E}(|X_i|) \leq a$ and $\mathbb{E}(|Y_i|) \leq b$ for all $i \in I$. Hence

$$\mathbb{E}(|X_i + Y_i|) \leq \mathbb{E}(|X_i| + |Y_i|) \leq \mathbb{E}(|X_i|) + \mathbb{E}(|Y_i|) \leq a + b, \quad i \in I \quad (4.12.9)$$

Next let $\epsilon > 0$. There exists $\delta_1 > 0$ such that if $A \in \mathcal{F}$ with $\mathbb{P}(A) < \delta_1$ then $\mathbb{E}(|X_i|; A) < \epsilon/2$ for all $i \in I$, and similarly, there exists $\delta_2 > 0$ such that if $A \in \mathcal{F}$ with $\mathbb{P}(A) < \delta_2$ then $\mathbb{E}(|Y_i|; A) < \epsilon/2$ for all $i \in I$. Hence if $A \in \mathcal{F}$ with $\mathbb{P}(A) < \delta_1 \wedge \delta_2$ then

$$\mathbb{E}(|X_i + Y_i|; A) \leq \mathbb{E}(|X_i| + |Y_i|; A) = \mathbb{E}(|X_i|; A) + \mathbb{E}(|Y_i|; A) < \epsilon/2 + \epsilon/2 = \epsilon, \quad i \in I \quad (4.12.10)$$

2. There exists $a \in (0, \infty)$ such that $\mathbb{E}(|X_i|) \leq a$ for all $i \in I$. Hence

$$\mathbb{E}(|cX_i|) = |c|\mathbb{E}(|X_i|) \leq ca, \quad i \in I \quad (4.12.11)$$

The second condition is trivial if $c = 0$, so suppose $c \neq 0$. For $\epsilon > 0$ there exists $\delta > 0$ such that if $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta$ then $\mathbb{E}(|X_i|; A) < \epsilon/c$ for all $i \in I$. Hence $\mathbb{E}(|cX_i|; A) = |c|\mathbb{E}(|X_i|; A) < \epsilon$.

The following corollary is trivial, but will be needed in our discussion of convergence below.

Suppose that $\{X_i : i \in I\}$ is uniformly integrable and that X is a random variable with $\mathbb{E}(|X|) < \infty$. Then $\{X_i - X : i \in I\}$ is uniformly integrable.

Proof

Let $Y_i = X$ for each $i \in I$. Then $\{Y_i : i \in I\}$ is uniformly integrable, so the result follows from the previous theorem.

Convergence

We now come to the main results, and the reason for the definition of uniform integrability in the first place. To set up the notation, suppose that X_n is a random variable for $n \in \mathbb{N}_+$ and that X is a random variable. We know that if $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean then $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability. The converse is also true if and only if the sequence is uniformly integrable. Here is the first half:

If $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean, then $\{X_n : n \in \mathbb{N}\}$ is uniformly integrable.

Proof

The hypothesis means that $X_n \rightarrow X$ as $n \rightarrow \infty$ in the vector space \mathcal{L}_1 . That is, $\mathbb{E}(|X_n|) < \infty$ for $n \in \mathbb{N}_+$, $\mathbb{E}(|X|) < \infty$, and $\mathbb{E}(|X_n - X|) \rightarrow 0$ as $n \rightarrow \infty$. From the last section, we know that this implies that $\mathbb{E}(|X_n|) \rightarrow \mathbb{E}(|X|)$ as $n \rightarrow \infty$, so $\mathbb{E}(|X_n|)$ is bounded in $n \in \mathbb{N}$. Let $\epsilon > 0$. Then there exists $N \in \mathbb{N}_+$ such that if $n > N$ then $\mathbb{E}(|X_n - X|) < \epsilon/2$. Since all of our variables are in \mathcal{L}_1 , for each $n \in \mathbb{N}_+$ there exists $\delta_n > 0$ such that if $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta_n$ then $\mathbb{E}(|X_n - X|; A) < \epsilon/2$. Similarly, there exists $\delta_0 > 0$ such that if $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta_0$ then $\mathbb{E}(|X|; A) < \epsilon/2$. Let $\delta = \min\{\delta_n : n \in \{0, 1, \dots, N\}\}$ so $\delta > 0$. If $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta$ then

$$\mathbb{E}(|X_n|; A) = \mathbb{E}(|X_n - X + X|; A) \leq \mathbb{E}(|X_n - X|; A) + \mathbb{E}(|X|; A), \quad n \in \mathbb{N}_+ \quad (4.12.12)$$

If $n \leq N$ then $\mathbb{E}(|X_n - X|; A) \leq \epsilon/2$ since $\delta \leq \delta_n$. If $n > N$ then $\mathbb{E}(|X_n - X|; A) \leq \mathbb{E}(|X_n - X|) < \epsilon/2$. For all n , $\mathbb{E}(|X|; A) < \epsilon/2$ since $\delta \leq \delta_0$. So for all $n \in \mathbb{N}_+$, $\mathbb{E}(|X_n|; A) < \epsilon$ and hence $\{X_n : n \in \mathbb{N}_+\}$ is uniformly integrable.

Here is the more important half, known as the *uniform integrability theorem*:

If $\{X_n : n \in \mathbb{N}_+\}$ is uniformly integrable and $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability, then $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean.

Proof

Since $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability, we know that there exists a subsequence $(X_{n_k} : k \in \mathbb{N}_+)$ of $(X_n : n \in \mathbb{N}_+)$ such that $X_{n_k} \rightarrow X$ as $k \rightarrow \infty$ with probability 1. By the uniform integrability, $\mathbb{E}(|X_n|)$ is bounded in $n \in \mathbb{N}_+$. Hence by Fatou's lemma

$$\mathbb{E}(|X|) = \mathbb{E}\left(\liminf_{k \rightarrow \infty} |X_{n_k}|\right) \leq \liminf_{n \rightarrow \infty} \mathbb{E}(|X_n|) \leq \limsup_{n \rightarrow \infty} \mathbb{E}(|X_n|) < \infty \quad (4.12.13)$$

Let $Y_n = X_n - X$ for $n \in \mathbb{N}_+$. From the [corollary above](#), we know that $\{Y_n : n \in \mathbb{N}_+\}$ is uniformly integrable, and we also know that Y_n converges to 0 as $n \rightarrow \infty$ in probability. Hence we need to show that $Y_n \rightarrow 0$ as $n \rightarrow \infty$ in mean. Let $\epsilon > 0$. By uniform

integrability, there exists $\delta > 0$ such that if $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta$ then $\mathbb{E}(|Y_n| : A) < \epsilon/2$ for all $n \in \mathbb{N}$. Since $Y_n \rightarrow 0$ as $n \rightarrow \infty$ in probability, there exists $N \in \mathbb{N}_+$ such that if $n > N$ then $\mathbb{P}(|Y_n| > \epsilon/2) < \delta$. Hence if $n > N$ then

$$\mathbb{E}(|Y_n|) = \mathbb{E}(|Y_n|; |Y_n| \leq \epsilon/2) + \mathbb{E}(|Y_n|; |Y_n| > \epsilon/2) < \epsilon/2 + \epsilon/2 = \epsilon \quad (4.12.14)$$

Hence $Y_n \rightarrow 0$ as $n \rightarrow \infty$ in mean.

As a corollary, recall that if $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1, then $X_n \rightarrow X$ as $n \rightarrow \infty$ in probability. Hence if $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is uniformly integrable then $X_n \rightarrow X$ as $n \rightarrow \infty$ in mean.

Examples

Our first example shows that bounded \mathcal{L}_1 norm is not sufficient for uniform integrability.

Suppose that U is uniformly distributed on the interval $(0, 1)$ (so U has the *standard uniform distribution*). For $n \in \mathbb{N}_+$, let $X_n = n\mathbf{1}(U \leq 1/n)$. Then

1. $\mathbb{E}(|X_n|) = 1$ for all $n \in \mathbb{N}_+$
2. $\mathbb{E}(|X_n|; |X_n| > x) = 1$ for $x > 0$, $n \in \mathbb{N}_+$ with $n > x$

Proof

First note that $|X_n| = X_n$ since $X_n \geq 0$.

1. By definition, $\mathbb{E}(X_n) = n\mathbb{P}(U \leq 1/n) = n/n = 1$ for $n \in \mathbb{N}_+$.
2. If $n > x > 0$ then $X_n > x$ if and only if $X_n = n$ if and only if $U \leq 1/n$. Hence $\mathbb{E}(X_n; X_n > x) = n\mathbb{P}(U \leq 1/n) = 1$ as before.

By part (b), $\mathbb{E}(|X_n|; |X_n| > x)$ does not converge to 0 as $x \rightarrow \infty$ uniformly in $n \in \mathbb{N}_+$, so $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is not uniformly integrable.

The next example gives an important application to conditional expected value. Recall that if X is a random variable with $\mathbb{E}(|X|) < \infty$ and \mathcal{G} is a sub σ -algebra of \mathcal{F} then $\mathbb{E}(X | \mathcal{G})$ is the expected value of X given the information in \mathcal{G} , and is the \mathcal{G} -measurable random variable closest to X in a sense. Indeed if $X \in \mathcal{L}_2(\mathcal{F})$ then $\mathbb{E}(X | \mathcal{G})$ is the projection of X onto $\mathcal{L}_2(\mathcal{G})$. The collection of all conditional expected values of X is uniformly integrable:

Suppose that X is a real-valued random variable with $\mathbb{E}(|X|) < \infty$. Then $\{\mathbb{E}(X | \mathcal{G}) : \mathcal{G} \text{ is a sub } \sigma\text{-algebra of } \mathcal{F}\}$ is uniformly integrable.

Proof

We use the [characterization above](#). Let \mathcal{G} be a sub σ -algebra of \mathcal{F} . Recall that $|\mathbb{E}(X | \mathcal{G})| \leq \mathbb{E}(|X| | \mathcal{G})$ and hence

$$\mathbb{E}[|\mathbb{E}(X | \mathcal{G})|] \leq \mathbb{E}[\mathbb{E}(|X| | \mathcal{G})] = \mathbb{E}(|X|) \quad (4.12.15)$$

So property (a) holds. Next let $\epsilon > 0$. Since $\mathbb{E}(|X|) < \infty$, there exists $\delta > 0$ such that if $A \in \mathcal{F}$ and $\mathbb{P}(A) < \delta$ then $\mathbb{E}(|X|; A) < \epsilon$. Suppose that $A \in \mathcal{G}$ with $\mathbb{P}(A) < \delta$. Then $|\mathbb{E}(X | \mathcal{G})|\mathbf{1}_A \leq \mathbb{E}(|X| | \mathcal{G})\mathbf{1}_A$ so

$$\mathbb{E}[|\mathbb{E}(X | \mathcal{G})|; A] \leq \mathbb{E}[\mathbb{E}(|X| | \mathcal{G}); A] = \mathbb{E}[\mathbb{E}(|X|\mathbf{1}_A | \mathcal{G})] = \mathbb{E}(|X|; A) < \epsilon \quad (4.12.16)$$

So condition (b) holds. Note that the first equality in the displayed equation holds since $A \in \mathcal{G}$.

Note that the collection of sub σ -algebras of \mathcal{F} , and so also the collection of conditional expected values above, might well be uncountable. The conditional expected values range from $\mathbb{E}(X)$, when $\mathcal{G} = \{\Omega, \emptyset\}$ to X itself, when $\mathcal{G} = \mathcal{F}$.

This page titled [4.12: Uniformly Integrable Variables](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

4.13: Kernels and Operators

The goal of this section is to study a type of mathematical object that arises naturally in the context of conditional expected value and parametric distributions, and is of fundamental importance in the study of stochastic processes, particularly Markov processes. In a sense, the main object of study in this section is a generalization of a matrix, and the operations generalizations of matrix operations. If you keep this in mind, this section may seem less abstract.

Basic Theory

Definitions

Recall that a measurable space (S, \mathcal{S}) consists of a set S and a σ -algebra \mathcal{S} of subsets of S . If μ is a positive measure on (S, \mathcal{S}) , then (S, \mathcal{S}, μ) is a measure space. The two most important special cases that we have studied frequently are

1. *Discrete*: S is countable, $\mathcal{S} = \mathcal{P}(S)$ is the collection of all subsets of S , and $\mu = \#$ is counting measure on (S, \mathcal{S}) .
2. *Euclidean*: S is a measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$, \mathcal{S} is the collection of subsets of S that are also measurable, and $\mu = \lambda_n$ is n -dimensional Lebesgue measure on (S, \mathcal{S}) .

More generally, S usually comes with a topology that is locally compact, Hausdorff, with a countable base (LCCB), and \mathcal{S} is the *Borel σ -algebra*, the σ -algebra generated by the topology (the collection of open subsets of S). The measure μ is usually a *Borel measure*, and so satisfies $\mu(C) < \infty$ if $C \subseteq S$ is compact. A discrete measure space is of this type, corresponding to the discrete topology. A Euclidean measure space is also of this type, corresponding to the Euclidean topology, if S is open or closed (which is usually the case). In the discrete case, every function from S to another measurable space is measurable, and every from function from S to another topological space is continuous, so the measure theory is not really necessary.

Recall also that the measure space (S, \mathcal{S}, μ) is σ -finite if there exists a countable collection $\{A_i : i \in I\} \subseteq \mathcal{S}$ such that $\mu(A_i) < \infty$ for $i \in I$ and $S = \bigcup_{i \in I} A_i$. If (S, \mathcal{S}, μ) is a Borel measure space corresponding to an LCCB topology, then it is σ -finite.

If $f : S \rightarrow \mathbb{R}$ is measurable, define $\|f\| = \sup\{|f(x)| : x \in S\}$. Of course we may well have $\|f\| = \infty$. Let $\mathcal{B}(S)$ denote the collection of bounded measurable functions $f : S \rightarrow \mathbb{R}$. Under the usual operations of pointwise addition and scalar multiplication, $\mathcal{B}(S)$ is a vector space, and $\|\cdot\|$ is the natural norm on this space, known as the *supremum norm*. This vector space plays an important role.

In this section, it is sometimes more natural to write integrals with respect to the positive measure μ with the differential *before* the integrand, rather than after. However, rest assured that this is mere notation, the meaning of the integral is the same. So if $f : S \rightarrow \mathbb{R}$ is measurable then we may write the integral of f with respect to μ in operator notation as

$$\mu f = \int_S \mu(dx) f(x) \quad (4.13.1)$$

assuming, as usual, that the integral exists. This will be the case if f is nonnegative, although ∞ is a possible value. More generally, the integral exists in $\mathbb{R} \cup \{-\infty, \infty\}$ if $\mu f^+ < \infty$ or $\mu f^- < \infty$ where f^+ and f^- are the positive and negative parts of f . If both are finite, the integral exists in \mathbb{R} (and f is *integrable* with respect to μ). If μ is a probability measure and we think of (S, \mathcal{S}) as the sample space of a random experiment, then we can think of f as a real-valued random variable, in which case our new notation is not too far from our traditional expected value $\mathbb{E}(f)$. Our main definition comes next.

Suppose that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. A *kernel* from (S, \mathcal{S}) to (T, \mathcal{T}) is a function $K : S \times \mathcal{T} \rightarrow [0, \infty]$ such that

1. $x \mapsto K(x, A)$ is a measurable function from S into $[0, \infty]$ for each $A \in \mathcal{T}$.
2. $A \mapsto K(x, A)$ is a positive measure on \mathcal{T} for each $x \in S$.

If $(T, \mathcal{T}) = (S, \mathcal{S})$, then K is said to be a *kernel on* (S, \mathcal{S}) .

There are several classes of kernels that deserve special names.

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) . Then

1. K is σ -finite if the measure $K(x, \cdot)$ is σ -finite for every $x \in S$.
2. K is finite if $K(x, T) < \infty$ for every $x \in S$.
3. K is bounded if $K(x, T)$ is bounded in $x \in S$.
4. K is a probability kernel if $K(x, T) = 1$ for every $x \in S$.

Define $\|K\| = \sup\{K(x, T) : x \in S\}$, so that $\|K\| < \infty$ if K is a bounded kernel and $\|K\| = 1$ if K is a probability kernel.

So a probability kernel is bounded, a bounded kernel is finite, and a finite kernel is σ -finite. The terms *stochastic kernel* and *Markov kernel* are also used for probability kernels, and for a probability kernel $\|K\| = 1$ of course. The terms are consistent with terms used for measures: K is a finite kernel if and only if $K(x, \cdot)$ is a finite measure for each $x \in S$, and K is a probability kernel if and only if $K(x, \cdot)$ is a probability measure for each $x \in S$. Note that $\|K\|$ is simply the supremum norm of the function $x \mapsto K(x, T)$.

A kernel defines two natural integral operators, by operating on the *left* with measures, and by operating on the *right* with functions. As usual, we are often a bit casual with the question of existence. Basically in this section, we assume that any integrals mentioned exist.

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) .

1. If μ is a positive measure on (S, \mathcal{S}) , then μK defined as follows is a positive measure on (T, \mathcal{T}) :

$$\mu K(A) = \int_S \mu(dx) K(x, A), \quad A \in \mathcal{T} \quad (4.13.2)$$

2. If $f : T \rightarrow \mathbb{R}$ is measurable, then $Kf : S \rightarrow \mathbb{R}$ defined as follows is measurable (assuming that the integrals exist in \mathbb{R}):

$$Kf(x) = \int_T K(x, dy) f(y), \quad x \in S \quad (4.13.3)$$

Proof

1. Clearly $\mu K(A) \geq 0$ for $A \in \mathcal{T}$. Suppose that $\{A_j : j \in J\}$ is a countable collection of disjoint sets in \mathcal{T} and $A = \bigcup_{j \in J} A_j$. Then

$$\begin{aligned} \mu K(A) &= \int_S \mu(dx) K(x, A) = \int_S \mu(dx) \left(\sum_{j \in J} K(x, A_j) \right) \\ &= \sum_{j \in J} \int_S \mu(dx) K(x, A_j) = \sum_{j \in J} \mu K(A_j) \end{aligned}$$

The interchange of sum and integral is justified since the terms are nonnegative.

2. The measurability of Kf follows from the measurability of f and of $x \mapsto K(x, A)$ for $A \in \mathcal{T}$, and from basic properties of the integral.

Thus, a kernel transforms measures on (S, \mathcal{S}) into measures on (T, \mathcal{T}) , and transforms certain measurable functions from T to \mathbb{R} into measurable functions from S to \mathbb{R} . Again, part (b) assumes that f is integrable with respect to the measure $K(x, \cdot)$ for every $x \in S$. In particular, the last statement will hold in the following important special case:

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) and that $f \in \mathcal{B}(T)$.

1. If K is finite then Kf is defined and $\|Kf\| = \|K\| \|f\|$.
2. If K is bounded then $Kf \in \mathcal{B}(T)$.

Proof

1. If K is finite then

$$K|f|(x) = \int_T K(x, dy) |f(y)| \leq \int_T K(x, dy) \|f\| = \|f\| K(x, T) < \infty \quad x \in S \quad (4.13.4)$$

Hence f is integrable with respect to $K(x, \cdot)$ for each $x \in S$ so Kf is defined. Continuing with our inequalities, we have $|Kf(x)| \leq K|f|(x) \leq \|f\|K(x, T) \leq \|f\|\|K\|$ so $\|Kf\| \leq \|K\|\|f\|$. Moreover equality holds when $f = \mathbf{1}_T$, the constant function 1 on T .

2. If K is bounded then $\|K\| < \infty$ so from (a), $\|Kf\| < \infty$.

The *identity kernel* I on the measurable space (S, \mathcal{S}) is defined by $I(x, A) = \mathbf{1}(x \in A)$ for $x \in S$ and $A \in \mathcal{S}$.

Thus, $I(x, A) = 1$ if $x \in A$ and $I(x, A) = 0$ if $x \notin A$. So $x \mapsto I(x, A)$ is the *indicator function* of $A \in \mathcal{S}$, while $A \mapsto I(x, A)$ is *point mass* at $x \in S$. Clearly the identity kernel is a probability kernel. If we need to indicate the dependence on the particular space, we will add a subscript. The following result justifies the name.

Let I denote the identity kernel on (S, \mathcal{S}) .

1. If μ is a positive measure on (S, \mathcal{S}) then $\mu I = \mu$.
2. If $f : S \rightarrow \mathbb{R}$ is measurable, then $If = f$.

Constructions

We can create a new kernel from two given kernels, by the usual operations of addition and scalar multiplication.

Suppose that K and L are kernels from (S, \mathcal{S}) to (T, \mathcal{T}) , and that $c \in [0, \infty)$. Then cK and $K + L$ defined below are also kernels from (S, \mathcal{S}) to (T, \mathcal{T}) .

1. $(cK)(x, A) = cK(x, A)$ for $x \in S$ and $A \in \mathcal{T}$.
2. $(K + L)(x, A) = K(x, A) + L(x, A)$ for $x \in S$ and $A \in \mathcal{T}$.

If K and L are σ -finite (finite) (bounded) then cK and $K + L$ are σ -finite (finite) (bounded), respectively.

Proof

These results are simple.

1. Since $x \mapsto K(x, A)$ is measurable for $A \in \mathcal{T}$, so is $x \mapsto cK(x, A)$. Since $A \mapsto K(x, A)$ is a positive measure on (T, \mathcal{T}) for $x \in S$, so is $A \mapsto cK(x, A)$ since $c \geq 0$.
2. Since $x \mapsto K(x, A)$ and $x \mapsto L(x, A)$ are measurable for $A \in \mathcal{T}$, so is $x \mapsto K(x, A) + L(x, A)$. Since $A \mapsto K(x, A)$ and $A \mapsto L(x, A)$ are positive measures on (T, \mathcal{T}) for $x \in S$, so is $A \mapsto K(x, A) + L(x, A)$.

A simple corollary of the last result is that if $a, b \in [0, \infty)$ then $aK + bL$ is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) . In particular, if K, L are probability kernels and $p \in (0, 1)$ then $pK + (1 - p)L$ is a probability kernel. A more interesting and important way to form a new kernel from two given kernels is via a “multiplication” operation.

Suppose that K is a kernel from (R, \mathcal{R}) to (S, \mathcal{S}) and that L is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) . Then KL defined as follows is a kernel from (R, \mathcal{R}) to (T, \mathcal{T}) :

$$KL(x, A) = \int_S K(x, dy)L(y, A), \quad x \in R, A \in \mathcal{T} \quad (4.13.5)$$

1. If K is finite and L is bounded then KL is finite.
2. If K and L are bounded then KL is bounded.
3. If K and L are stochastic then KL is stochastic

Proof

The measurability of $x \mapsto (KL)(x, A)$ for $A \in \mathcal{T}$ follows from basic properties of the integral. For the second property, fix $x \in R$. Clearly $KL(x, A) \geq 0$ for $A \in \mathcal{T}$. Suppose that $\{A_j : j \in J\}$ is a countable collection of disjoint sets in \mathcal{T} and $A = \bigcup_{j \in J} A_j$. Then

$$\begin{aligned}
 KL(x, A) &= \int_S K(x, dy) L(x, A) = \int_S K(x, dy) \left(\sum_{j \in J} L(y, A_j) \right) \\
 &= \sum_{j \in J} \int_S K(x, dy) L(y, A_j) = \sum_{j \in J} KL(x, A_j)
 \end{aligned}$$

The interchange of sum and integral is justified since the terms are nonnegative.

Once again, the identity kernel lives up to its name:

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) . Then

1. $I_S K = K$
2. $K I_T = K$

The next several results show that the operations are *associative* whenever they make sense.

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , μ is a positive measure on \mathcal{S} , $c \in [0, \infty)$, and $f : T \rightarrow \mathbb{R}$ is measurable. Then, assuming that the appropriate integrals exist,

1. $c(\mu K) = (c\mu)K$
2. $c(Kf) = (cK)f$
3. $(\mu K)f = \mu(Kf)$

Proof

These results follow easily from the definitions.

1. The common measure on \mathcal{T} is $c\mu K(A) = c \int_S \mu(dx) K(x, A)$ for $A \in \mathcal{T}$.
2. The common function from S to \mathbb{R} is $cKf(x) = c \int_S K(x, dy) f(y)$ for $x \in S$, assuming that the integral exists for $x \in S$.
3. The common real number is $\mu Kf = \int_S \mu(dx) \int_T K(x, dy) f(y)$, assuming that the integrals exist.

Suppose that K is a kernel from (R, \mathcal{R}) to (S, \mathcal{S}) and L is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) . Suppose also that μ is a positive measure on (R, \mathcal{R}) , $f : T \rightarrow \mathbb{R}$ is measurable, and $c \in [0, \infty)$. Then, assuming that the appropriate integrals exist,

1. $(\mu K)L = \mu(KL)$
2. $K(Lf) = (KL)f$
3. $c(KL) = (cK)L$

Proof

These results follow easily from the definitions.

1. The common measure on (T, \mathcal{T}) is $\mu KL(A) = \int_R \mu(dx) \int_S K(x, dy) L(y, A)$ for $A \in \mathcal{T}$.
2. The common measurable function from R to \mathbb{R} is $KLf(x) = \int_S K(x, dy) \int_T L(y, dz) f(z)$ for $x \in R$, assuming that the integral exists for $x \in S$.
3. The common kernel from (R, \mathcal{R}) to (T, \mathcal{T}) is $cKL(x, A) = c \int_S K(x, dy) L(y, A)$ for $x \in R$ and $A \in \mathcal{T}$.

Suppose that K is a kernel from (R, \mathcal{R}) to (S, \mathcal{S}) , L is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , and M is a kernel from (T, \mathcal{T}) to (U, \mathcal{U}) . Then $(KL)M = K(LM)$.

Proof

This results follow easily from the definitions. The common kernel from (R, \mathcal{R}) to (U, \mathcal{U}) is

$$KLM(x, A) = \int_S K(x, dy) \int_T L(y, dz) M(z, A), \quad x \in R, A \in \mathcal{U} \quad (4.13.6)$$

The next several results show that the *distributive property* holds whenever the operations makes sense.

Suppose that K and L are kernels from (R, \mathcal{R}) to (S, \mathcal{S}) and that M and N are kernels from (S, \mathcal{S}) to (T, \mathcal{T}) . Suppose also that μ is a positive measure on (R, \mathcal{R}) and that $f : S \rightarrow \mathbb{R}$ is measurable. Then, assuming that the appropriate integrals exist,

1. $(K + L)M = KM + LM$
2. $K(M + N) = KM + KN$
3. $\mu(K + L) = \mu K + \mu L$
4. $(K + L)f = Kf + Lf$

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , and that μ and ν are positive measures on (S, \mathcal{S}) , and that f and g are measurable functions from T to \mathbb{R} . Then, assuming that the appropriate integrals exist,

1. $(\mu + \nu)K = \mu K + \nu K$
2. $K(f + g) = Kf + Kg$
3. $\mu(f + g) = \mu f + \mu g$
4. $(\mu + \nu)f = \mu f + \nu f$

In particular, note that if K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , then the transformation $\mu \mapsto \mu K$ defined for positive measures on (S, \mathcal{S}) , and the transformation $f \mapsto Kf$ defined for measurable functions $f : T \rightarrow \mathbb{R}$ (for which Kf exists), are both *linear* operators. If μ is a positive measure on (S, \mathcal{S}) , then the integral operator $f \mapsto \mu f$ defined for measurable $f : S \rightarrow \mathbb{R}$ (for which μf exists) is also linear, but of course, we already knew that. Finally, note that the operator $f \mapsto Kf$ is *positive*: if $f \geq 0$ then $Kf \geq 0$. Here is the important summary of our results when the kernel is bounded.

If K is a bounded kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , then $f \mapsto Kf$ is a bounded, linear transformation from $\mathcal{B}(T)$ to $\mathcal{B}(S)$ and $\|K\|$ is the norm of the transformation.

The *commutative property* for the product of kernels fails with a passion. If K and L are kernels, then depending on the measurable spaces, KL may be well defined, but not LK . Even if both products are defined, they may be kernels from or to different measurable spaces. Even if both are defined from and to the same measurable spaces, it may well happen that $KL \neq LK$. Some [examples](#) are given below

If K is a kernel on (S, \mathcal{S}) and $n \in \mathbb{N}$, we let $K^n = KK \cdots K$, the n -fold power of K . By convention, $K^0 = I$, the identity kernel on S .

Fixed points of the operators associated with a kernel turn out to be very important.

Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) .

1. A positive measure μ on (S, \mathcal{S}) such that $\mu K = \mu$ is said to be *invariant* for K .
2. A measurable function $f : T \rightarrow \mathbb{R}$ such that $Kf = f$ is said to be *invariant* for K

So in the language of linear algebra (or functional analysis), an invariant measure is a *left eigenvector* of the kernel, while an invariant function is a *right eigenvector* of the kernel, both corresponding to the eigenvalue 1. By our results above, if μ and ν are invariant measures and $c \in [0, \infty)$, then $\mu + \nu$ and $c\mu$ are also invariant. Similarly, if f and g are invariant functions and $c \in \mathbb{R}$, then $f + g$ and cf are also invariant.

Of course we are particularly interested in *probability* kernels.

Suppose that P is a probability kernel from (R, \mathcal{R}) to (S, \mathcal{S}) and that Q is a probability kernel from (S, \mathcal{S}) to (T, \mathcal{T}) . Suppose also that μ is a probability measure on (R, \mathcal{R}) . Then

1. PQ is a probability kernel from (R, \mathcal{R}) to (T, \mathcal{T}) .
2. μP is a probability measure on (S, \mathcal{S}) .

Proof

1. We know that PQ is a kernel from (R, \mathcal{R}) to (T, \mathcal{T}) . So we just need to note that

$$PQ(T) = \int_S P(x, dy)Q(y, T) = \int_S P(x, dy) = P(x, S) = 1, \quad x \in R \quad (4.13.7)$$

2. We know that μP is a positive measure on (S, \mathcal{S}) . So we just need to note that

$$\mu P(S) = \int_R \mu(dx) P(x, S) = \int_R \mu(dx) = \mu(R) = 1 \quad (4.13.8)$$

As a corollary, it follows that if P is a probability kernel on (S, \mathcal{S}) , then so is P^n for $n \in \mathbb{N}$.

The operators associated with a kernel are of fundamental importance, and we can easily recover the kernel from the operators. Suppose that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , and let $x \in S$ and $A \in \mathcal{T}$. Then trivially, $K \mathbf{1}_A(x) = K(x, A)$ where as usual, $\mathbf{1}_A$ is the indicator function of A . Trivially also $\delta_x K(A) = K(x, A)$ where δ_x is point mass at x .

Kernel Functions

Usually our *measurable* spaces are in fact *measure* spaces, with natural measures associated with the spaces, as in the special cases described in (1). When we start with measure spaces, kernels are usually constructed from *density functions* in much the same way that positive measures are defined from density functions.

Suppose that $(S, \mathcal{S}, \lambda)$ and (T, \mathcal{T}, μ) are measure spaces. As usual, $S \times T$ is given the product σ -algebra $\mathcal{S} \otimes \mathcal{T}$. If $k : S \times T \rightarrow [0, \infty)$ is measurable, then the function K defined as follows is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) :

$$K(x, A) = \int_A k(x, y) \mu(dy), \quad x \in S, A \in \mathcal{T} \quad (4.13.9)$$

Proof

The measurability of $x \mapsto K(x, A) = \int_A k(x, y) \mu(dy)$ for $A \in \mathcal{T}$ follows from a basic property of the integral. The fact that $A \mapsto K(x, A) = \int_A k(x, y) \mu(dy)$ is a positive measure on \mathcal{T} for $x \in S$ also follows from a basic property of the integral. In fact, $y \mapsto k(x, y)$ is the density of this measure with respect to μ .

Clearly the kernel K depends on the positive measure μ on (T, \mathcal{T}) as well as the function k , while the measure λ on (S, \mathcal{S}) plays no role (and so is not even necessary). But again, our point of view is that the spaces have fixed, natural measures. Appropriately enough, the function k is called a *kernel density function* (with respect to μ), or simply a *kernel function*.

Suppose again that $(S, \mathcal{S}, \lambda)$ and (T, \mathcal{T}, μ) are measure spaces. Suppose also K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) with kernel function k . If $f : T \rightarrow \mathbb{R}$ is measurable, then, assuming that the integrals exists,

$$Kf(x) = \int_S k(x, y) f(y) \mu(dy), \quad x \in S \quad (4.13.10)$$

Proof

This follows since the function $y \mapsto k(x, y)$ is the density of the measure $A \mapsto K(x, A)$ with respect to μ :

$$Kf(x) = \int_S K(x, dy) f(y) = \int_S k(x, y) f(y) \mu(dy), \quad x \in S \quad (4.13.11)$$

A kernel function defines an operator on the left with functions on S in a completely analogous way to the operator on the right above with functions on T .

Suppose again that $(S, \mathcal{S}, \lambda)$ and (T, \mathcal{T}, μ) are measure spaces, and that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) with kernel function k . If $f : S \rightarrow \mathbb{R}$ is measurable, then the function $fK : T \rightarrow \mathbb{R}$ defined as follows is also measurable, assuming that the integrals exists

$$fK(y) = \int_S \lambda(dx) f(x) k(x, y), \quad y \in T \quad (4.13.12)$$

The operator defined above depends on the measure λ on (S, \mathcal{S}) as well as the kernel function k , while the measure μ on (T, \mathcal{T}) plays no role (and so is not even necessary). But again, our point of view is that the spaces have fixed, natural measures. Here is how our new operation on the left with *functions* relates to our old operation on the left with *measures*.

Suppose again that $(S, \mathcal{S}, \lambda)$ and (T, \mathcal{T}, μ) are measure spaces, and that K is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) with kernel function k . Suppose also that $f : S \rightarrow [0, \infty)$ is measurable, and let ρ denote the measure on (S, \mathcal{S}) that has density f with respect to λ . Then fK is the density of the measure ρK with respect to μ .

Proof

The main tool, as usual, is an interchange of integrals. For $B \in \mathcal{T}$,

$$\begin{aligned}\rho K(B) &= \int_S \rho(dx) K(x, B) = \int_S f(x) K(x, B) \lambda(dx) = \int_S f(x) \left[\int_B k(x, y) \mu(dy) \right] \lambda(dx) \\ &= \int_B \left[\int_S f(x) k(x, y) \lambda(dx) \right] \mu(dy) = \int_B fK(y) \mu(dy)\end{aligned}$$

As always, we are particularly interested in stochastic kernels. With a kernel function, we can have *doubly* stochastic kernels.

Suppose again that $(S, \mathcal{S}, \lambda)$ and (T, \mathcal{T}, μ) are measure spaces and that $k : S \times T \rightarrow [0, \infty)$ is measurable. Then k is a *double stochastic kernel function* if

1. $\int_T k(x, y) \mu(dy) = 1$ for $x \in S$
2. $\int_S \lambda(dx) k(x, y) = 1$ for $y \in S$

Of course, condition (a) simply means that the kernel associated with k is a stochastic kernel according to our original definition.

The most common and important special case is when the two spaces are the same. Thus, if $(S, \mathcal{S}, \lambda)$ is a measure space and $k : S \times S \rightarrow [0, \infty)$ is measurable, then we have an operator K that operates on the left and on the right with measurable functions $f : S \rightarrow \mathbb{R}$:

$$\begin{aligned}fK(y) &= \int_S \lambda(dx) f(x) k(x, y), \quad y \in S \\ Kf(x) &= \int_S k(x, y) f(y) \lambda(dy), \quad x \in S\end{aligned}$$

If f is nonnegative and μ is the measure on with density function f , then fK is the density function of the measure μK (both with respect to λ).

Suppose again that $(S, \mathcal{S}, \lambda)$ is a measure space and $k : S \times S \rightarrow [0, \infty)$ is measurable. Then k is *symmetric* if $k(x, y) = k(y, x)$ for all $(x, y) \in S^2$.

Of course, if k is a symmetric, stochastic kernel function on $(S, \mathcal{S}, \lambda)$ then k is doubly stochastic, but the converse is not true.

Suppose that $(R, \mathcal{R}, \lambda)$, (S, \mathcal{S}, μ) , and (T, \mathcal{T}, ρ) are measure spaces. Suppose also that K is a kernel from (R, \mathcal{R}) to (S, \mathcal{S}) with kernel function k , and that L is a kernel from (S, \mathcal{S}) to (T, \mathcal{T}) with kernel function l . Then the kernel KL from (R, \mathcal{R}) to (T, \mathcal{T}) has density kl given by

$$kl(x, z) = \int_S k(x, y) l(y, z) \mu(dy), \quad (x, z) \in R \times T \quad (4.13.13)$$

Proof

Once again, the main tool is an interchange of integrals via Fubini's theorem. Let $x \in R$ and $B \in \mathcal{T}$. Then

$$\begin{aligned}KL(x, B) &= \int_S K(x, dy) L(y, B) = \int_S k(x, y) L(y, B) \mu(dy) \\ &= \int_S k(x, y) \left[\int_B l(y, z) \rho(dz) \right] \mu(dy) = \int_B \left[\int_S k(x, y) l(y, z) \mu(dy) \right] \rho(dz) = \int_B kl(x, z) \rho(dz)\end{aligned}$$

Examples and Special Cases

The Discrete Case

In this subsection, we assume that the measure spaces are discrete, as described in (1). Since the σ -algebra (all subsets) and the measure (counting measure) are understood, we don't need to reference them. Recall that integrals with respect to counting measure are sums. Suppose now that K is a kernel from the discrete space S to the discrete space T . For $x \in S$ and $y \in T$, let $K(x, y) = K(x, \{y\})$. Then more generally,

$$K(x, A) = \sum_{y \in A} K(x, y), \quad x \in S, A \subseteq T \quad (4.13.14)$$

The function $(x, y) \mapsto K(x, y)$ is simply the kernel function of the kernel K , as defined above, but in this case we usually don't bother with using a different symbol for the function as opposed to the kernel. The function K can be thought of as a *matrix*, with rows indexed by S and columns indexed by T (and so an infinite matrix if S or T is countably infinite). With this interpretation, all of the operations defined above can be thought of as matrix operations. If $f : T \rightarrow \mathbb{R}$ and f is thought of as a column vector indexed by T , then Kf is simply the ordinary product of the matrix K and the vector f ; the product is a column vector indexed by S :

$$Kf(x) = \sum_{y \in S} K(x, y)f(y), \quad x \in S \quad (4.13.15)$$

Similarly, if $f : S \rightarrow \mathbb{R}$ and f is thought of as a row vector indexed by S , then fK is simple the ordinary product of the vector f and the matrix K ; the product is a row vector indexed by T :

$$fK(y) = \sum_{x \in S} f(x)K(x, y), \quad y \in T \quad (4.13.16)$$

If L is another kernel from T to another discrete space U , then as functions, KL is the simply the matrix product of K and L :

$$KL(x, z) = \sum_{y \in T} K(x, y)L(y, z), \quad (x, z) \in S \times U \quad (4.13.17)$$

Let $S = \{1, 2, 3\}$ and $T = \{1, 2, 3, 4\}$. Define the kernel K from S to T by $K(x, y) = x + y$ for $(x, y) \in S \times T$. Define the function f on S by $f(x) = x!$ for $x \in S$, and define the function g on T by $g(y) = y^2$ for $y \in T$. Compute each of the following using matrix algebra:

1. fK
2. Kg

Answer

In matrix form,

$$K = \begin{bmatrix} 2 & 3 & 4 & 5 \\ 3 & 4 & 5 & 6 \\ 4 & 5 & 6 & 7 \end{bmatrix}, \quad f = [1 \quad 2 \quad 6], \quad g = \begin{bmatrix} 1 \\ 4 \\ 9 \\ 16 \end{bmatrix} \quad (4.13.18)$$

1. As a row vector indexed by T , the product is $fK = [32 \quad 41 \quad 50 \quad 59]$
2. As a column vector indexed by S ,

$$Kg = \begin{bmatrix} 130 \\ 160 \\ 190 \end{bmatrix} \quad (4.13.19)$$

Let $R = \{0, 1\}$, $S = \{a, b\}$, and $T = \{1, 2, 3\}$. Define the kernel K from R to S , the kernel L from S to S and the kernel M from S to T in matrix form as follows:

$$K = \begin{bmatrix} 1 & 4 \\ 2 & 3 \end{bmatrix}, \quad L = \begin{bmatrix} 2 & 2 \\ 1 & 5 \end{bmatrix}, \quad M = \begin{bmatrix} 1 & 0 & 2 \\ 0 & 3 & 1 \end{bmatrix} \quad (4.13.20)$$

Compute each of the following kernels, or explain why the operation does not make sense:

1. KL
2. LK
3. K^2
4. L^2
5. KM
6. LM

Proof

Note that these are not just abstract matrices, but rather have rows and columns indexed by the appropriate spaces. So the products make sense only when the spaces match appropriately; it's not just a matter of the number of rows and columns.

1. KL is the kernel from R to S given by

$$KL = \begin{bmatrix} 6 & 22 \\ 7 & 19 \end{bmatrix} \quad (4.13.21)$$

2. LK is not defined since the column space S of L is not the same as the row space R of K .
3. K^2 is not defined since the row space R is not the same as the column space S .
4. L^2 is the kernel from S to S given by

$$L^2 = \begin{bmatrix} 6 & 14 \\ 7 & 27 \end{bmatrix} \quad (4.13.22)$$

5. KM is the kernel from R to T given by

$$KM = \begin{bmatrix} 1 & 12 & 6 \\ 2 & 9 & 7 \end{bmatrix} \quad (4.13.23)$$

6. LM is the kernel from S to T given by

$$LM = \begin{bmatrix} 2 & 6 & 6 \\ 1 & 15 & 7 \end{bmatrix} \quad (4.13.24)$$

Conditional Probability

An important class of probability kernels arises from the distribution of one random variable, conditioned on the value of another random variable. In this subsection, suppose that $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, and that (S, \mathcal{S}) and (T, \mathcal{T}) are measurable spaces. Further, suppose that X and Y are random variables defined on the probability space, with X taking values in S and that Y taking values in T . Informally, X and Y are random variables defined on the same underlying random experiment.

The function P defined as follows is a probability kernel from (S, \mathcal{S}) to (T, \mathcal{T}) , known as the *conditional probability kernel* of Y given X .

$$P(x, A) = \mathbb{P}(Y \in A \mid X = x), \quad x \in S, A \in \mathcal{T} \quad (4.13.25)$$

Proof

Recall that for $A \in \mathcal{T}$, the conditional probability $\mathbb{P}(Y \in A \mid X)$ is itself a random variable, and is measurable with respect to $\sigma(X)$. That is, $\mathbb{P}(Y \in A \mid X) = P(X, A)$ for some measurable function $x \mapsto P(x, A)$ from S to $[0, 1]$. Then, by definition, $\mathbb{P}(Y \in A \mid X = x) = P(x, A)$. Trivially, of course, $A \mapsto P(x, A)$ is a probability measure on (T, \mathcal{T}) for $x \in S$.

The operators associated with this kernel have natural interpretations.

Let P be the conditional probability kernel of Y given X .

1. If $f : T \rightarrow \mathbb{R}$ is measurable, then $Pf(x) = \mathbb{E}[f(Y) \mid X = x]$ for $x \in S$ (assuming as usual that the expected value exists).
2. If μ is the probability distribution of X then μP is the probability distribution of Y .

Proof

These are basic results that we have already studied, dressed up in new notation.

1. Since $A \mapsto P(x, A)$ is the conditional distribution of Y given $X = x$,

$$\mathbb{E}[f(Y) \mid X = x] = \int_S P(x, dy) f(y) = P f(x) \quad (4.13.26)$$

2. Let $A \in \mathcal{T}$. Conditioning on X gives

$$\mathbb{P}(Y \in A) = \mathbb{E}[\mathbb{P}(Y \in A \mid X)] = \int_S \mu(dx) P(Y \in A \mid X = x) = \int_S \mu(dx) P(x, A) = \mu P(A) \quad (4.13.27)$$

As in the general discussion above, the measurable spaces (S, \mathcal{S}) and (T, \mathcal{T}) are usually *measure* spaces with natural measures attached. So the conditional probability distributions are often given via conditional probability density functions, which then play the role of kernel functions. The next two exercises give examples.

Suppose that X and Y are random variables for an experiment, taking values in \mathbb{R} . For $x \in \mathbb{R}$, the conditional distribution of Y given $X = x$ is normal with mean x and standard deviation 1. Use the notation and operations of this section for the following computations:

1. Give the kernel function for the conditional distribution of Y given X .
2. Find $\mathbb{E}(Y^2 \mid X = x)$.
3. Suppose that X has the standard normal distribution. Find the probability density function of Y .

Answer

1. The kernel function (with respect to Lebesgue measure, of course) is

$$p(x, y) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(y-x)^2}, \quad x, y \in \mathbb{R} \quad (4.13.28)$$

2. Let $g(y) = y^2$ for $y \in \mathbb{R}$. Then $E(Y^2 \mid X = x) = P g(x) = 1 + x^2$ for $x \in \mathbb{R}$
3. The standard normal PDF f is given $f(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ for $x \in \mathbb{R}$. Thus Y has PDF fP .

$$fP(y) \int_{-\infty}^{\infty} f(x) p(x, y) dx = \frac{1}{2\sqrt{\pi}} e^{-\frac{1}{4}y^2}, \quad y \in \mathbb{R} \quad (4.13.29)$$

This is the PDF of the normal distribution with mean 0 and variance 2.

Suppose that X and Y are random variables for an experiment, with X taking values in $\{a, b, c\}$ and Y taking values in $\{1, 2, 3, 4\}$. The kernel function of Y given X is as follows: $P(a, y) = 1/4$, $P(b, y) = y/10$, and $P(c, y) = y^2/30$, each for $y \in \{1, 2, 3, 4\}$.

1. Give the kernel P in matrix form and verify that it is a probability kernel.
2. Find fP where $f(a) = f(b) = f(c) = 1/3$. The result is the density function of Y given that X is uniformly distributed.
3. Find Pg where $g(y) = y$ for $y \in \{1, 2, 3, 4\}$. The resulting function is $\mathbb{E}(Y \mid X = x)$ for $x \in \{a, b, c\}$.

Answer

1. P is given in matrix form below. Note that the row sums are 1.

$$P = \begin{bmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{10} & \frac{2}{10} & \frac{3}{10} & \frac{4}{10} \\ \frac{1}{30} & \frac{4}{30} & \frac{9}{30} & \frac{16}{30} \end{bmatrix} \quad (4.13.30)$$

2. In matrix form, $f = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$ and $fP = \begin{bmatrix} \frac{23}{180} & \frac{35}{180} & \frac{51}{180} & \frac{71}{180} \end{bmatrix}$.
3. In matrix form,

$$g = \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \end{bmatrix}, \quad Pg = \begin{bmatrix} \frac{5}{2} \\ 3 \\ \frac{10}{3} \end{bmatrix} \quad (4.13.31)$$

Parametric Distributions

A parametric probability distribution also defines a probability kernel in a natural way, with the parameter playing the role of the kernel variable, and the distribution playing the role of the measure. Such distributions are usually defined in terms of a parametric density function which then defines a kernel function, again with the parameter playing the role of the first argument and the variable the role of the second argument. If the parameter is thought of as a given value of another random variable, as in Bayesian analysis, then there is considerable overlap with the previous subsection. In most cases, (and in particular in the examples below), the spaces involved are either discrete or Euclidean, as described in (1).

Consider the parametric family of exponential distributions. Let f denote the identity function on $(0, \infty)$.

1. Give the probability density function as a probability kernel function p on $(0, \infty)$.
2. Find Pf .
3. Find fP .
4. Find p^2 , the kernel function corresponding to the product kernel P^2 .

Answer

1. $p(r, x) = re^{-rx}$ for $r, x \in (0, \infty)$.
2. For $r \in (0, \infty)$,

$$Pf(r) = \int_0^\infty p(r, x)f(x) dx = \int_0^\infty xre^{-rx} dx = \frac{1}{r} \quad (4.13.32)$$

This is the mean of the exponential distribution.

3. For $x \in (0, \infty)$,

$$fP(x) = \int_0^\infty f(r)p(r, x) dr = \int_0^\infty r^2e^{-rx} dr = \frac{2}{x^3} \quad (4.13.33)$$

4. For $r, y \in (0, \infty)$,

$$p^2(r, y) = \int_0^\infty p(r, x)p(x, y) dx = \int_0^\infty rxe^{-(r+y)x} dx = \frac{r}{(r+y)^2} \quad (4.13.34)$$

Consider the parametric family of Poisson distributions. Let f be the identity function on \mathbb{N} and let g be the identity function on $(0, \infty)$.

1. Give the probability density function p as a probability kernel function from $(0, \infty)$ to \mathbb{N} .
2. Show that $Pf = g$.
3. Show that $gP = f$.

Answer

1. $p(r, n) = e^{-r} \frac{r^n}{n!}$ for $r \in (0, \infty)$ and $n \in \mathbb{N}$.
2. For $r \in (0, \infty)$, $Pf(r)$ is the mean of the Poisson distribution with parameter r :

$$Pf(r) = \sum_{n=0}^\infty p(r, n)f(n) = \sum_{n=0}^\infty ne^{-r} \frac{r^n}{n!} = r \quad (4.13.35)$$

3. For $n \in \mathbb{N}$,

$$gP(n) = \int_0^\infty g(r)p(r, n) dr = \int_0^\infty e^{-r} \frac{r^{n+1}}{n!} dr = n \quad (4.13.36)$$

Clearly the Poisson distribution has some very special and elegant properties. The next family of distributions also has some very special properties. Compare this exercise with the exercise (30).

Consider the family of normal distributions, parameterized by the mean and with variance 1.

1. Give the probability density function as a probability kernel function p on \mathbb{R} .

2. Show that p is symmetric.
3. Let f be the identity function on \mathbb{R} . Show that $Pf = f$ and $fP = f$.
4. For $n \in \mathbb{N}$, find p^n the kernel function for the operator P^n .

Answer

1. For $\mu, x \in \mathbb{R}$,

$$p(\mu, x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu)^2} \quad (4.13.37)$$

That is, $x \mapsto p(x, \mu)$ is the normal probability density function with mean μ and variance 1.

2. Note that $p(\mu, x) = p(x, \mu)$ for $\mu, x \in \mathbb{R}$. So $\mu \mapsto p(\mu, x)$ is the normal probability density function with mean x and variance 1.
3. Since $f(x) = x$ for $x \in \mathbb{R}$, this follows from the previous two parts: $Pf(\mu) = \mu$ for $\mu \in \mathbb{R}$ and $fP(x) = x$ for $x \in \mathbb{R}$
4. For $\mu, y \in \mathbb{R}$,

$$p^2(\mu, x) = \int_{-\infty}^{\infty} p(\mu, t)p(t, y) dt = \frac{1}{\sqrt{4\pi}} e^{-\frac{1}{4}(x-\mu)^2} \quad (4.13.38)$$

so that $x \mapsto p^2(\mu, x)$ is the normal PDF with mean μ and variance 2. By induction,

$$p^n(\mu, x) = \frac{1}{\sqrt{2\pi n}} e^{-\frac{1}{2n}(x-\mu)^2} \quad (4.13.39)$$

for $n \in \mathbb{N}_+$ and $\mu, x \in \mathbb{R}$. Thus $x \mapsto p^n(\mu, x)$ is the normal PDF with mean μ and variance n .

For each of the following special distributions, express the probability density function as a probability kernel function. Be sure to specify the parameter spaces.

1. The general normal distribution on \mathbb{R} .
2. The beta distribution on $(0, 1)$.
3. The negative binomial distribution on \mathbb{N} .

Answer

1. The normal distribution with mean μ and standard deviation σ defines a kernel function p from $\mathbb{R} \times (0, \infty)$ to \mathbb{R} given by

$$p[(\mu, \sigma), x] = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\left(\frac{x-\mu}{\sigma}\right)^2\right] \quad (4.13.40)$$

2. The beta distribution with left parameter a and right parameter b defines a kernel function p from $(0, \infty)^2$ to $(0, 1)$ given by

$$p[(a, b), x] = \frac{1}{B(a, b)} x^{a-1} y^{b-1} \quad (4.13.41)$$

where B is the beta function.

3. The negative binomial distribution with stopping parameter k and success parameter α defines a kernel function p from $(0, \infty) \times (0, 1)$ to \mathbb{N} given by

$$p[(n, \alpha), k] = \binom{n+k-1}{n} \alpha^k (1-\alpha)^n \quad (4.13.42)$$

This page titled [4.13: Kernels and Operators](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

5: Special Distributions

In this chapter, we study several general families of probability distributions and a number of special parametric families of distributions. Unlike the other expository chapters in this text, the sections are not linearly ordered and so this chapter serves primarily as a reference. You may want to study these topics as the need arises.

First, we need to discuss what makes a probability distribution *special* in the first place. In some cases, a distribution may be important because it is connected with other special distributions in interesting ways (via transformations, limits, conditioning, etc.). In some cases, a parametric family may be important because it can be used to model a wide variety of random phenomena. This may be the case because of fundamental underlying principles, or simply because the family has a rich collection of probability density functions with a small number of parameters (usually 3 or less). As a general philosophical principle, we try to model a random process with as few parameters as possible; this is sometimes referred to as the principle of *parsimony of parameters*. In turn, this is a special case of *Ockham's razor*, named in honor of William of Ockham, the principle that states that one should use the simplest model that adequately describes a given phenomenon. Parsimony is important because often the parameters are not known and must be estimated.

In many cases, a special parametric family of distributions will have one or more distinguished *standard members*, corresponding to specified values of some of the parameters. Usually the standard distributions will be mathematically simplest, and often other members of the family can be constructed from the standard distributions by simple transformations on the underlying standard random variable.

An incredible variety of special distributions have been studied over the years, and new ones are constantly being added to the literature. To truly deserve the adjective *special*, a distribution should have a certain level of mathematical elegance and economy, and should arise in interesting and diverse applications.

[5.1: Location-Scale Families](#)

[5.2: General Exponential Families](#)

[5.3: Stable Distributions](#)

[5.4: Infinitely Divisible Distributions](#)

[5.5: Power Series Distributions](#)

[5.6: The Normal Distribution](#)

[5.7: The Multivariate Normal Distribution](#)

[5.8: The Gamma Distribution](#)

[5.9: Chi-Square and Related Distribution](#)

[5.10: The Student t Distribution](#)

[5.11: The F Distribution](#)

[5.12: The Lognormal Distribution](#)

[5.13: The Folded Normal Distribution](#)

[5.14: The Rayleigh Distribution](#)

[5.15: The Maxwell Distribution](#)

[5.16: The Lévy Distribution](#)

[5.17: The Beta Distribution](#)

[5.18: The Beta Prime Distribution](#)

[5.19: The Arcsine Distribution](#)

[5.20: General Uniform Distributions](#)

[5.21: The Uniform Distribution on an Interval](#)

[5.22: Discrete Uniform Distributions](#)

[5.23: The Semicircle Distribution](#)

[5.24: The Triangle Distribution](#)

- [5.25: The Irwin-Hall Distribution](#)
- [5.26: The U-Power Distribution](#)
- [5.27: The Sine Distribution](#)
- [5.28: The Laplace Distribution](#)
- [5.29: The Logistic Distribution](#)
- [5.30: The Extreme Value Distribution](#)
- [5.31: The Hyperbolic Secant Distribution](#)
- [5.32: The Cauchy Distribution](#)
- [5.33: The Exponential-Logarithmic Distribution](#)
- [5.34: The Gompertz Distribution](#)
- [5.35: The Log-Logistic Distribution](#)
- [5.36: The Pareto Distribution](#)
- [5.37: The Wald Distribution](#)
- [5.38: The Weibull Distribution](#)
- [5.39: Benford's Law](#)
- [5.40: The Zeta Distribution](#)
- [5.41: The Logarithmic Series Distribution](#)

This page titled [5: Special Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.1: Location-Scale Families

General Theory

As usual, our starting point is a random experiment modeled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so that Ω is the set of outcomes, \mathcal{F} the collection of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . In this section, we assume that we fixed random variable Z defined on the probability space, taking values in \mathbb{R} .

Definition

For $a \in \mathbb{R}$ and $b \in (0, \infty)$, let $X = a + bZ$. The two-parameter family of distributions associated with X is called the *location-scale family* associated with the given distribution of Z . Specifically, a is the *location parameter* and b the *scale parameter*.

Thus a *linear transformation*, with positive slope, of the underlying random variable Z creates a *location-scale family* for the underlying distribution. In the special case that $b = 1$, the one-parameter family is called the *location family* associated with the given distribution, and in the special case that $a = 0$, the one-parameter family is called the *scale family* associated with the given distribution. Scale transformations, as the name suggests, occur naturally when physical units are changed. For example, if a random variable represents the length of an object, then a change of units from meters to inches corresponds to a scale transformation. Location transformations often occur when the zero reference point is changed, in measuring distance or time, for example. Location-scale transformations can also occur with a change of physical units. For example, if a random variable represents the temperature of an object, then a change of units from Fahrenheit to Celsius corresponds to a location-scale transformation.

Distribution Functions

Our goal is to relate various functions that determine the distribution of $X = a + bZ$ to the corresponding functions for Z . First we consider the (cumulative) distribution function.

If Z has distribution function G then X has distribution function F given by

$$F(x) = G\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.1.1)$$

Proof

For $x \in \mathbb{R}$

$$F(x) = \mathbb{P}(X \leq x) = \mathbb{P}(a + bZ \leq x) = \mathbb{P}\left(Z \leq \frac{x-a}{b}\right) = G\left(\frac{x-a}{b}\right) \quad (5.1.2)$$

Next we consider the probability density function. The results are a bit different for discrete distributions and continuous distribution, not surprising since the density function has different meanings in these two cases.

If Z has a discrete distribution with probability density function g then X also has a discrete distribution, with probability density function f given by

$$f(x) = g\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.1.3)$$

Proof

Z takes values in a countable subset $S \subset \mathbb{R}$ and hence X takes values in $T = \{a + bz : z \in S\}$, which is also countable. Moreover

$$f(x) = \mathbb{P}(X = x) = \mathbb{P}\left(Z = \frac{x-a}{b}\right) = g\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.1.4)$$

If Z has a continuous distribution with probability density function g , then X also has a continuous distribution, with probability density function f given by

$$f(x) = \frac{1}{b} g\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.1.5)$$

1. For the location family associated with g , the graph of f is obtained by shifting the graph of g , a units to the right if $a > 0$ and $-a$ units to the left if $a < 0$.
2. For the scale family associated with g , if $b > 1$, the graph of f is obtained from the graph of g by stretching horizontally and compressing vertically, by a factor of b . If $0 < b < 1$, the graph of f is obtained from the graph of g by compressing horizontally and stretching vertically, by a factor of b .

Proof

First note that $\mathbb{P}(X = x) = \mathbb{P}\left(Z = \frac{x-a}{b}\right) = 0$, so X has a continuous distribution. Typically, Z takes values in an interval of \mathbb{R} and thus so does X . The formula for the density function follows by taking derivatives of the [distribution function](#) above, since $f = F'$ and $g = G'$.

If Z has a mode at z , then X has a mode at $x = a + bz$.

Proof

This follows from [density function](#) in the discrete case or the [density function](#) in the continuous case. If g has a maximum at z then f has a maximum at $x = a + bz$.

Next we relate the quantile functions of Z and X .

If G and F are the distribution functions of Z and X , respectively, then

1. $F^{-1}(p) = a + b G^{-1}(p)$ for $p \in (0, 1)$
2. If z is a quantile of order p for Z then $x = a + bz$ is a quantile of order p for X .

Proof

These results follow from the [distribution function](#) above.

Suppose now that Z has a continuous distribution on $[0, \infty)$, and that we think of Z as the failure time of a device (or the time of death of an organism). Let $X = bZ$ where $b \in [0, \infty)$, so that the distribution of X is the scale family associated with the distribution of Z . Then X also has a continuous distribution on $[0, \infty)$ and can also be thought of as the failure time of a device (perhaps in different units).

Let G^c and F^c denote the reliability functions of Z and X respectively, and let r and R denote the failure rate functions of Z and X , respectively. Then

1. $F^c(x) = G^c(x/b)$ for $x \in [0, \infty)$
2. $R(x) = \frac{1}{b} r\left(\frac{x}{b}\right)$ for $x \in [0, \infty)$

Proof

Recall that $G^c = 1 - G$, $F^c = 1 - F$, $r = g/\bar{G}$, and $R = f/\bar{F}$. Thus the results follow from the [distribution function](#) and the [density function](#) above.

Moments

The following theorem relates the mean, variance, and standard deviation of Z and X .

As before, suppose that $X = a + bZ$. Then

1. $\mathbb{E}(X) = a + b \mathbb{E}(Z)$
2. $\text{var}(X) = b^2 \text{var}(Z)$
3. $\text{sd}(X) = b \text{sd}(Z)$

Proof

These result follow immediately from basic properties of expected value and variance.

Recall that the *standard score* of a random variable is obtained by subtracting the mean and dividing by the standard deviation. The standard score is dimensionless (that is, has no physical units) and measures the distance from the mean to the random variable in standard deviations. Since location-scale families essentially correspond to a change of units, it's not surprising that the standard score is unchanged by a location-scale transformation.

The standard scores of X and Z are the same:

$$\frac{X - \mathbb{E}(X)}{\text{sd}(X)} = \frac{Z - \mathbb{E}(Z)}{\text{sd}(Z)} \quad (5.1.6)$$

Proof

From the [mean and variance](#) above:

$$\frac{X - \mathbb{E}(X)}{\text{sd}(X)} = \frac{a + bZ - [a + b\mathbb{E}(Z)]}{b\text{sd}(Z)} = \frac{Z - \mathbb{E}(Z)}{\text{sd}(Z)} \quad (5.1.7)$$

Recall that the skewness and kurtosis of a random variable are the third and fourth moments, respectively, of the standard score. Thus it follows from the previous result that skewness and kurtosis are unchanged by location-scale transformations: $\text{skew}(X) = \text{skew}(Z)$, $\text{kurt}(X) = \text{kurt}(Z)$.

We can represent the moments of X (about 0) to those of Z by means of the binomial theorem:

$$\mathbb{E}(X^n) = \sum_{k=0}^n \binom{n}{k} b^k a^{n-k} \mathbb{E}(Z^k), \quad n \in \mathbb{N} \quad (5.1.8)$$

Of course, the moments of X about the location parameter a have a simple representation in terms of the moments of Z about 0:

$$\mathbb{E}[(X - a)^n] = b^n \mathbb{E}(Z^n), \quad n \in \mathbb{N} \quad (5.1.9)$$

The following exercise relates the moment generating functions of Z and X .

If Z has moment generating function m then X has moment generating function M given by

$$M(t) = e^{at} m(bt) \quad (5.1.10)$$

Proof

$$M(t) = \mathbb{E}(e^{tX}) = \mathbb{E}[e^{t(a+bZ)}] = e^{ta} \mathbb{E}(e^{tbZ}) = e^{at} m(bt) \quad (5.1.11)$$

Type

As we noted earlier, two probability distributions that are related by a location-scale transformation can be thought of as governing the same underlying random quantity, but in different physical units. This relationship is important enough to deserve a name.

Suppose that P and Q are probability distributions on \mathbb{R} with distribution functions F and G , respectively. Then P and Q are of the same *type* if there exist constants $a \in \mathbb{R}$ and $b \in (0, \infty)$ such that

$$F(x) = G\left(\frac{x - a}{b}\right), \quad x \in \mathbb{R} \quad (5.1.12)$$

Being of the same type is an equivalence relation on the collection of probability distributions on \mathbb{R} . That is, if P , Q , and R are probability distribution on \mathbb{R} then

1. P is the same type as P (the *reflexive* property).
2. If P is the same type as Q then Q is the same type as P (the *symmetric* property).
3. If P is the same type as Q , and Q is the same type as R , then P is the same type as R (the *transitive* property).

Proof

Let F , G , and H denote the distribution functions of P , Q , and R respectively.

1. This is trivial, of course, since we can take $a = 0$ and $b = 1$.

2. Suppose there exists $a \in \mathbb{R}$ and $b \in (0, \infty)$ such that $F(x) = G\left(\frac{x-a}{b}\right)$ for $x \in \mathbb{R}$. Then

$$G(x) = F(a + bx) = F\left(\frac{x - (-a/b)}{1/b}\right) \text{ for } x \in \mathbb{R}.$$

3. Suppose there exists $a, c \in \mathbb{R}$ and $b, d \in (0, \infty)$ such that $F(x) = G\left(\frac{x-a}{b}\right)$ and $G(x) = H\left(\frac{x-c}{d}\right)$ for $x \in \mathbb{R}$. Then

$$F(x) = H\left(\frac{x - (a+bc)}{bd}\right) \text{ for } x \in \mathbb{R}.$$

So, the collection of probability distributions on \mathbb{R} is partitioned into mutually exclusive equivalence classes, where the distributions in each class are all of the same type.

Examples and Applications

Special Distributions

Many of the special parametric families of distributions studied in this chapter and elsewhere in this text are location and/or scale families.

The [arcsine distribution](#) is a location-scale family.

The [Cauchy distribution](#) is a location-scale family.

The [exponential distribution](#) is a scale family.

The [exponential-logarithmic distribution](#) is a scale family for each value of the shape parameter.

The [extreme value distribution](#) is a location-scale family.

The [gamma distribution](#) is a scale family for each value of the shape parameter.

The [Gompertz distribution](#) is a scale family for each value of the shape parameter.

The [half-normal distribution](#) is a scale family.

The [hyperbolic secant](#) distribution is a location-scale family.

The [Lévy distribution](#) is a location scale family.

The [logistic distribution](#) is a location-scale family.

The [log-logistic distribution](#) is a scale family for each value of the shape parameter.

The [Maxwell distribution](#) is a scale family.

The [normal distribution](#) is a location-scale family.

The [Pareto distribution](#) is a scale family for each value of the shape parameter.

The [Rayleigh distribution](#) is a scale family.

The [semicircle distribution](#) is a location-scale family.

The [triangle distribution](#) is a location-scale family for each value of the shape parameter.

The [uniform distribution](#) on an interval is a location-scale family.

The [U-power distribution](#) is a location-scale family for each value of the shape parameter.

The [Weibull distribution](#) is a scale family for each value of the shape parameter.

The [Wald distribution](#) is a scale family, although in the usual formulation, neither of the parameters is a scale parameter.

This page titled [5.1: Location-Scale Families](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.2: General Exponential Families

Basic Theory

Definition

We start with a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ as a model for a random experiment. So as usual, Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . For the general formulation that we want in this section, we need two additional spaces, a measure space (S, \mathcal{S}, μ) (where the probability distributions will live) and a measurable space (T, \mathcal{T}) (serving the role of a *parameter space*). Typically, these spaces fall into our two standard categories. Specifically, the measure space is usually one of the following:

- *Discrete.* S is countable, \mathcal{S} is the collection of all subsets of S , and $\mu = \#$ is counting measure.
- *Euclidean.* S is a sufficiently nice Borel measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$, \mathcal{S} is the σ -algebra of Borel measurable subsets of S , and $\mu = \lambda_n$ is n -dimensional Lebesgue measure.

Similarly, the parameter space (T, \mathcal{T}) is usually either discrete, so that T is countable and \mathcal{T} the collection of all subsets of T , or Euclidean so that T is a sufficiently nice Borel measurable subset of \mathbb{R}^m for some $m \in \mathbb{N}_+$ and \mathcal{T} is the σ -algebra of Borel measurable subsets of T .

Suppose now that X is random variable defined on the probability space, taking values in S , and that the distribution of X depends on a parameter $\theta \in T$. For $\theta \in T$ we assume that the distribution of X has probability density function f_θ with respect to μ .

for $k \in \mathbb{N}_+$, the family of distributions of X is a k -parameter *exponential family* if

$$f_\theta(x) = \alpha(\theta) g(x) \exp\left(\sum_{i=1}^k \beta_i(\theta) h_i(x)\right); \quad x \in S, \theta \in T \quad (5.2.1)$$

where α and $(\beta_1, \beta_2, \dots, \beta_k)$ are measurable functions from T into \mathbb{R} , and where g and (h_1, h_2, \dots, h_k) are measurable functions from S into \mathbb{R} . Moreover, k is assumed to be the smallest such integer.

1. The parameters $(\beta_1(\theta), \beta_2(\theta), \dots, \beta_k(\theta))$ are called the *natural parameters* of the distribution.
2. the random variables $(h_1(X), h_2(X), \dots, h_k(X))$ are called the *natural statistics* of the distribution.

Although the definition may look intimidating, exponential families are useful because many important theoretical results in statistics hold for exponential families, and because many special parametric families of distributions turn out to be exponential families. It's important to emphasize that the representation of $f_\theta(x)$ given in the definition must hold for all $x \in S$ and $\theta \in T$. If the representation only holds for a set of $x \in S$ that depends on the particular $\theta \in T$, then the family of distributions is not a general exponential family.

The next result shows that if we sample from the distribution of an exponential family, then the distribution of the random sample is itself an exponential family with the same natural parameters.

Suppose that the distribution of random variable X is a k -parameter exponential family with natural parameters $(\beta_1(\theta), \beta_2(\theta), \dots, \beta_k(\theta))$, and natural statistics $(h_1(X), h_2(X), \dots, h_k(X))$. Let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be a sequence of n independent random variables, each with the same distribution as X . Then \mathbf{X} is a k -parameter exponential family with natural parameters $(\beta_1(\theta), \beta_2(\theta), \dots, \beta_k(\theta))$, and natural statistics

$$u_j(\mathbf{X}) = \sum_{i=1}^n h_j(X_i), \quad j \in \{1, 2, \dots, k\} \quad (5.2.2)$$

Proof

Let f_θ denote the PDF of X corresponding to the parameter value $\theta \in T$, so that $f_\theta(x)$ has the representation given in the [definition](#) for $x \in S$ and $\theta \in T$. Then for $\theta \in T$, $\mathbf{X} = (X_1, X_2, \dots, X_n)$ has PDF g_θ given by

$$g_\theta(x_1, x_2, \dots, x_n) = f_\theta(x_1) f_\theta(x_2) \cdots f_\theta(x_n), \quad (x_1, x_2, \dots, x_n) \in S^n \quad (5.2.3)$$

Substituting and simplifying gives the result.

Examples and Special Cases

Special Distributions

Many of the special distributions studied in this chapter are general exponential families, at least with respect to some of their parameters. On the other hand, most commonly, a parametric family fails to be a general exponential family because the support set depends on the parameter. The following theorems give a number of examples. Proofs will be provided in the individual sections.

The Bernoulli distribution is a one parameter exponential family in the success parameter $p \in [0, 1]$

The beta distribution is a two-parameter exponential family in the shape parameters $a \in (0, \infty)$, $b \in (0, \infty)$.

The beta prime distribution is a two-parameter exponential family in the shape parameters $a \in (0, \infty)$, $b \in (0, \infty)$.

The binomial distribution is a one-parameter exponential family in the success parameter $p \in [0, 1]$ for a fixed value of the trial parameter $n \in \mathbb{N}_+$.

The chi-square distribution is a one-parameter exponential family in the degrees of freedom $n \in (0, \infty)$.

The exponential distribution is a one-parameter exponential family (appropriately enough), in the rate parameter $r \in (0, \infty)$.

The gamma distribution is a two-parameter exponential family in the shape parameter $k \in (0, \infty)$ and the scale parameter $b \in (0, \infty)$.

The geometric distribution is a one-parameter exponential family in the success probability $p \in (0, 1)$.

The half normal distribution is a one-parameter exponential family in the scale parameter $\sigma \in (0, \infty)$

The Laplace distribution is a one-parameter exponential family in the scale parameter $b \in (0, \infty)$ for a fixed value of the location parameter $a \in \mathbb{R}$.

The Lévy distribution is a one-parameter exponential family in the scale parameter $b \in (0, \infty)$ for a fixed value of the location parameter $a \in \mathbb{R}$.

The logarithmic distribution is a one-parameter exponential family in the shape parameter $p \in (0, 1)$

The lognormal distribution is a two parameter exponential family in the shape parameters $\mu \in \mathbb{R}$, $\sigma \in (0, \infty)$.

The Maxwell distribution is a one-parameter exponential family in the scale parameter $b \in (0, \infty)$.

The k -dimensional multinomial distribution is a k -parameter exponential family in the probability parameters (p_1, p_2, \dots, p_k) for a fixed value of the trial parameter $n \in \mathbb{N}_+$.

The k -dimensional multivariate normal distribution is a $\frac{1}{2}(k^2 + 3k)$ -parameter exponential family with respect to the mean vector $\boldsymbol{\mu}$ and the variance-covariance matrix \mathbf{V} .

The negative binomial distribution is a one-parameter exponential family in the success parameter $p \in (0, 1)$ for a fixed value of the stopping parameter $k \in \mathbb{N}_+$.

The normal distribution is a two-parameter exponential family in the mean $\mu \in \mathbb{R}$ and the standard deviation $\sigma \in (0, \infty)$.

The Pareto distribution is a one-parameter exponential family in the shape parameter for a fixed value of the scale parameter.

The Poisson distribution is a one-parameter exponential family.

The Rayleigh distribution is a one-parameter exponential family.

The U-power distribution is a one-parameter exponential family in the shape parameter, for fixed values of the location and scale parameters.

The Weibull distribution is a one-parameter exponential family in the scale parameter for a fixed value of the shape parameter.

The zeta distribution is a one-parameter exponential family.

The Wald distribution is a two-parameter exponential family.

This page titled [5.2: General Exponential Families](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.3: Stable Distributions

This section discusses a theoretical topic that you may want to skip if you are a new student of probability.

Basic Theory

Stable distributions are an important general class of probability distributions on \mathbb{R} that are defined in terms of location-scale transformations. Stable distributions occur as limits (in distribution) of scaled and centered sums of independent, identically distributed variables. Such limits generalize the central limit theorem, and so stable distributions generalize the normal distribution in a sense. The pioneering work on stable distributions was done by Paul Lévy.

Definition

In this section, we consider real-valued random variables whose distributions are not degenerate (that is, not concentrated at a single value). After all, a random variable with a degenerate distribution is not really *random*, and so is not of much interest.

Random variable X has a *stable distribution* if the following condition holds: If $n \in \mathbb{N}_+$ and (X_1, X_2, \dots, X_n) is a sequence of independent variables, each with the same distribution as X , then $X_1 + X_2 + \dots + X_n$ has the same distribution as $a_n + b_n X$ for some $a_n \in \mathbb{R}$ and $b_n \in (0, \infty)$. If $a_n = 0$ for $n \in \mathbb{N}_+$ then the distribution of X is *strictly stable*.

1. The parameters a_n for $n \in \mathbb{N}_+$ are the *centering parameters*.
2. The parameters b_n for $n \in \mathbb{N}_+$ are the *norming parameters*.

Details

Since the distribution of X is not point mass at 0, note that if the distribution of $a + bX$ is the same as the distribution of $c + dX$ for some $a, c \in \mathbb{R}$ and $b, d \in (0, \infty)$, then $a = c$ and $b = d$. Thus, the centering parameters a_n and the norming parameters b_n are uniquely defined for $n \in \mathbb{N}_+$.

Recall that two distributions on \mathbb{R} that are related by a location-scale transformation are said to be of the *same type*, and that *being of the same type* defines an equivalence relation on the class of distributions on \mathbb{R} . With this terminology, the definition of stability has a more elegant expression: X has a stable distribution if the sum of a finite number of independent copies of X is of the same type as X . As we will see, the norming parameters are more important than the centering parameters, and in fact, only certain norming parameters can occur.

Basic Properties

We start with some very simple results that follow easily from the definition, before moving on to the deeper results.

Suppose that X has a stable distribution with mean μ and finite variance. Then the norming parameters are \sqrt{n} and the centering parameters are $(n - \sqrt{n})\mu$ for $n \in \mathbb{N}_+$.

Proof

As usual, let a_n and b_n denote the centering and norming parameters of X for $n \in \mathbb{N}_+$, and let σ^2 denote the (finite) variance of X . Suppose that $n \in \mathbb{N}_+$ and that (X_1, X_2, \dots, X_n) is a sequence of independent variables, each with the distribution of X . Then $X_1 + X_2 + \dots + X_n$ has the same distribution as $a_n + b_n X$. Taking variances gives $n\sigma^2 = b_n^2 \sigma^2$ and hence $b_n = \sqrt{n}$. Taking expected values now gives $n\mu = a_n + \sqrt{n}\mu$.

It will turn out that the only stable distribution with finite variance is the normal distribution, but the result above is useful as an intermediate step. Next, it seems fairly clear from the definition that the family of stable distributions is itself a location-scale family.

Suppose that the distribution of X is stable, with centering parameters $a_n \in \mathbb{R}$ and norming parameters $b_n \in (0, \infty)$ for $n \in \mathbb{N}_+$. If $c \in \mathbb{R}$ and $d \in (0, \infty)$, then the distribution of $Y = c + dX$ is also stable, with centering parameters $da_n + (n - b_n)c$ and norming parameters b_n for $n \in \mathbb{N}_+$.

Proof

Suppose that $n \in \mathbb{N}_+$ and that (Y_1, Y_2, \dots, Y_n) is a sequence of independent variables, each with the same distribution as Y . Then $Y_1 + Y_2 + \dots + Y_n$ has the same distribution as $nc + d(X_1 + X_2 + \dots + X_n)$ where (X_1, X_2, \dots) is a sequence of independent variables, each with the same distribution as X . By stability, $X_1 + X_2 + \dots + X_n$ has the same distribution as $a_n + b_n X$. Hence $Y_1 + Y_2 + \dots + Y_n$ has the same distribution as $(nc + da_n) + db_n X$, which in turn has the same distribution as $[da_n + (n - b_n)c] + b_n Y$.

An important point is the the norming parameters are unchanged under a location-scale transformation.

Suppose that the distribution of X is stable, with centering parameters $a_n \in \mathbb{R}$ and norming parameters $b_n \in (0, \infty)$ for $n \in \mathbb{N}_+$. Then the distribution of $-X$ is stable, with centering parameters $-a_n$ and norming parameters b_n for $n \in \mathbb{N}_+$.

Proof

If $n \in \mathbb{N}_+$ and (X_1, X_2, \dots, X_n) is a sequence of independent variables, each with the same distribution as X then $(-X_1, -X_2, \dots, -X_n)$ is a sequence of independent variables each with the same distribution as $-X$. By stability, $-\sum_{i=1}^n X_i$ has the same distribution as $-(a_n + b_n X) = -a_n + b_n(-X)$.

From the last two results, if X has a stable distribution, then so does $c + dX$, with the same norming parameters, for every $c, d \in \mathbb{R}$ with $d \neq 0$. Stable distributions are also closed under convolution (corresponding to sums of independent variables) if the norming parameters are the same.

Suppose that X and Y are independent variables. Assume also that X has a stable distribution with centering parameters $a_n \in \mathbb{R}$ and norming parameters $b_n \in (0, \infty)$ for $n \in \mathbb{N}_+$, and that Y has a stable distribution with centering parameters $c_n \in \mathbb{R}$ and the same norming parameters b_n for $n \in \mathbb{N}_+$. Then $Z = X + Y$ has a stable distribution with centering parameters $a_n + c_n$ and norming parameters b_n for $n \in \mathbb{N}_+$.

Proof

Suppose that $n \in \mathbb{N}_+$ and that (Z_1, X_2, \dots, Z_n) is a sequence of independent variables, each with the same distribution as Z . Then $\sum_{i=1}^n Z_i$ has the same distribution as $\sum_{i=1}^n X_i + \sum_{i=1}^n Y_i$ where $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent variables, each with the same distribution as X , and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ is a sequence of independent variables, each with the same distribution as Y , and where \mathbf{X} and \mathbf{Y} are independent. By stability, this is the same as the distribution of $(a_n + b_n X) + (c_n + b_n Y) = (a_n + c_n) + b_n(X + Y)$.

We can now give another characterization of stability that just involves two independent copies of X .

Random variable X has a stable distribution if and only if the following condition holds: If X_1, X_2 are independent variables, each with the same distribution as X and $d_1, d_2 \in (0, \infty)$ then $d_1 X_1 + d_2 X_2$ has the same distribution as $a + bX$ for some $a \in \mathbb{R}$ and $b \in (0, \infty)$.

Proof

Suppose that the condition in the theorem holds. We will show by induction that the condition in the [definition](#) holds. For $n = 2$, the stability condition is a special case of the condition in the theorem, with $d_1 = d_2 = 1$. Suppose that the stability condition holds for a given $n \in \mathbb{N}_+$. Suppose that $(X_1, X_2, \dots, X_n, X_{n+1})$ is a sequence of independent random variables, each with the distribution of X . By the induction hypothesis, $Y_n = X_1 + X_2 + \dots + X_n$ has the same distribution as $a_n + b_n X$ for some $a_n \in \mathbb{R}$ and $b_n \in (0, \infty)$. By independence, $Y_{n+1} = X_1 + X_2 + \dots + X_n + X_{n+1}$ has the same distribution as $a_n + b_n X_1 + X_{n+1}$. By another application of the condition above, $b_n X_1 + X_{n+1}$ has the same distribution as $c + b_{n+1} X$ for some $c \in \mathbb{R}$ and $b_{n+1} \in (0, \infty)$. But then Y_{n+1} has the same distribution as $(a_n + c) + b_{n+1} X$.

As a corollary of a couple of the results above, we have the following:

Suppose that X and Y are independent with the same stable distribution. Then the distribution of $X - Y$ is strictly stable, with the same norming parameters.

Note that the distribution of $X - Y$ is symmetric (about 0). The last result is useful because it allows us to get rid of the centering parameters when proving facts about the norming parameters. Here is the most important of those facts:

Suppose that X has a stable distribution. Then the norming parameters have the form $b_n = n^{1/\alpha}$ for $n \in \mathbb{N}_+$, for some $\alpha \in (0, 2]$. The parameter α is known as the *index* or *characteristic exponent* of the distribution.

Proof

The proof is in several steps, and is based on the proof in [An Introduction to Probability Theory and Its Applications](#), Volume II, by William Feller. The proof uses the basic trick of writing a sum of independent copies of X in different ways in order to obtain relationships between the norming constants b_n .

First we can assume from our last result that the distribution of X is symmetric and strictly stable. Let (X_1, X_2, \dots) be a sequence of independent variables, each with the distribution of X . Let $Y_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}_+$. Now let $n, m \in \mathbb{N}_+$ and consider Y_{mn} . Directly from stability, Y_{mn} has the same distribution as $b_{mn}X$. On the other hand, Y_{mn} can be thought of as a sum of m “blocks”, where each block is a sum of n independent copies of X . Each block has the same distribution as b_nX , and since the blocks are independent, it follows that Y_{mn} has the same distribution as

$$b_nX_1 + b_nX_2 + \dots + b_nX_m = b_n(X_1 + X_2 + \dots + X_m) \quad (5.3.1)$$

But by another application of stability, the random variable on the right has the same distribution as b_nb_mX . It then follows that $b_{mn} = b_nb_m$ for all $m, n \in \mathbb{N}_+$ which in turn leads to $b_{n^k} = b_n^k$ for all $n, k \in \mathbb{N}_+$.

We use the same trick again, this time with a sum. Let $m, n \in \mathbb{N}_+$ and consider Y_{m+n} . Directly from stability, Y_{m+n} has the same distribution as $b_{m+n}X$. On the other hand, Y_{m+n} can be thought of as the sum of two blocks. The first is the sum of m independent copies of X and hence has the same distribution as b_mX , while the second is the sum of n independent copies of X and hence has the same distribution as b_nX . Since the blocks are independent, it follows that $b_{m+n}X$ has the same distribution as $b_mX_1 + b_nX_2$, or equivalently, X has the same distribution as

$$U = \frac{b_m}{b_{m+n}}X_1 + \frac{b_n}{b_{m+n}}X_2 \quad (5.3.2)$$

Next note that for $x > 0$,

$$\left\{ X_1 \geq 0, X_2 > \frac{b_{m+n}}{b_n}x \right\} \subseteq \{U > x\} \quad (5.3.3)$$

and so by independence,

$$\mathbb{P}(U > x) \geq \mathbb{P}\left(X_1 \geq 0, X_2 > \frac{b_{m+n}}{b_n}x\right) = \mathbb{P}(X_1 \geq 0)\mathbb{P}\left(X_2 > \frac{b_{m+n}}{b_n}x\right) \quad (5.3.4)$$

But by symmetry, $\mathbb{P}(X_1 \geq 0) \geq \frac{1}{2}$. Also X_2 and U have the same distribution as X , so we conclude that

$$\mathbb{P}(X > x) \geq \frac{1}{2}\mathbb{P}\left(X > \frac{b_{m+n}}{b_n}x\right), \quad x > 0 \quad (5.3.5)$$

It follows that the ratios b_n/b_{m+n} are bounded for $m, n \in \mathbb{N}_+$. If that were not the case, we could find a sequence of integers m, n with $b_{m+n}/b_n \rightarrow 0$, in which case the displayed equation above would give the contradiction $\mathbb{P}(X > x) \geq \frac{1}{4}$ for all $x > 0$. Restating, the ratios b_k/b_n are bounded for $k, n \in \mathbb{N}_+$ with $k < n$.

Fix $r \in \mathbb{N}_+$. There exists a unique $\alpha \in (0, \infty)$ with $b_r = r^{1/\alpha}$. It then follows from step 1 above that $b_n = n^{1/\alpha}$ for every $n = r^j$ with $j \in \mathbb{N}_+$. Similarly, if $s \in \mathbb{N}_+$, there exists $\beta \in (0, \infty)$ with $b_s = s^{1/\beta}$ and then $b_m = m^{1/\beta}$ for every $m = s^k$ with $k \in \mathbb{N}_+$. For our next step, we show that $\alpha = \beta$ and it then follows that $b_n = n^{1/\alpha}$ for every $n \in \mathbb{N}_+$. Towards that end, note that if $m = s^k$ with $k \in \mathbb{N}_+$ there exists $n = r^j$ with $j \in \mathbb{N}_+$ with $n \leq m \leq rn$. Hence

$$b_m = m^{1/\beta} \leq (rn)^{1/\beta} = r^{1/\beta}b_n^{\alpha/\beta} \quad (5.3.6)$$

Therefore

$$\frac{b_m}{b_n} \leq r^{1/\beta}b_n^{\alpha/\beta-1} \quad (5.3.7)$$

Since the coefficients b_n are unbounded in $n \in \mathbb{N}_+$, but the ratios b_n/b_m are bounded for $m, n \in \mathbb{N}_+$ with $m > n$, the last inequality implies that $\beta \leq \alpha$. Reversing the roles of m and n then gives $\alpha \leq \beta$ and hence $\alpha = \beta$.

All that remains to show is that $\alpha \leq 2$. We will do this by showing that if $\alpha > 2$, then X must have finite variance, in which case the [finite variance property](#) above leads to the contradiction $\alpha = 2$. Since X^2 is nonnegative,

$$\mathbb{E}(X^2) = \int_0^\infty \mathbb{P}(X^2 > x) dx = \int_0^\infty \mathbb{P}(|X| > \sqrt{x}) dx = \sum_{k=1}^\infty \int_{2^{k-1}}^{2^k} \mathbb{P}(|X| > \sqrt{x}) dx \quad (5.3.8)$$

So the idea is to find bounds on the integrals on the right so that the sum converges. Towards that end, note that for $t > 0$ and $n \in \mathbb{N}_+$

$$\mathbb{P}(|Y_n| > tb_n) = \mathbb{P}(b_n |X| > tb_n) = \mathbb{P}(|X| > t) \quad (5.3.9)$$

Hence we can choose t so that $\mathbb{P}(|Y_n| > tb_n) \leq \frac{1}{4}$. On the other hand, using a special inequality for symmetric distributions,

$$\frac{1}{2}(1 - \exp[-n\mathbb{P}(|X| > tb_n)]) \leq \mathbb{P}(|Y_n| > tb_n) \quad (5.3.10)$$

This implies that $n\mathbb{P}(|X| > tb_n)$ is bounded in n or otherwise the two inequalities together would lead to $\frac{1}{2} \leq \frac{1}{4}$. Substituting $x = tb_n = tn^{1/\alpha}$ leads to $\mathbb{P}(|X| > x) \leq Mx^{-\alpha}$ for some $M > 0$. It then follows that

$$\int_{2^{k-1}}^{2^k} \mathbb{P}(|X| > \sqrt{x}) dx \leq M2^{k(1-\alpha/2)} \quad (5.3.11)$$

If $\alpha > 2$, the series with the terms on the right converges and we have $\mathbb{E}(X^2) < \infty$.

Every stable distribution is continuous.

Proof

As in the proof of the previous theorem, suppose that X has a symmetric stable distribution with norming parameters b_n for $n \in \mathbb{N}_+$. As a special case of the last proof, for $n \in \mathbb{N}_+$, X has the same distribution as

$$\frac{1}{b_{n+1}}X_1 + \frac{b_n}{b_{n+1}}X_2 \quad (5.3.12)$$

where X_1 and X_2 are independent and also have this distribution. Suppose now that $\mathbb{P}(X = x) = p$ for some $x \neq 0$ where $p > 0$. Then

$$\mathbb{P}\left(X = \frac{1+b_n}{b_{1+n}}x\right) \geq \mathbb{P}(X_1 = x)\mathbb{P}(X_2 = x) = p^2 > 0 \quad (5.3.13)$$

If the index $\alpha \neq 1$, the points

$$\frac{(1+b_n)}{b_{1+n}}x = \frac{1+n^{1/\alpha}}{(1+n)^{1/\alpha}}x, \quad n \in \mathbb{N}_+ \quad (5.3.14)$$

are distinct, which gives us infinitely many atoms, each with probability at least p^2 —clearly a contradiction.

Next, suppose that the only atom is $x = 0$ and that $\mathbb{P}(X = 0) = p$ where $p \in (0, 1)$. Then $X_1 + X_2$ has the same distribution as b_2X . But $\mathbb{P}(X_1 + X_2 = 0) = \mathbb{P}(X_1 = 0)\mathbb{P}(X_2 = 0) = p^2$ while $\mathbb{P}(b_2X = 0) = \mathbb{P}(X = 0) = p$, another contradiction.

The next result is a precise statement of the limit theorem alluded to in the introductory paragraph.

Suppose that (X_1, X_2, \dots) is a sequence of independent, identically distributed random variables, and let $Y_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}_+$. If there exist constants $a_n \in \mathbb{R}$ and $b_n \in (0, \infty)$ for $n \in \mathbb{N}_+$ such that $(Y_n - a_n)/b_n$ has a (non-degenerate) limiting distribution as $n \rightarrow \infty$, then the limiting distribution is stable.

The following theorem completely characterizes stable distributions in terms of the characteristic function.

Suppose that X has a stable distribution. The characteristic function of X has the following form, for some $\alpha \in (0, 2]$, $\beta \in [-1, 1]$, $c \in \mathbb{R}$, and $d \in (0, \infty)$

$$\chi(t) = \mathbb{E}(e^{itX}) = \exp(itc - d^\alpha |t|^\alpha [1 + i\beta \operatorname{sgn}(t)u_\alpha(t)]), \quad t \in \mathbb{R} \quad (5.3.15)$$

where sgn is the usual sign function, and where

$$u_\alpha(t) = \begin{cases} \tan(\frac{\pi\alpha}{2}), & \alpha \neq 1 \\ \frac{2}{\pi} \ln(|t|), & \alpha = 1 \end{cases} \quad (5.3.16)$$

1. The parameter α is the *index*, as before.
2. The parameter β is the *skewness parameter*.
3. The parameter c is the *location parameter*.
4. The parameter d is the *scale parameter*.

Thus, the family of stable distributions is a 4 parameter family. The index parameter α and the skewness parameter β can be considered *shape parameters*. When the location parameter $c = 0$ and the scale parameter $d = 1$, we get the *standard form* of the stable distributions, with characteristic function

$$\chi(t) = \mathbb{E}(e^{itX}) = \exp(-|t|^\alpha [1 + i\beta \operatorname{sgn}(t)u_\alpha(t)]), \quad t \in \mathbb{R} \quad (5.3.17)$$

The characteristic function gives another proof that stable distributions are closed under convolution (corresponding to sums of independent variables), if the index is fixed.

Suppose that X_1 and X_2 are independent random variables, and that X_1 and X_2 have the stable distribution with common index $\alpha \in (0, 2]$, skewness parameter $\beta_k \in [-1, 1]$, location parameter $c_k \in \mathbb{R}$, and scale parameter $d_k \in (0, \infty)$. Then $X_1 + X_2$ has the stable distribution with index α , location parameter $c = c_1 + c_2$, scale parameter $d = (d_1^\alpha + d_2^\alpha)^{1/\alpha}$, and skewness parameter

$$\beta = \frac{\beta_1 d_1^\alpha + \beta_2 d_2^\alpha}{d_1^\alpha + d_2^\alpha} \quad (5.3.18)$$

Proof

Let χ_k denote the characteristic function of X_k for $k \in \{1, 2\}$. Then $X_1 + X_2$ has characteristic function $\chi = \chi_1 \chi_2$. The result follows from using the form of the characteristic function [above](#) and some algebra.

Special Cases

Three special parametric families of distributions studied in this chapter are stable. In the proofs in this subsection, we use the definition of stability and various important properties of the distributions. These properties, in turn, are verified in the sections devoted to the distributions. We also give proofs based on the characteristic function, which allows us to identify the skewness parameter.

The normal distribution is stable with index $\alpha = 2$. There is no skewness parameter.

Proof

If $n \in \mathbb{N}_+$ and (Z_1, Z_2, \dots, Z_n) is a sequence of independent variables, each with the standard normal distribution, then $Z_1 + Z_2 + \dots + Z_n$ has the normal distribution with mean 0 and variance n . But this is also the distribution of $\sqrt{n}Z$ where Z has the standard normal distribution. Hence the standard normal distribution is strictly stable, with index $\alpha = 2$. The normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$ is the distribution of $\mu + \sigma Z$. From our basic properties above, this distribution is stable with index $\alpha = 2$ and centering parameters $(n - \sqrt{n})\mu$ for $n \in \mathbb{N}$.

In terms of the characteristic function, note that if $\alpha = 2$ then $u_\alpha(t) = \tan(\pi) = 0$ so the skewness parameter β drops out completely. The characteristic function in standard form $\chi(t) = e^{-t^2}$ for $t \in \mathbb{R}$, which is the characteristic function of the normal distribution with mean 0 and variance 2.

Of course, the normal distribution has finite variance, so once we know that it is stable, it follows from the [finite variance property](#) above that the index must be 2. Moreover, the characteristic function shows that the normal distribution is the only stable distribution with index 2, and hence the only stable distribution with finite variance.

Open the special distribution simulator and select the normal distribution. Vary the parameters and note the shape and location of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The Cauchy distribution is stable with index $\alpha = 1$ and skewness parameter $\beta = 0$.

Proof

If $n \in \mathbb{N}_+$ and (Z_1, Z_2, \dots, Z_n) is a sequence of independent variables, each with the standard Cauchy distribution, then $Z_1 + Z_2 + \dots + Z_n$ has the Cauchy distribution scale parameter n . By definition this is the same as the distribution of nZ where Z has the standard Cauchy distribution. Hence the standard Cauchy distribution is strictly stable, with index $\alpha = 1$. The Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$ is the distribution of $a + bZ$. From our basic properties above, this distribution is strictly stable with index $\alpha = 1$.

When $\alpha = 1$ and $\beta = 0$ the characteristic function in standard form is $\chi(t) = \exp(-|t|)$ for $t \in \mathbb{R}$, which is the characteristic function of the standard Cauchy distribution.

Open the special distribution simulator and select the Cauchy distribution. Vary the parameters and note the shape and location of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The Lévy distribution is stable with index $\alpha = \frac{1}{2}$ and skewness parameter $\beta = 1$.

Proof

If $n \in \mathbb{N}_+$ and (Z_1, Z_2, \dots, Z_n) is a sequence of independent variables, each with the standard Lévy distribution, then $Z_1 + Z_2 + \dots + Z_n$ has the Lévy distribution scale parameter n^2 . By definition this is the same as the distribution of n^2Z where Z has the standard Lévy distribution. Hence the standard Lévy distribution is strictly stable, with index $\alpha = \frac{1}{2}$. The Lévy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$ is the distribution of $a + bZ$. From our basic properties above, this distribution is stable with index $\alpha = \frac{1}{2}$ and centering parameters $(n - n^2)a$ for $n \in \mathbb{N}_+$.

When $\alpha = \frac{1}{2}$ note that $u_\alpha(t) = \tan\left(\frac{\pi}{4}\right) = 1$. So the characteristic function in standard form with $\alpha = \frac{1}{2}$ and $\beta = 1$ is

$$\chi(t) = \exp\left(-|t|^{1/2} [1 + i \operatorname{sgn}(t)]\right) \quad (5.3.19)$$

which is the characteristic function of the standard Lévy distribution.

Open the special distribution simulator and select the Lévy distribution. Vary the parameters and note the shape and location of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The normal, Cauchy, and Lévy distributions are the only stable distributions for which the probability density function is known in closed form.

This page titled [5.3: Stable Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.4: Infinitely Divisible Distributions

This section discusses a theoretical topic that you may want to skip if you are a new student of probability.

Basic Theory

Infinitely divisible distributions form an important class of distributions on \mathbb{R} that includes the stable distributions, the compound Poisson distributions, as well as several of the most important special parametric families of distributions. Basically, the distribution of a real-valued random variable is infinitely divisible if for each $n \in \mathbb{N}_+$, the variable can be decomposed into the sum of n independent copies of another variable. Here is the precise definition.

The distribution of a real-valued random variable X is *infinitely divisible* if for every $n \in \mathbb{N}_+$, there exists a sequence of independent, identically distributed variables (X_1, X_2, \dots, X_n) such that $X_1 + X_2 + \dots + X_n$ has the same distribution as X .

If the distribution of X is stable then the distribution is infinitely divisible.

Proof

Let $n \in \mathbb{N}_+$ and let (X_1, X_2, \dots, X_n) be a sequence of independent variables, each with the same distribution as X . By the definition of stability, there exists $a_n \in \mathbb{R}$ and $b_n \in (0, \infty)$ such that $\sum_{i=1}^n X_i$ has the same distribution as $a_n + b_n X$. But then

$$\frac{1}{b_n} \left(\sum_{i=1}^n X_i - a_n \right) = \sum_{i=1}^n \frac{X_i - a_n/n}{b_n} \quad (5.4.1)$$

has the same distribution as X . But $\left(\frac{X_i - a_n/n}{b_n} : i \in \{1, 2, \dots, n\} \right)$ is an IID sequence, and hence the distribution of X is infinitely divisible.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent, identically distributed random variables, and that N has a Poisson distribution and is independent of \mathbf{X} . Recall that the distribution of $\sum_{i=1}^N X_i$ is said to be compound Poisson. Like the stable distributions, the compound Poisson distributions form another important class of infinitely divisible distributions.

Suppose that Y is a random variable.

1. If Y is compound Poisson then Y is infinitely divisible.
2. If Y is infinitely divisible and takes values in \mathbb{N} then Y is compound Poisson.

Proof

1. Suppose that Y is compound Poisson, so that we can take $Y = \sum_{i=1}^N X_i$ where $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent, identically distributed random variables with common characteristic function ϕ , and where N is independent of \mathbf{X} and has the Poisson distribution with parameter $\lambda \in (0, \infty)$. The characteristic function χ of Y is given by $\chi(t) = \exp(\lambda[\phi(t) - 1])$ for $t \in \mathbb{R}$. But then for $n \in \mathbb{N}_+$,

$$\chi(t) = \left[\exp\left(\frac{\lambda}{n}[\phi(t) - 1]\right) \right]^n, \quad t \in \mathbb{R} \quad (5.4.2)$$

But $t \mapsto \exp\left(\frac{\lambda}{n}[\phi(t) - 1]\right)$ is the characteristic function of the compound Poisson distribution corresponding to \mathbf{X} but with Poisson parameter λ/n . Restated in terms of random variables, $Y = \sum_{i=1}^n Y_i$ where Y_i has the compound Poisson distribution corresponding to \mathbf{X} with Poisson parameter λ/n .

2. The proof is from [An Introduction to Probability Theory and Its Applications](#) by William Feller, and requires some additional notation. Recall that the symbol \asymp is used to connect functions that are *asymptotically the same* in the sense that the ratio converges to 1. Suppose now that Y takes values in \mathbb{N} and is infinitely divisible. In this case we can use probability generating functions rather than characteristic functions, so let P denote the PGF of Y . By definition, $P(t) = \sum_{k=0}^{\infty} p_k t^k$ where $p_k = \mathbb{P}(Y = k)$ for $k \in \mathbb{N}$. Since Y is infinitely divisible, $P^{1/n}$ is also a PGF for every

$n \in \mathbb{N}_+$, so let $P^{1/n}(t) = \sum_{k=0}^{\infty} p_{nk} t^k$ where $p_{nk} \geq 0$ for $k \in \mathbb{N}$ and $n \in \mathbb{N}_+$ and $\sum_{k=0}^{\infty} p_{nk} = 1$ for $n \in \mathbb{N}_+$. As with all PGFs, the series for $P(t)$ and for $P^{1/n}(t)$ converge at least for $t \in [0, 1]$, and this interval is sufficient for a PGF to completely determine the underlyingly distribution. For $n \in \mathbb{N}_+$, we have

$$\sum_{k=0}^{\infty} p_k t^k = \left(\sum_{k=0}^{\infty} p_{nk} t^k \right)^n \quad (5.4.3)$$

Expanding the series on the right and then equating coefficients of the two series term by term, we see that if $p_0 = 0$ then $p_{n0} = 0$ which in turn would imply $p_1 = \dots = p_{n-1} = 0$. Since this is true for all $n \in \mathbb{N}_+$, we would have P identically 0, which is a contradiction. Hence $p_0 > 0$ and so $P(t) > 0$ for $t \in [0, 1]$ and therefore $[P(t)/p_0]^{1/n} \rightarrow 1$ as $n \rightarrow \infty$ for $t \in [0, 1]$. Next recall that $\ln(1+x) \asymp x$ as $x \downarrow 0$. It follows that for $t \in [0, 1]$,

$$\ln \left(\left[\frac{P(t)}{p_0} \right]^{1/n} \right) = \ln \left\{ 1 + \left(\left[\frac{P(t)}{p_0} \right]^{1/n} - 1 \right) \right\} \asymp \left[\frac{P(t)}{p_0} \right]^{1/n} - 1 \text{ as } n \rightarrow \infty \quad (5.4.4)$$

As a special case, when $t = 1$, we have $\ln \left[(1/p_0)^{1/n} \right] \asymp (1/p_0)^{1/n} - 1$ as $n \rightarrow \infty$. Hence using properties of logarithms and a bit of algebra,

$$\frac{\ln[P(t)] - \ln(p_0)}{-\ln(p_0)} = \frac{\ln \left([P(t)/p_0]^{1/n} \right)}{\ln \left[(1/p_0)^{1/n} \right]} \asymp \frac{P^{1/n}(t) - p_0^{1/n}}{1 - p_0^{1/n}} \text{ as } n \rightarrow \infty \quad (5.4.5)$$

The power series (about 0) for the expression on the right has positive coefficients, and the expression takes the value 1 when $t = 1$. Thus, the expression on the right is a PGF for each $n \in \mathbb{N}_+$. By the continuity theorem for convergence in distribution, it follows that the left side, which we will denote by $Q(t)$, is also a PGF. Solving we have

$$P(t) = \exp(\lambda[Q(t) - 1]), \quad t \in [0, 1] \quad (5.4.6)$$

where $\lambda = -\ln(p_0)$. This is the PGF of the distribution obtained by compounding the distribution with PGF Q with the Poisson distribution with parameter λ .

Special Cases

A number of special distributions are infinitely divisible. Proofs of the results stated below are given in the individual sections.

Stable Distributions

First, the normal distribution, the Cauchy distribution, and the Lévy distribution are stable, so they are infinitely divisible. However, direct arguments give more information, because we can identify the distribution of the component variables.

The normal distribution is infinitely divisible. If X has the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \dots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the normal distribution with mean μ/n and standard deviation σ/\sqrt{n} for each $i \in \{1, 2, \dots, n\}$.

The Cauchy distribution is infinitely divisible. If X has the Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \dots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the Cauchy distribution with location parameter a/n and scale parameter b/n for each $i \in \{1, 2, \dots, n\}$.

Other Special Distributions

On the other hand, there are distributions that are infinitely divisible but not stable.

The gamma distribution is infinitely divisible. If X has the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \dots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the gamma distribution with shape parameter k/n and scale parameter b for each $i \in \{1, 2, \dots, n\}$.

The chi-square distribution is infinitely divisible. If X has the chi-square distribution with $k \in (0, \infty)$ degrees of freedom, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \cdots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the chi-square distribution with k/n degrees of freedom for each $i \in \{1, 2, \dots, n\}$.

The Poisson distribution is infinitely divisible. If X has the Poisson distribution with rate parameter $\lambda \in (0, \infty)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \cdots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the Poisson distribution with rate parameter λ/n for each $i \in \{1, 2, \dots, n\}$.

The general negative binomial distribution on \mathbb{N} is infinitely divisible. If X has the negative binomial distribution on \mathbb{N} with parameters $k \in (0, \infty)$ and $p \in (0, 1)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \cdots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the negative binomial distribution on \mathbb{N} with parameters k/n and p for each $i \in \{1, 2, \dots, n\}$.

Since the Poisson distribution and the negative binomial distributions are distributions on \mathbb{N} , it follows from the [characterization above](#) that these distributions must be compound Poisson. Of course it is completely trivial that the Poisson distribution is compound Poisson, but it's far from obvious that the negative binomial distribution has this property. It turns out that the negative binomial distribution can be obtained by compounding the logarithmic series distribution with the Poisson distribution.

The Wald distribution is infinitely divisible. If X has the Wald distribution with shape parameter $\lambda \in (0, \infty)$ and mean $\mu \in (0, \infty)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \cdots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and X_i has the Wald distribution with shape parameter λ/n^2 and mean μ/n for each $i \in \{1, 2, \dots, n\}$.

This page titled [5.4: Infinitely Divisible Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.5: Power Series Distributions

Power Series Distributions are discrete distributions on (a subset of) \mathbb{N} constructed from power series. This class of distributions is important because most of the special, discrete distributions are power series distributions.

Basic Theory

Power Series

Suppose that $\mathbf{a} = (a_0, a_1, a_2, \dots)$ is a sequence of nonnegative real numbers. We are interested in the power series with \mathbf{a} as the sequence of coefficients. Recall first that the *partial sum* of order $n \in \mathbb{N}$ is

$$g_n(\theta) = \sum_{k=0}^n a_k \theta^k, \quad \theta \in \mathbb{R} \quad (5.5.1)$$

The *power series* g is then defined by $g(\theta) = \lim_{n \rightarrow \infty} g_n(\theta)$ for $\theta \in \mathbb{R}$ for which the limit exists, and is denoted

$$g(\theta) = \sum_{n=0}^{\infty} a_n \theta^n \quad (5.5.2)$$

Note that the series converges when $\theta = 0$, and $g(0) = a_0$. Beyond this trivial case, recall that there exists $r \in [0, \infty]$ such that the series converges absolutely for $|\theta| < r$ and diverges for $|\theta| > r$. The number r is the *radius of convergence*. From now on, we assume that $r > 0$. If $r < \infty$, the series may converge (absolutely) or may diverge to ∞ at the endpoint r . At $-r$, the series may converge absolutely, may converge conditionally, or may diverge.

Distributions

From now on, we restrict θ to the interval $[0, r)$; this interval is our *parameter space*. Some of the results below may hold when $r < \infty$ and $\theta = r$, but dealing with this case explicitly makes the exposition unnecessarily cumbersome.

Suppose that N is a random variable with values in \mathbb{N} . Then N has the *power series distribution* associated with the function g (or equivalently with the sequence \mathbf{a}) and with parameter $\theta \in [0, r)$ if N has probability density function f_θ given by

$$f_\theta(n) = \frac{a_n \theta^n}{g(\theta)}, \quad n \in \mathbb{N} \quad (5.5.3)$$

Proof

To show that f_θ is a valid discrete probability density function, note that $a_n \theta^n$ is nonnegative for each $n \in \mathbb{N}$ and $g(\theta)$, by definition, is the normalizing constant for the sequence $(a_n \theta^n : n \in \mathbb{N})$.

Note that when $\theta = 0$, the distribution is simply the point mass distribution at 0; that is, $f_0(0) = 1$.

The distribution function F_θ is given by

$$F_\theta(n) = \frac{g_n(\theta)}{g(\theta)}, \quad n \in \mathbb{N} \quad (5.5.4)$$

Proof

This follows immediately from the definitions since $F_\theta(n) = \sum_{k=0}^n f_\theta(k)$ for $n \in \mathbb{N}$

Of course, the probability density function f_θ is most useful when the power series $g(\theta)$ can be given in closed form, and similarly the distribution function F_θ is most useful when the power series and the partial sums can be given in closed form

Moments

The moments of N can be expressed in terms of the underlying power series function g , and the nicest expression is for the factorial moments. Recall that the permutation formula is $t^{(k)} = t(t-1) \cdots (t-k+1)$ for $t \in \mathbb{R}$ and $k \in \mathbb{N}$, and the *factorial moment* of N of order $k \in \mathbb{N}$ is $\mathbb{E}(N^{(k)})$.

For $\theta \in [0, r)$, the factorial moments of N are as follows, where $g^{(k)}$ is the k th derivative of g .

$$\mathbb{E}(N^{(k)}) = \frac{\theta^k g^{(k)}(\theta)}{g(\theta)}, \quad k \in \mathbb{N} \quad (5.5.5)$$

Proof

Recall that a power series is infinitely differentiable in the open interval of convergence, and that the derivatives can be taken term by term. Thus

$$\mathbb{E}(N^{(k)}) = \sum_{n=0}^{\infty} n^{(k)} \frac{a_n \theta^n}{g(\theta)} = \frac{\theta^k}{g(\theta)} \sum_{n=k}^{\infty} a_n n^{(k)} \theta^{n-k} = \frac{\theta^k}{g(\theta)} g^{(k)}(\theta) \quad (5.5.6)$$

The mean and variance of N are

1. $\mathbb{E}(N) = \theta g'(\theta) / g(\theta)$
2. $\text{var}(N) = \theta^2 (g''(\theta) / g(\theta) - [g'(\theta) / g(\theta)]^2)$

Proof

1. This follows from the previous result on factorial moments with $k = 1$.
2. This also follows from the previous result since $\text{var}(N) = \mathbb{E}(N^{(2)}) + \mathbb{E}(N)[1 - \mathbb{E}(N)]$.

The probability generating function of N also has a simple expression in terms of g .

For $\theta \in (0, r)$, the probability generating function P of N is given by

$$P(t) = \mathbb{E}(t^N) = \frac{g(\theta t)}{g(\theta)}, \quad t < \frac{r}{\theta} \quad (5.5.7)$$

Proof

For $t \in (0, r/\theta)$,

$$P(t) = \sum_{n=0}^{\infty} t^n f_{\theta}(n) = \frac{1}{g(\theta)} \sum_{n=0}^{\infty} a_n (t\theta)^n = \frac{g(t\theta)}{g(\theta)} \quad (5.5.8)$$

Relations

Power series distributions are closed with respect to sums of independent variables.

Suppose that N_1 and N_2 are independent, and have power series distributions relative to the functions g_1 and g_2 , respectively, each with parameter value $\theta < \min\{r_1, r_2\}$. Then $N_1 + N_2$ has the power series distribution relative to the function $g_1 g_2$, with parameter value θ .

Proof

A direct proof is possible, but there is an easy proof using probability generating functions. Recall that the PGF of the sum of independent variables is the product of the PGFs. Hence the PGF of $N_1 + N_2$ is

$$P(t) = P_1(t)P_2(t) = \frac{g_1(\theta t)}{g_1(\theta)} \frac{g_2(\theta t)}{g_2(\theta)} = \frac{g_1(\theta t)g_2(\theta t)}{g_1(\theta)g_2(\theta)}, \quad t < \min\left\{\frac{r_1}{\theta}, \frac{r_2}{\theta}\right\} \quad (5.5.9)$$

The last expression is the PGF of the power series distribution relative to the function $g_1 g_2$, at θ .

Here is a simple corollary.

Suppose that (N_1, N_2, \dots, N_k) is a sequence of independent variables, each with the same power series distribution, relative to the function g and with parameter value $\theta < r$. Then $N_1 + N_2 + \dots + N_k$ has the power series distribution relative to the function g^k and with parameter θ .

In the context of this result, recall that (N_1, N_2, \dots, N_k) is a *random sample* of size k from the common distribution.

Examples and Special Cases

Special Distributions

The Poisson distribution with rate parameter $\lambda \in [0, \infty)$ is a power series distribution relative to the function $g(\lambda) = e^\lambda$ for $\lambda \in [0, \infty)$.

Proof

This follows directly from the definition, since the PDF of the Poisson distribution with parameter λ is $f(n) = e^{-\lambda} \lambda^n / n!$ for $n \in \mathbb{N}$.

The geometric distribution on \mathbb{N} with success parameter $p \in (0, 1]$ is a power series distribution relative to the function $g(\theta) = 1/(1 - \theta)$ for $\theta \in [0, 1)$, where $\theta = 1 - p$.

Proof

This follows directly from the definition, since the PDF of the geometric distribution on \mathbb{N} is $f(n) = (1 - p)^n p = (1 - \theta)^n \theta$ for $n \in \mathbb{N}$.

For fixed $k \in (0, \infty)$, the negative binomial distribution on \mathbb{N} with stopping parameter k and success parameter $p \in (0, 1]$ is a power series distribution relative to the function $g(\theta) = 1/(1 - \theta)^k$ for $\theta \in [0, 1)$, where $\theta = 1 - p$.

Proof

This follows from the result above on [sums of IID variables](#), but can also be seen directly, since the PDF is

$$f(n) = \binom{n+k-1}{k-1} p^k (1-p)^n = (1-\theta)^k \binom{n+k-1}{k-1} \theta^n, \quad n \in \mathbb{N} \quad (5.5.10)$$

For fixed $n \in \mathbb{N}_+$, the binomial distribution with trial parameter n and success parameter $p \in [0, 1]$ is a power series distribution relative to the function $g(\theta) = (1 + \theta)^n$ for $\theta \in [0, \infty)$, where $\theta = p/(1 - p)$.

Proof

Note that the PDF of the binomial distribution is

$$f(k) = \binom{n}{k} p^k (1-p)^{n-k} = (1-p)^n \binom{n}{k} \left(\frac{p}{1-p} \right)^k = \frac{1}{(1+\theta)^n} \binom{n}{k} \theta^k, \quad k \in \{0, 1, \dots, n\} \quad (5.5.11)$$

where $\theta = p/(1 - p)$. This shows that the distribution is a power series distribution corresponding to the function $g(\theta) = (1 + \theta)^n$.

The logarithmic distribution with parameter $p \in (0, 1)$ is a power series distribution relative to the function $g(p) = -\ln(1 - p)$ for $p \in [0, 1)$.

Proof

This follows directly from the definition, since the PDF is

$$f(n) = \frac{1}{-\ln(1-p)} \frac{1}{n} p^n, \quad n \in \mathbb{N} \quad (5.5.12)$$

This page titled [5.5: Power Series Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.6: The Normal Distribution

The normal distribution holds an honored role in probability and statistics, mostly because of the central limit theorem, one of the fundamental theorems that forms a bridge between the two subjects. In addition, as we will see, the normal distribution has many nice mathematical properties. The normal distribution is also called the *Gaussian distribution*, in honor of Carl Friedrich Gauss, who was among the first to use the distribution.

The Standard Normal Distribution

Distribution Functions

The *standard normal distribution* is a continuous distribution on \mathbb{R} with probability density function ϕ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad z \in \mathbb{R} \quad (5.6.1)$$

Proof that ϕ is a probability density function

Let $c = \int_{-\infty}^{\infty} e^{-z^2/2} dz$. We need to show that $c = \sqrt{2\pi}$. That is, $\sqrt{2\pi}$ is the normalizing constant for the function $z \mapsto e^{-z^2/2}$. The proof uses a nice trick:

$$c^2 = \int_{-\infty}^{\infty} e^{-x^2/2} dx \int_{-\infty}^{\infty} e^{-y^2/2} dy = \int_{\mathbb{R}^2} e^{-(x^2+y^2)/2} d(x, y) \quad (5.6.2)$$

We now convert the double integral to polar coordinates: $x = r \cos \theta$, $y = r \sin \theta$ where $r \in [0, \infty)$ and $\theta \in [0, 2\pi)$. So, $x^2 + y^2 = r^2$ and $d(x, y) = r d(r, \theta)$. Thus, converting back to iterated integrals,

$$c^2 = \int_0^{2\pi} \int_0^{\infty} r e^{-r^2/2} dr d\theta \quad (5.6.3)$$

Substituting $u = r^2/2$ in the inner integral gives $\int_0^{\infty} e^{-u} du = 1$ and then the outer integral is $\int_0^{2\pi} 1 d\theta = 2\pi$. Thus, $c^2 = 2\pi$ and so $c = \sqrt{2\pi}$.

The standard normal probability density function has the famous “bell shape” that is known to just about everyone.

The standard normal density function ϕ satisfies the following properties:

1. ϕ is symmetric about $z = 0$.
2. ϕ increases and then decreases, with mode $z = 0$.
3. ϕ is concave upward and then downward and then upward again, with inflection points at $z = \pm 1$.
4. $\phi(z) \rightarrow 0$ as $z \rightarrow \infty$ and as $z \rightarrow -\infty$.

Proof

These results follow from standard calculus. Note that $\phi'(z) = -z\phi(z)$ (which gives (b)) and hence also $\phi''(z) = (z^2 - 1)\phi(z)$ (which gives (c)).

In the Special Distribution Simulator, select the normal distribution and keep the default settings. Note the shape and location of the standard normal density function. Run the simulation 1000 times, and compare the empirical density function to the probability density function.

The standard normal distribution function Φ , given by

$$\Phi(z) = \int_{-\infty}^z \phi(t) dt = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt \quad (5.6.4)$$

and its inverse, the quantile function Φ^{-1} , cannot be expressed in closed form in terms of elementary functions. However approximate values of these functions can be obtained from the special distribution calculator, and from most mathematics and statistics software. Indeed these functions are so important that they are considered *special functions* of mathematics.

The standard normal distribution function Φ satisfies the following properties:

1. $\Phi(-z) = 1 - \Phi(z)$ for $z \in \mathbb{R}$
2. $\Phi^{-1}(p) = -\Phi^{-1}(1-p)$ for $p \in (0, 1)$
3. $\Phi(0) = \frac{1}{2}$, so the median is 0.

Proof

Part (a) follows from the symmetry of ϕ . Part (b) follows from part (a). Part (c) follows from part (a) with $z = 0$.

In the special distribution calculator, select the normal distribution and keep the default settings.

1. Note the shape of the density function and the distribution function.
2. Find the first and third quartiles.
3. Compute the interquartile range.

In the special distribution calculator, select the normal distribution and keep the default settings. Find the quantiles of the following orders for the standard normal distribution:

1. $p = 0.001, p = 0.999$
2. $p = 0.05, p = 0.95$
3. $p = 0.1, p = 0.9$

Moments

Suppose that random variable Z has the standard normal distribution.

The mean and variance of Z are

1. $\mathbb{E}(Z) = 0$
2. $\text{var}(Z) = 1$

Proof

1. Of course, by symmetry, if Z has a mean, the mean must be 0, but we have to argue that the mean exists. Actually it's not hard to compute the mean directly. Note that

$$\mathbb{E}(Z) = \int_{-\infty}^{\infty} z \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = \int_{-\infty}^0 z \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz + \int_0^{\infty} z \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz \quad (5.6.5)$$

The integrals on the right can be evaluated explicitly using the simple substitution $u = z^2/2$. The result is $\mathbb{E}(Z) = -1/\sqrt{2\pi} + 1/\sqrt{2\pi} = 0$.

2. Note that

$$\text{var}(Z) = \mathbb{E}(Z^2) = \int_{-\infty}^{\infty} z^2 \phi(z) dz \quad (5.6.6)$$

Integrate by parts, using the parts $u = z$ and $dv = z\phi(z) dz$. Thus $du = dz$ and $v = -\phi(z)$. Note that $z\phi(z) \rightarrow 0$ as $z \rightarrow \infty$ and as $z \rightarrow -\infty$. Thus, the integration by parts formula gives $\text{var}(Z) = \int_{-\infty}^{\infty} \phi(z) dz = 1$.

- 3.

In the Special Distribution Simulator, select the normal distribution and keep the default settings. Note the shape and size of the mean \pm standard deviation bar. Run the simulation 1000 times, and compare the empirical mean and standard deviation to the true mean and standard deviation.

More generally, we can compute all of the moments. The key is the following recursion formula.

For $n \in \mathbb{N}_+$, $\mathbb{E}(Z^{n+1}) = n\mathbb{E}(Z^{n-1})$

Proof

First we use the differential equation in the proof of the [PDF properties](#) above, namely $\phi'(z) = -z\phi(z)$.

$$\mathbb{E}(Z^{n+1}) = \int_{-\infty}^{\infty} z^{n+1} \phi(z) dz = \int_{-\infty}^{\infty} z^n z \phi(z) dz = - \int_{-\infty}^{\infty} z^n \phi'(z) dz \quad (5.6.7)$$

Now we integrate by parts, with $u = z^n$ and $dv = \phi'(z) dz$ to get

$$\mathbb{E}(Z^{n+1}) = -z^n \phi(z) \Big|_{-\infty}^{\infty} + \int_{-\infty}^{\infty} n z^{n-1} \phi(z) dz = 0 + n \mathbb{E}(Z^{n-1}) \quad (5.6.8)$$

The moments of the standard normal distribution are now easy to compute.

For $n \in \mathbb{N}$,

1. $\mathbb{E}(Z^{2n}) = 1 \cdot 3 \cdots (2n-1) = (2n)! / (n! 2^n)$
2. $\mathbb{E}(Z^{2n+1}) = 0$

Proof

The result follows from the [mean and variance](#) and [recursion relation](#) above.

1. Since $\mathbb{E}(Z) = 0$ it follows that $\mathbb{E}(Z^n) = 0$ for every odd $n \in \mathbb{N}$.
2. Since $\mathbb{E}(Z^2) = 1$, it follows that $\mathbb{E}(Z^4) = 1 \cdot 3$ and then $\mathbb{E}(Z^6) = 1 \cdot 3 \cdot 5$, and so forth. You can use induction, if you like, for a more formal proof.

Of course, the fact that the odd-order moments are 0 also follows from the symmetry of the distribution. The following theorem gives the skewness and kurtosis of the standard normal distribution.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = 3$

Proof

1. This follows immediately from the symmetry of the distribution. Directly, since Z has mean 0 and variance 1, $\text{skew}(Z) = \mathbb{E}(Z^3) = 0$.
2. Since $\mathbb{E}(Z) = 0$ and $\text{var}(Z) = 1$, $\text{kurt}(Z) = \mathbb{E}(Z^4) = 3$.

Because of the last result, (and the use of the standard normal distribution literally as a *standard*), the *excess kurtosis* of a random variable is defined to be the ordinary kurtosis minus 3. Thus, the excess kurtosis of the normal distribution is 0.

Many other important properties of the normal distribution are most easily obtained using the moment generating function or the characteristic function.

The moment generating function m and characteristic function χ of Z are given by

1. $m(t) = e^{t^2/2}$ for $t \in \mathbb{R}$.
2. $\chi(t) = e^{-t^2/2}$ for $t \in \mathbb{R}$.

Proof

1. Note that

$$m(t) = \mathbb{E}(e^{tZ}) = \int_{-\infty}^{\infty} e^{tz} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}z^2 + tz\right) dz \quad (5.6.9)$$

We complete the square in z to get $-\frac{1}{2}z^2 + tz = -\frac{1}{2}(z-t)^2 + \frac{1}{2}t^2$. Thus we have

$$\mathbb{E}(e^{tZ}) = e^{\frac{1}{2}t^2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(z-t)^2\right] dz \quad (5.6.10)$$

In the integral, if we use the simple substitution $u = z - t$ then the integral becomes $\int_{-\infty}^{\infty} \phi(u) du = 1$. Hence

$$\mathbb{E}(e^{tZ}) = e^{\frac{1}{2}t^2},$$

2. This follows from (a) since $\chi(t) = m(it)$.

Thus, the standard normal distribution has the curious property that the characteristic function is a multiple of the probability density function:

$$\chi = \sqrt{2\pi}\phi \quad (5.6.11)$$

The moment generating function can be used to give another derivation of the moments of Z , since we know that $\mathbb{E}(Z^n) = m^{(n)}(0)$.

The General Normal Distribution

The general normal distribution is the location-scale family associated with the standard normal distribution.

Suppose that $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ and that Z has the standard normal distribution. Then $X = \mu + \sigma Z$ has the *normal distribution* with location parameter μ and scale parameter σ .

Distribution Functions

Suppose that X has the normal distribution with location parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$. The basic properties of the density function and distribution function of X follow from general results for location scale families.

The probability density function f of X is given by

$$f(x) = \frac{1}{\sigma} \phi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (5.6.12)$$

Proof

This follows from the change of variables formula corresponding to the transformation $x = \mu + \sigma z$.

The probability density function f satisfies the following properties:

1. f is symmetric about $x = \mu$.
2. f increases and then decreases with mode $x = \mu$.
3. f is concave upward then downward then upward again, with inflection points at $x = \mu \pm \sigma$.
4. $f(x) \rightarrow 0$ as $x \rightarrow \infty$ and as $x \rightarrow -\infty$.

Proof

These properties follow from the corresponding [properties of \$\phi\$](#) .

In the special distribution simulator, select the normal distribution. Vary the parameters and note the shape and location of the probability density function. With your choice of parameter settings, run the simulation 1000 times and compare the empirical density function to the true probability density function.

Let F denote the distribution function of X , and as above, let Φ denote the standard normal distribution function.

The distribution function F and quantile function F^{-1} satisfy the following properties:

1. $F(x) = \Phi\left(\frac{x-\mu}{\sigma}\right)$ for $x \in \mathbb{R}$.
2. $F^{-1}(p) = \mu + \sigma \Phi^{-1}(p)$ for $p \in (0, 1)$.
3. $F(\mu) = \frac{1}{2}$ so the median occurs at $x = \mu$.

Proof

Part (a) follows since $X = \mu + \sigma Z$. Parts (b) and (c) follow from (a).

In the special distribution calculator, select the normal distribution. Vary the parameters and note the shape of the density function and the distribution function.

Moments

Suppose again that X has the normal distribution with location parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$. As the notation suggests, the location and scale parameters are also the mean and standard deviation, respectively.

The mean and variance of X are

1. $\mathbb{E}(X) = \mu$
2. $\text{var}(X) = \sigma^2$

Proof

This follows from the representation $X = \mu + \sigma Z$ and basic properties of expected value and variance.

So the parameters of the normal distribution are usually referred to as the *mean* and *standard deviation* rather than *location* and *scale*. The central moments of X can be computed easily from the moments of the standard normal distribution. The ordinary (raw) moments of X can be computed from the central moments, but the formulas are a bit messy.

For $n \in \mathbb{N}$,

1. $\mathbb{E}[(X - \mu)^{2n}] = 1 \cdot 3 \cdots (2n - 1)\sigma^{2n} = (2n)!\sigma^{2n} / (n!2^n)$
2. $\mathbb{E}[(X - \mu)^{2n+1}] = 0$

All of the odd central moments of X are 0, a fact that also follows from the symmetry of the probability density function.

In the special distribution simulator select the normal distribution. Vary the mean and standard deviation and note the size and location of the mean/standard deviation bar. With your choice of parameter settings, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The following exercise gives the skewness and kurtosis.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = 3$

Proof

The skewness and kurtosis of a variable are defined in terms of the standard score, so these results follows from the [corresponding result for \$Z\$](#) .

The moment generating function M and characteristic function χ of X are given by

1. $M(t) = \exp(\mu t + \frac{1}{2}\sigma^2 t^2)$ for $t \in \mathbb{R}$.
2. $\chi(t) = \exp(i\mu t - \frac{1}{2}\sigma^2 t^2)$ for $t \in \mathbb{R}$

Proof

1. This follows from the representation $X = \mu + \sigma Z$, basic properties of expected value, and the MGF of Z in (12):

$$\mathbb{E}(e^{tX}) = \mathbb{E}(e^{t\mu + t\sigma Z}) = e^{t\mu} \mathbb{E}(e^{t\sigma Z}) = e^{t\mu} e^{\frac{1}{2}t^2\sigma^2} = e^{t\mu + \frac{1}{2}\sigma^2 t^2} \quad (5.6.13)$$

2. This follows from (a) since $\chi(t) = M(it)$.

Relations

The normal family of distributions satisfies two very important properties: invariance under linear transformations of the variable and invariance with respect to sums of independent variables. The first property is essentially a restatement of the fact that the normal distribution is a location-scale family.

Suppose that X is normally distributed with mean μ and variance σ^2 . If $a \in \mathbb{R}$ and $b \in \mathbb{R} \setminus \{0\}$, then $a + bX$ is normally distributed with mean $a + b\mu$ and variance $b^2\sigma^2$.

Proof

The MGF of $a + bX$ is

$$\mathbb{E} \left[e^{t(a+bX)} \right] = e^{ta} \mathbb{E} \left[e^{(tb)X} \right] = e^{ta} e^{\mu(tb) + \sigma^2(tb)^2/2} = e^{(a+b\mu)t + b^2\sigma^2 t^2/2} \quad (5.6.14)$$

which we recognize as the MGF of the normal distribution with mean $a + b\mu$ and variance $b^2\sigma^2$.

Recall that in general, if X is a random variable with mean μ and standard deviation $\sigma > 0$, then $Z = (X - \mu)/\sigma$ is the *standard score* of X . A corollary of the last result is that if X has a normal distribution then the standard score Z has a standard normal distribution. Conversely, any normally distributed variable can be constructed from a standard normal variable.

Standard score.

1. If X has the normal distribution with mean μ and standard deviation σ then $Z = \frac{X - \mu}{\sigma}$ has the standard normal distribution.
2. If Z has the standard normal distribution and if $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, then $X = \mu + \sigma Z$ has the normal distribution with mean μ and standard deviation σ .

Suppose that X_1 and X_2 are independent random variables, and that X_i is normally distributed with mean μ_i and variance σ_i^2 for $i \in \{1, 2\}$. Then $X_1 + X_2$ is normally distributed with

1. $\mathbb{E}(X_1 + X_2) = \mu_1 + \mu_2$
2. $\text{var}(X_1 + X_2) = \sigma_1^2 + \sigma_2^2$

Proof

The MGF of $X_1 + X_2$ is the product of the MGFs, so

$$\mathbb{E}(\exp[t(X_1 + X_2)]) = \exp(\mu_1 t + \sigma_1^2 t^2/2) \exp(\mu_2 t + \sigma_2^2 t^2/2) = \exp[(\mu_1 + \mu_2)t + (\sigma_1^2 + \sigma_2^2)t^2/2] \quad (5.6.15)$$

which we recognize as the MGF of the normal distribution with mean $\mu_1 + \mu_2$ and variance $\sigma_1^2 + \sigma_2^2$.

This theorem generalizes to a sum of n independent, normal variables. The important part is that the sum is still normal; the expressions for the mean and variance are standard results that hold for the sum of independent variables generally. As a consequence of this result and the one for [linear transformations](#), it follows that the normal distribution is stable.

The normal distribution is stable. Specifically, suppose that X has the normal distribution with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$. If (X_1, X_2, \dots, X_n) are independent copies of X , then $X_1 + X_2 + \dots + X_n$ has the same distribution as $(n - \sqrt{n})\mu + \sqrt{n}X$, namely normal with mean $n\mu$ and variance $n\sigma^2$.

Proof

As a consequence of the result for [sums](#) $X_1 + X_2 + \dots + X_n$ has the normal distribution with mean $n\mu$ and variance $n\sigma^2$. As a consequence of the result for [linear transformations](#), $(n - \sqrt{n})\mu + \sqrt{n}X$ has the normal distribution with mean $(n - \sqrt{n})\mu + \sqrt{n}\mu = n\mu$ and variance $(\sqrt{n})^2\sigma^2 = n\sigma^2$.

All stable distributions are infinitely divisible, so the normal distribution belongs to this family as well. For completeness, here is the explicit statement:

The normal distribution is infinitely divisible. Specifically, if X has the normal distribution with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$, then for $n \in \mathbb{N}_+$, X has the same distribution as $X_1 + X_2 + \dots + X_n$ where (X_1, X_2, \dots, X_n) are independent, and each has the normal distribution with mean μ/n and variance σ^2/n .

Finally, the normal distribution belongs to the family of general exponential distributions.

Suppose that X has the normal distribution with mean μ and variance σ^2 . The distribution is a two-parameter exponential family with natural parameters $\left(\frac{\mu}{\sigma^2}, -\frac{1}{2\sigma^2}\right)$, and natural statistics (X, X^2) .

Proof

Expanding the square, the normal PDF can be written in the form

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \exp\left(\frac{\mu}{\sigma^2}x - \frac{1}{2\sigma^2}x^2\right), \quad x \in \mathbb{R} \quad (5.6.16)$$

so the result follows from the definition of the general exponential family.

A number of other special distributions studied in this chapter are constructed from normally distributed variables. These include

- The lognormal distribution
- The folded normal distribution, which includes the half normal distribution as a special case
- The Rayleigh distribution
- The Maxwell distribution
- The Lévy distribution

Also, as mentioned at the beginning of this section, the importance of the normal distribution stems in large part from the central limit theorem, one of the fundamental theorems of probability. By virtue of this theorem, the normal distribution is connected to many other distributions, by means of limits and approximations, including the special distributions in the following list. Details are given in the individual sections.

- The binomial distribution
- The negative binomial distribution
- The Poisson distribution
- The gamma distribution
- The chi-square distribution
- The student t distribution
- The Irwin-Hall distribution

Computational Exercises

Suppose that the volume of beer in a bottle of a certain brand is normally distributed with mean 0.5 liter and standard deviation 0.01 liter.

1. Find the probability that a bottle will contain at least 0.48 liter.
2. Find the volume that corresponds to the 95th percentile

Answer

Let X denote the volume of beer in liters

1. $\mathbb{P}(X > 0.48) = 0.9772$
2. $x_{0.95} = 0.51645$

A metal rod is designed to fit into a circular hole on a certain assembly. The radius of the rod is normally distributed with mean 1 cm and standard deviation 0.002 cm. The radius of the hole is normally distributed with mean 1.01 cm and standard deviation 0.003 cm. The machining processes that produce the rod and the hole are independent. Find the probability that the rod is too big for the hole.

Answer

Let X denote the radius of the rod and Y the radius of the hole. $\mathbb{P}(Y - X < 0) = 0.0028$

The weight of a peach from a certain orchard is normally distributed with mean 8 ounces and standard deviation 1 ounce. Find the probability that the combined weight of 5 peaches exceeds 45 ounces.

Answer

Let X denote the combined weight of the 5 peaches, in ounces. $\mathbb{P}(X > 45) = 0.0127$

A Further Generalization

In some settings, it's convenient to consider a constant as having a normal distribution (with mean being the constant and variance 0, of course). This convention simplifies the statements of theorems and definitions in these settings. Of course, the formulas for

the [probability density function](#) and the [distribution function](#) do not hold for a constant, but the other results involving the [moment generating function](#), [linear transformations](#), and [sums](#) are still valid. Moreover, the result for [linear transformations](#) would hold for all a and b .

This page titled [5.6: The Normal Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.7: The Multivariate Normal Distribution

The multivariate normal distribution is among the most important of multivariate distributions, particularly in statistical inference and the study of Gaussian processes such as Brownian motion. The distribution arises naturally from linear transformations of independent normal variables. In this section, we consider the bivariate normal distribution first, because explicit results can be given and because graphical interpretations are possible. Then, with the aid of matrix notation, we discuss the general multivariate distribution.

The Bivariate Normal Distribution

The Standard Distribution

Recall that the probability density function ϕ of the standard normal distribution is given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad z \in \mathbb{R} \quad (5.7.1)$$

The corresponding distribution function is denoted Φ and is considered a special function in mathematics:

$$\Phi(z) = \int_{-\infty}^z \phi(x) dx = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx, \quad z \in \mathbb{R} \quad (5.7.2)$$

Finally, the moment generating function m is given by

$$m(t) = \mathbb{E}(e^{tZ}) = \exp\left[\frac{1}{2}\text{var}(tZ)\right] = e^{t^2/2}, \quad t \in \mathbb{R} \quad (5.7.3)$$

Suppose that Z and W are independent random variables, each with the standard normal distribution. The distribution of (Z, W) is known as the *standard bivariate normal distribution*.

The basic properties of the standard bivariate normal distribution follow easily from independence and properties of the (univariate) normal distribution. Recall first that the graph of a function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ is a surface. For $c \in \mathbb{R}$, the set of points $\{(x, y) \in \mathbb{R}^2 : f(x, y) = c\}$ is the *level curve* of f at level c . The graph of f can be understood by means of the level curves.

The probability density function ϕ_2 of the standard bivariate normal distribution is given by

$$\phi_2(z, w) = \frac{1}{2\pi} e^{-\frac{1}{2}(z^2 + w^2)}, \quad (z, w) \in \mathbb{R}^2 \quad (5.7.4)$$

1. The level curves of ϕ_2 are circles centered at the origin.
2. The mode of the distribution is $(0, 0)$.
3. ϕ_2 is concave downward on $\{(z, w) \in \mathbb{R}^2 : z^2 + w^2 < 1\}$

Proof

By independence, $\phi_2(z, w) = \phi(z)\phi(w)$ for $(z, w) \in \mathbb{R}^2$. Parts (a) and (b) are clear. For part (c), the second derivative matrix of ϕ_2 is

$$\begin{bmatrix} \phi_2(z, w)(z^2 - 1) & \phi_2(z, w)zw \\ \phi_2(z, w)zw & \phi_2(z, w)(w^2 - 1) \end{bmatrix} \quad (5.7.5)$$

with determinant $\phi_2^2(z, w)(1 - z^2 - w^2)$. The determinant is positive and the diagonal entries negative on the circular region $\{(z, w) \in \mathbb{R}^2 : z^2 + w^2 < 1\}$, so the matrix is negative definite on this region.

Clearly ϕ has a number of symmetry properties as well: $\phi_2(z, w)$ is symmetric in z about 0 so that $\phi_2(-z, w) = \phi_2(z, w)$; $\phi_2(z, w)$ is symmetric in w about 0 so that $\phi_2(z, -w) = \phi_2(z, w)$; $\phi_2(z, w)$ is symmetric in (z, w) so that $\phi_2(z, w) = \phi_2(w, z)$. In short, ϕ_2 has the classical “bell shape” that we associate with normal distributions.

Open the bivariate normal experiment, keep the default settings to get the standard bivariate normal distribution. Run the experiment 1000 times. Observe the cloud of points in the scatterplot, and compare the empirical density functions to the probability density functions.

Suppose that (Z, W) has the standard bivariate normal distribution. The moment generating function m_2 of (Z, W) is given by

$$m_2(s, t) = \mathbb{E}[\exp(sZ + tW)] = \exp\left[\frac{1}{2}\text{var}(sZ + tW)\right] = \exp\left[\frac{1}{2}(s^2 + t^2)\right], \quad (s, t) \in \mathbb{R}^2 \quad (5.7.6)$$

Proof

By independence, $m_2(s, t) = m(s)m(t)$ for $(s, t) \in \mathbb{R}^2$ where m is the standard normal MGF.

The General Distribution

The general bivariate normal distribution can be constructed by means of an affine transformation on a standard bivariate normal vector. The distribution has 5 parameters. As we will see, two are location parameters, two are scale parameters, and one is a correlation parameter.

Suppose that (Z, W) has the standard bivariate normal distribution. Let $\mu, \nu \in \mathbb{R}$; $\sigma, \tau \in (0, \infty)$; and $\rho \in (-1, 1)$, and let X and Y be new random variables defined by

$$X = \mu + \sigma Z \quad (5.7.7)$$

$$Y = \nu + \tau \rho Z + \tau \sqrt{1 - \rho^2} W \quad (5.7.8)$$

The joint distribution of (X, Y) is called the *bivariate normal distribution* with parameters $(\mu, \nu, \sigma, \tau, \rho)$

We can use the change of variables formula to find the joint probability density function.

Suppose that (X, Y) has the bivariate normal distribution with the parameters $(\mu, \nu, \sigma, \tau, \rho)$ as specified above. The joint probability density function f of (X, Y) is given by

$$f(x, y) = \frac{1}{2\pi\sigma\tau\sqrt{1-\rho^2}} \exp\left\{-\frac{1}{2(1-\rho^2)}\left[\frac{(x-\mu)^2}{\sigma^2} - 2\rho\frac{(x-\mu)(y-\nu)}{\sigma\tau} + \frac{(y-\nu)^2}{\tau^2}\right]\right\}, \quad (x, y) \in \mathbb{R}^2 \quad (5.7.9)$$

1. The level curves of f are ellipses centered at (μ, ν) .
2. The mode of the distribution is (μ, ν) .

Proof

Consider the transformation that defines (x, y) from (z, w) in the [definition](#). The inverse transformation is given by

$$z = \frac{x - \mu}{\sigma} \quad (5.7.10)$$

$$w = \frac{x - \mu}{\sigma\sqrt{1-\rho^2}} - \rho\frac{y - \nu}{\tau\sqrt{1-\rho^2}} \quad (5.7.11)$$

The Jacobian of the inverse transformation is

$$\frac{\partial(z, w)}{\partial(x, y)} = \frac{1}{\sigma\tau\sqrt{1-\rho^2}} \quad (5.7.12)$$

Note that the Jacobian is a constant, because the transformation is affine. The result now follows from the independence of Z and W , and the change of variables formula

1. Note that f has the form $f(x, y) = a \exp[-bg(x, y)]$ where a and b are positive constants and

$$g(x, y) = \frac{(x - \mu)^2}{\sigma^2} - 2\rho\frac{(x - \mu)(y - \nu)}{\sigma\tau} + \frac{(y - \nu)^2}{\tau^2}, \quad (x, y) \in \mathbb{R}^2 \quad (5.7.13)$$

The graph of g is a paraboloid opening upward. The level curves of f are the same as the level curves of g (but at different levels of course).

2. The maximum of f occurs at the minimum of g , at the point (μ, ν) .

The following theorem gives fundamental properties of the bivariate normal distribution.

Suppose that (X, Y) has the bivariate normal distribution with parameters $(\mu, \nu, \sigma, \tau, \rho)$ as specified above. Then

1. X is normally distributed with mean μ and standard deviation σ .
2. Y is normally distributed with mean ν and standard deviation τ .
3. $\text{cor}(X, Y) = \rho$.
4. X and Y are independent if and only if $\rho = 0$.

Proof

These result can be proved from the probability density function, but it's easier and more helpful to use the transformation definition. So, assume that (X, Y) is defined in terms of the standard bivariate normal pair (Z, W) as in the [definition](#).

1. $X = \mu + \sigma Z$ so X has the normal distribution with mean μ and standard deviation σ . This is a basic property of the normal distribution, and indeed is the way that the general normal variable is constructed from a standard normal variable.

2. Since Z and W are independent and each has the standard normal distribution, $Y = \nu + \tau\rho Z + \tau\sqrt{1-\rho^2}W$ is normally distributed by another basic property. Because Z and W have mean 0, it follows from the linear property of expected value that $\mathbb{E}(Y) = \nu$. Similarly, since Z and W have variance 1, it follows from basic properties of variance that $\text{var}(Y) = \tau^2\rho^2 + \tau^2(1-\rho^2) = \tau^2$.
3. Using the bi-linear property of covariance and independence we have $\text{cov}(X, Y) = \rho\tau\sigma \text{cov}(Z, Z) = \rho\tau\sigma$, and hence from (a) and (b), $\text{cor}(X, Y) = \rho$.
4. As a general property, recall that if X and Y are independent then $\text{cor}(X, Y) = 0$. Conversely, if $\rho = 0$ then $X = \mu + \sigma Z$ and $Y = \nu + \tau W$. Since Z and W are independent, so are X and Y .

Thus, two random variables with a joint normal distribution are independent if and only if they are uncorrelated.

In the bivariate normal experiment, change the standard deviations of X and Y with the scroll bars. Watch the change in the shape of the probability density functions. Now change the correlation with the scroll bar and note that the probability density functions do *not* change. For various values of the parameters, run the experiment 1000 times. Observe the cloud of points in the scatterplot, and compare the empirical density functions to the probability density functions.

In the case of perfect correlation ($\rho = 1$ or $\rho = -1$), the distribution of (X, Y) is also said to be bivariate normal, but *degenerate*. In this case, we know from our study of covariance and correlation that (X, Y) takes values on the regression line $\{(x, y) \in \mathbb{R}^2 : y = \nu + \rho\frac{\tau}{\sigma}(x - \mu)\}$, and hence does not have a probability density function (with respect to Lebesgue measure on \mathbb{R}^2). Degenerate normal distributions will be discussed in more detail [below](#).

In the bivariate normal experiment, run the experiment 1000 times with the values of ρ given below and selected values of σ and τ . Observe the cloud of points in the scatterplot and compare the empirical density functions to the probability density functions.

1. $\rho \in \{0, 0.3, 0.5, 0.7, 1\}$
2. $\rho \in \{-0.3, -0.5, -0.7, -1\}$

The conditional distributions are also normal.

Suppose that (X, Y) has the bivariate normal distribution with parameters $(\mu, \nu, \sigma, \tau, \rho)$ as specified above.

1. For $x \in \mathbb{R}$, the conditional distribution of Y given $X = x$ is normal with mean $\mathbb{E}(Y | X = x) = \nu + \rho\frac{\tau}{\sigma}(x - \mu)$ and variance $\text{var}(Y | X = x) = \tau^2(1 - \rho^2)$.
2. For $y \in \mathbb{R}$, the conditional distribution of X given $Y = y$ is normal with mean $\mathbb{E}(X | Y = y) = \mu + \rho\frac{\sigma}{\tau}(y - \nu)$ and variance $\text{var}(X | Y = y) = \sigma^2(1 - \rho^2)$.

Proof from density functions

By symmetry, we need only prove (a). The conditional PDF of Y given $X = x$ is $y \mapsto f(x, y)/g(x)$ where f is the [joint PDF](#), and where g is the PDF of X , namely the normal PDF with mean μ and standard deviation σ . The result then follows after some algebra.

Proof from random variables

Again, we only need to prove (a). We can assume that (X, Y) is defined in terms of a standard normal pair (Z, W) as in the [definition](#). Hence

$$Y = \nu + \rho\tau\frac{X - \mu}{\sigma} + \tau\sqrt{1 - \rho^2}W \quad (5.7.14)$$

Since that X and W are independent, the conditional distribution of Y given $X = x$ is the distribution of $\nu + \rho\tau\frac{x - \mu}{\sigma} + \tau\sqrt{1 - \rho^2}W$. The latter distribution is normal, with mean and variance specified in the theorem.

Note that the conditional variances do not depend on the value of the given variable.

In the bivariate normal experiment, set the standard deviation of X to 1.5, the standard deviation of Y to 0.5, and the correlation to 0.7.

1. Run the experiment 100 times.
2. For each run, compute $\mathbb{E}(Y | X = x)$ the predicted value of Y for the given the value of X .
3. Over all 100 runs, compute the square root of the average of the squared errors between the predicted value of Y and the true value of Y .

You may be perplexed by the lack of symmetry in how (X, Y) is defined in terms of (Z, W) in the original [definition](#). Note however that the distribution is completely determined by the 5 parameters. If we define $X' = \mu + \sigma\rho Z + \sigma\sqrt{1 - \rho^2}W$ and $Y' = \nu + \tau Z$ then (X', Y') has the same distribution as (X, Y) , namely the bivariate normal distribution with parameters $(\mu, \nu, \sigma, \tau, \rho)$ (although, of course (X', Y') and (X, Y) are different random vectors). There are other ways to define the same distribution as an affine transformation of (Z, W) —the situation will be clarified in the next subsection.

Suppose that (X, Y) has the bivariate normal distribution with parameters $(\mu, \nu, \sigma, \tau, \rho)$. Then (X, Y) has moment generating function M given by

$$M(s, t) = \mathbb{E}[\exp(sX + tY)] = \exp\left[\mathbb{E}(sX + tY) + \frac{1}{2}\text{var}(sX + tY)\right] = \exp\left[\mu s + \nu t + \frac{1}{2}(\sigma^2 s^2 + 2\rho\sigma\tau st + \tau^2 t^2)\right], \quad (s, t) \in \mathbb{R}^2 \quad (5.7.15)$$

Proof

Using the representation of (X, Y) in terms of the standard bivariate normal vector (Z, W) in the [definition](#) and collecting terms gives

$$M(s, t) = \mathbb{E}\left(\exp\left[(\mu s + \nu t) + (\sigma s + \rho\tau t)Z + \tau\sqrt{1 - \rho^2}tW\right]\right) \quad (5.7.16)$$

Hence from independence we have

$$M(s, t) = \exp(\mu s + \nu t)m(\sigma s + \rho\tau t)m\left(\tau\sqrt{1 - \rho^2}t\right) \quad (5.7.17)$$

where m is the standard normal MGF. Substituting and simplifying gives the result.

We showed [above](#) that if (X, Y) has a bivariate normal distribution then the marginal distributions of X and Y are also normal. The converse is not true.

Suppose that (X, Y) has probability density function f given by

$$f(x, y) = \frac{1}{2\pi}e^{-(x^2+y^2)/2}\left[1 + xy e^{-(x^2+y^2-2)/2}\right], \quad (x, y) \in \mathbb{R}^2 \quad (5.7.18)$$

1. X and Y each have standard normal distributions.
2. (X, Y) does not have a bivariate normal distribution.

Proof

Note that $f(x, y) = \phi_2(x, y)[1 + u(x)u(y)]$ for $(x, y) \in \mathbb{R}^2$, where ϕ_2 is the [bivariate standard normal PDF](#) and where u is given by $u(t) = te^{-(t^2-1)/2}$ for $t \in \mathbb{R}$. From simple calculus, u is symmetric about 0, has a local maximum at $t = 1$, and $u(t) \rightarrow 0$ as $t \rightarrow \infty$. In particular, $|u(t)| \leq 1$ for $t \in \mathbb{R}$ and hence $f(x, y) \geq 0$ for $(x, y) \in \mathbb{R}^2$. Next, a helpful trick is that we can write integrals of f as expected values of functions of a standard normal pair (Z, W) . In particular,

$$\int_{\mathbb{R}^2} f(x, y)dx dy = \mathbb{E}[1 + u(Z)u(W)] = 1 + \mathbb{E}[u(Z)]\mathbb{E}[u(W)] = 1 \quad (5.7.19)$$

since $\mathbb{E}[u(Z)] = \mathbb{E}[u(W)] = 0$ by the symmetry of the standard normal distribution and the symmetry of u about 0. Hence f is a valid PDF on \mathbb{R}^2 . Suppose now that (X, Y) has PDF f .

1. The PDF of X at $x \in \mathbb{R}$ is

$$\int_{\mathbb{R}} f(x, y)dy = \int_{\mathbb{R}} \phi_2(x, y)dy + u(x)\mathbb{E}[u(W)] = \phi(x) \quad (5.7.20)$$

where as usual, ϕ is the standard normal PDF on \mathbb{R} . By symmetry, Y also has the standard normal distribution.

2. f does not have the form of a [bivariate normal PDF](#) and hence (X, Y) does not have a bivariate normal distribution.

Transformations

Like its univariate counterpart, the family of bivariate normal distributions is preserved under two types of transformations on the underlying random vector: affine transformations and sums of independent vectors. We start with a preliminary result on affine transformations that should help clarify the original [definition](#). Throughout this discussion, we assume that the parameter vector $(\mu, \nu, \sigma, \tau, \rho)$ satisfies the usual conditions: $\mu, \nu \in \mathbb{R}$, and $\sigma, \tau \in (0, \infty)$, and $\rho \in (-1, 1)$.

Suppose that (Z, W) has the standard bivariate normal distribution. Let $X = a_1 + b_1Z + c_1W$ and $Y = a_2 + b_2Z + c_2W$ where the coefficients are in \mathbb{R} and $b_1c_2 - c_1b_2 \neq 0$. Then (X, Y) has a bivariate normal distribution with parameters given by

1. $\mathbb{E}(X) = a_1$
2. $\mathbb{E}(Y) = a_2$
3. $\text{var}(X) = b_1^2 + c_1^2$
4. $\text{var}(Y) = b_2^2 + c_2^2$
5. $\text{cov}(X, Y) = b_1b_2 + c_1c_2$

Proof

A direct proof using the change of variables formula is possible, but our goal is to show that (X, Y) can be written in the form given above in the [definition](#). First, parts (a)–(e) follow from basic properties of expected value, variance, and covariance. So, in the notation used in the definition, we have $\mu = a_1$, $\nu = a_2$, $\sigma = \sqrt{b_1^2 + c_1^2}$, $\tau = \sqrt{b_2^2 + c_2^2}$, and

$$\rho = \frac{b_1 b_2 + c_1 c_2}{\sqrt{b_1^2 + c_1^2} \sqrt{b_2^2 + c_2^2}} \quad (5.7.21)$$

(Note from the assumption on the coefficients that $b_1^2 + c_1^2 > 0$ and $b_2^2 + c_2^2 > 0$). Simple algebra shows that

$$\sqrt{1 - \rho^2} = \frac{b_1 c_2 - c_1 b_2}{\sqrt{b_1^2 + c_1^2} \sqrt{b_2^2 + c_2^2}} \quad (5.7.22)$$

Next we define

$$U = \frac{b_1 Z + c_1 W}{\sqrt{b_1^2 + c_1^2}} \quad (5.7.23)$$

$$V = \frac{c_1 Z - b_1 W}{\sqrt{b_1^2 + c_1^2}} \quad (5.7.24)$$

The transformation that defines (u, v) from (z, w) is its own inverse, and has Jacobian 1. Hence it follows that (U, V) has the same joint distribution as (Z, W) , namely the standard bivariate normal distribution. Simple algebra shows that

$$X = a_1 + \sqrt{b_1^2 + c_1^2} U = \mu + \sigma U \quad (5.7.25)$$

$$Y = a_2 + \frac{b_1 b_2 + c_1 c_2}{\sqrt{b_1^2 + c_1^2}} U + \frac{c_1 b_2 - b_1 c_2}{\sqrt{b_1^2 + c_1^2}} V = \nu + \tau \rho U + \tau \sqrt{1 - \rho^2} V \quad (5.7.26)$$

This is the form given in the [definition](#), so it follows that (X, Y) has a bivariate normal distribution.

Now it is easy to show more generally that the bivariate normal distribution is closed with respect to affine transformations.

Suppose that (X, Y) has the bivariate normal distribution with parameters $(\mu, \nu, \sigma, \tau, \rho)$. Define $U = a_1 + b_1 X + c_1 Y$ and $V = a_2 + b_2 X + c_2 Y$, where the coefficients are in \mathbb{R} and $b_1 c_2 - c_1 b_2 \neq 0$. Then (U, V) has a bivariate normal distribution with parameters as follows:

1. $\mathbb{E}(U) = a_1 + b_1 \mu + c_1 \nu$
2. $\mathbb{E}(V) = a_2 + b_2 \mu + c_2 \nu$
3. $\text{var}(U) = b_1^2 \sigma^2 + c_1^2 \tau^2 + 2b_1 c_1 \rho \sigma \tau$
4. $\text{var}(V) = b_2^2 \sigma^2 + c_2^2 \tau^2 + 2b_2 c_2 \rho \sigma \tau$
5. $\text{cov}(U, V) = b_1 b_2 \sigma^2 + c_1 c_2 \tau^2 + (b_1 c_2 + b_2 c_1) \rho \sigma \tau$

Proof

From our original [construction](#), we can write $X = \mu + \sigma Z$ and $Y = \nu + \tau \rho Z + \tau \sqrt{1 - \rho^2} W$ where (Z, W) has the standard bivariate normal distribution. Then by simple substitution, $U = A_1 + B_1 Z + C_1 W$ and $V = A_2 + B_2 Z + C_2 W$ where $A_i = a_i + b_i \mu + c_i \nu$, $B_i = b_i \sigma + c_i \tau \rho$, $C_i = c_i \tau \sqrt{1 - \rho^2}$ for $i \in \{1, 2\}$. By simple algebra,

$$B_1 C_2 - C_1 B_2 = \sigma \tau \sqrt{1 - \rho^2} (b_1 c_2 - c_1 b_2) \neq 0 \quad (5.7.27)$$

Hence (U, V) has a bivariate normal distribution from the previous theorem. Parts (a)–(e) follow from basic properties of expected value, variance, and covariance.

The bivariate normal distribution is preserved with respect to sums of independent variables.

Suppose that (X_i, Y_i) has the bivariate normal distribution with parameters $(\mu_i, \nu_i, \sigma_i, \tau_i, \rho_i)$ for $i \in \{1, 2\}$, and that (X_1, Y_1) and (X_2, Y_2) are independent. Then $(X_1 + X_2, Y_1 + Y_2)$ has the bivariate normal distribution with parameters given by

1. $\mathbb{E}(X_1 + X_2) = \mu_1 + \mu_2$
2. $\mathbb{E}(Y_1 + Y_2) = \nu_1 + \nu_2$
3. $\text{var}(X_1 + X_2) = \sigma_1^2 + \sigma_2^2$
4. $\text{var}(Y_1 + Y_2) = \tau_1^2 + \tau_2^2$

$$5. \operatorname{cov}(X_1 + X_2, Y_1 + Y_2) = \rho_1 \sigma_1 \tau_1 + \rho_2 \sigma_2 \tau_2$$

Proof

Let M_i denote the MGF of (X_i, Y_i) for $i \in \{1, 2\}$ and let M denote the MGF of $(X_1 + X_2, Y_1 + Y_2)$. By independence, $M(s, t) = M_1(s, t)M_2(s, t)$ for $(s, t) \in \mathbb{R}^2$. Using the [bivariate normal MGF](#), and basic properties of the exponential function,

$$M(s, t) = \exp\left(\mathbb{E}[s(X_1 + X_2) + t(Y_1 + Y_2)] + \frac{1}{2}\operatorname{var}[s(X_1 + X_2) + t(Y_1 + Y_2)]\right), \quad (s, t) \in \mathbb{R}^2 \quad (5.7.28)$$

Of course from basic properties of expected value, variance, and covariance,

$$\mathbb{E}[s(X_1 + X_2) + t(Y_1 + Y_2)] = s(\mu_1 + \mu_2) + t(\nu_1 + \nu_2) \quad (5.7.29)$$

$$\operatorname{var}[s(X_1 + X_2) + t(Y_1 + Y_2)] = s(\sigma_1^2 + \sigma_2^2) + t(\tau_1^2 + \tau_2^2) + 2st(\rho_1 \sigma_1 \tau_1 + \rho_2 \sigma_2 \tau_2) \quad (5.7.30)$$

Substituting gives the result.

The following result is important in the simulation of normal variables.

Suppose that (Z, W) has the standard bivariate normal distribution. Define the polar coordinates (R, Θ) of (Z, W) by the equations $Z = R \cos \Theta$, $W = R \sin \Theta$ where $R \geq 0$ and $0 \leq \Theta < 2\pi$. Then

1. R has probability density function g given by $g(r) = r e^{-\frac{1}{2}r^2}$ for $r \in [0, \infty)$.
2. Θ is uniformly distributed on $[0, 2\pi)$.
3. R and Θ are independent.

Proof

The Jacobian of the polar coordinate transformation that gives (z, w) from (r, θ) is r , as we all remember from calculus. Hence by the change of variables theorem, the PDF g of (R, Θ) in terms of the from [standard normal PDF](#) ϕ_2 is given by

$$g(r, \theta) = \phi_2(r \cos \theta, r \sin \theta) r = \frac{1}{2\pi} r e^{-r^2/2}, \quad (r, \theta) \in [0, \infty) \times [0, 2\pi) \quad (5.7.31)$$

The result then follows from the factorization theorem for independent random variables.

The distribution of R is known as the *standard Rayleigh distribution*, named for William Strutt, Lord Rayleigh. The Rayleigh distribution studied in more detail in a separate section.

Since the quantile function Φ^{-1} of the normal distribution cannot be given in a simple, closed form, we cannot use the usual random quantile method of simulating a normal random variable. However, the quantile method works quite well to simulate a Rayleigh variable, and of course simulating uniform variables is trivial. Hence we have a way of simulating a standard bivariate normal vector with a pair of random numbers (which, you will recall are independent random variables, each with the standard uniform distribution, that is, the uniform distribution on $[0, 1)$).

Suppose that U and V are independent random variables, each with the standard uniform distribution. Let $R = \sqrt{-2 \ln U}$ and $\Theta = 2\pi V$. Define $Z = R \cos \Theta$ and $W = R \sin \Theta$. Then (Z, W) has the standard bivariate normal distribution.

Proof

The Rayleigh distribution function F is given by $F(r) = 1 - e^{-r^2/2}$ for $r \in [0, \infty)$ and hence the quantile function is given by $F^{-1}(p) = \sqrt{-2 \ln(1-p)}$ for $p \in [0, 1)$. Hence if U has the standard uniform distribution, then $\sqrt{-2 \ln(1-U)}$ has the Rayleigh distribution. But $1 - U$ also has the standard uniform distribution so $R = \sqrt{-2 \ln U}$ also has the Rayleigh distribution. If V has the standard uniform distribution then of course $2\pi V$ is uniformly distributed on $[0, 2\pi)$. If U and V are independent, then so are R and Θ . By the previous theorem, if $Z = R \cos \Theta$ and $W = R \sin \Theta$, then (Z, W) has the standard bivariate normal distribution.

Of course, if we can simulate (Z, W) with a standard bivariate normal distribution, then we can simulate (X, Y) with the general bivariate normal distribution, with parameter $(\mu, \nu, \sigma, \tau, \rho)$ by definition (5), namely $X = \mu + \sigma Z$, $Y = \nu + \tau \rho Z + \tau \sqrt{1 - \rho^2} W$.

The General Multivariate Normal Distribution

The general multivariate normal distribution is a natural generalization of the bivariate normal distribution studied above. The exposition is very compact and elegant using expected value and covariance matrices, and would be horribly complex without these tools. Thus, this section requires some prerequisite knowledge of linear algebra. In particular, recall that \mathbf{A}^T denotes the transpose of a matrix \mathbf{A} and that we identify a vector in \mathbb{R}^n with the corresponding $n \times 1$ column vector.

The Standard Distribution

Suppose that $\mathbf{Z} = (Z_1, Z_2, \dots, Z_n)$ is a vector of independent random variables, each with the standard normal distribution. Then \mathbf{Z} is said to have the n -dimensional *standard normal distribution*.

1. $\mathbb{E}(\mathbf{Z}) = \mathbf{0}$ (the zero vector in \mathbb{R}^n).
2. $\text{vc}(\mathbf{Z}) = I$ (the $n \times n$ identity matrix).

\mathbf{Z} has probability density function ϕ_n given by

$$\phi_n(\mathbf{z}) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2} \mathbf{z} \cdot \mathbf{z}\right) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2} \sum_{i=1}^n z_i^2\right), \quad \mathbf{z} = (z_1, z_2, \dots, z_n) \in \mathbb{R}^n \quad (5.7.32)$$

where as usual, ϕ is the standard normal PDF.

Proof

By independence, $\phi_n(\mathbf{z}) = \phi(z_1)\phi(z_2) \cdots \phi(z_n)$.

\mathbf{Z} has moment generating function m_n given by

$$m_n(\mathbf{t}) = \mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{Z})] = \exp\left[\frac{1}{2} \text{var}(\mathbf{t} \cdot \mathbf{Z})\right] = \exp\left(\frac{1}{2} \mathbf{t} \cdot \mathbf{t}\right) = \exp\left(\frac{1}{2} \sum_{i=1}^n t_i^2\right), \quad \mathbf{t} = (t_1, t_2, \dots, t_n) \in \mathbb{R}^n \quad (5.7.33)$$

Proof

By independence, $\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{Z})] = m(t_1)m(t_2) \cdots m(t_n)$ where m is the standard normal MGF.

The General Distribution

Suppose that \mathbf{Z} has the n -dimensional standard normal distribution. Suppose also that $\boldsymbol{\mu} \in \mathbb{R}^n$ and that $\mathbf{A} \in \mathbb{R}^{n \times n}$ is invertible. The random vector $\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$ is said to have an n -dimensional normal distribution.

1. $\mathbb{E}(\mathbf{X}) = \boldsymbol{\mu}$.
2. $\text{vc}(\mathbf{X}) = \mathbf{A}\mathbf{A}^T$.

Proof

1. From the linear property of expected value, $\mathbb{E}(\mathbf{X}) = \boldsymbol{\mu} + \mathbf{A}\mathbb{E}(\mathbf{Z}) = \boldsymbol{\mu}$.
2. From basic properties of the variance-covariance matrix, $\text{vc}(\mathbf{X}) = \mathbf{A}\mathbf{A}^T \text{vc}(\mathbf{Z}) = \mathbf{A}\mathbf{A}^T$.

In the context of this result, recall that the variance-covariance matrix $\text{vc}(\mathbf{X}) = \mathbf{A}\mathbf{A}^T$ is symmetric and positive definite (and hence also invertible). We will now see that the multivariate normal distribution is completely determined by the expected value vector $\boldsymbol{\mu}$ and the variance-covariance matrix \mathbf{V} , and hence these give the basic parameters of the distribution.

Suppose that \mathbf{X} has an n -dimensional normal distribution with expected value vector $\boldsymbol{\mu}$ and variance-covariance matrix \mathbf{V} . The probability density function f of \mathbf{X} is given by

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{n/2} \sqrt{\det(\mathbf{V})}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}) \cdot \mathbf{V}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right], \quad \mathbf{x} \in \mathbb{R}^n \quad (5.7.34)$$

Proof

From the definition can assume that $\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$ where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is invertible and \mathbf{Z} has the n -dimensional standard normal distribution, so that $\mathbf{V} = \mathbf{A}\mathbf{A}^T$. The inverse of the transformation $\mathbf{x} = \boldsymbol{\mu} + \mathbf{A}\mathbf{z}$ is $\mathbf{z} = \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu})$ and hence the Jacobian of the inverse transformation is $\det(\mathbf{A}^{-1}) = 1/\det(\mathbf{A})$. Using the multivariate change of variables theorem,

$$f(\mathbf{x}) = \frac{1}{|\det(\mathbf{A})|} \phi_n[\mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu})] = \frac{1}{(2\pi)^{n/2} |\det(\mathbf{A})|} \exp\left[-\frac{1}{2} \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \cdot \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right], \quad \mathbf{x} \in \mathbb{R}^n \quad (5.7.35)$$

But $\det(\mathbf{V}) = \det(\mathbf{A}\mathbf{A}^T) = \det(\mathbf{A}) \det(\mathbf{A}^T) = [\det(\mathbf{A})]^2$ and hence $|\det(\mathbf{A})| = \sqrt{\det(\mathbf{V})}$. Also,

$$\mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \cdot \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu}) = [\mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu})]^T \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu}) = (\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{A}^{-1})^T \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \quad (5.7.36)$$

$$= (\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{A}^T)^{-1} \mathbf{A}^{-1}(\mathbf{x} - \boldsymbol{\mu}) = (\mathbf{x} - \boldsymbol{\mu})^T (\mathbf{A}\mathbf{A}^T)^{-1}(\mathbf{x} - \boldsymbol{\mu}) \quad (5.7.37)$$

$$= (\mathbf{x} - \boldsymbol{\mu})^T \mathbf{V}^{-1}(\mathbf{x} - \boldsymbol{\mu}) = (\mathbf{x} - \boldsymbol{\mu}) \cdot \mathbf{V}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \quad (5.7.38)$$

Suppose again that \mathbf{X} has an n -dimensional normal distribution with expected value vector $\boldsymbol{\mu}$ and variance-covariance matrix \mathbf{V} . The moment generating function M of \mathbf{X} is given by

$$M(\mathbf{t}) = \mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{X})] = \exp\left[\mathbb{E}(\mathbf{t} \cdot \mathbf{X}) + \frac{1}{2} \text{var}(\mathbf{t} \cdot \mathbf{X})\right] = \exp\left(\mathbf{t} \cdot \boldsymbol{\mu} + \frac{1}{2} \mathbf{t} \cdot \mathbf{V}\mathbf{t}\right), \quad \mathbf{t} \in \mathbb{R}^n \quad (5.7.39)$$

Proof

Once again we start with the [definition](#) and assume that $\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$ where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is invertible. we have $\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{X})] = \exp(\mathbf{t} \cdot \boldsymbol{\mu}) \mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{A}\mathbf{Z})]$. But $\mathbf{t} \cdot \mathbf{A}\mathbf{Z} = (\mathbf{A}^T \mathbf{t}) \cdot \mathbf{Z}$ so using the [MGF](#) of \mathbf{Z} we have

$$\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{A}\mathbf{Z})] = \exp\left[\frac{1}{2}(\mathbf{A}^T \mathbf{t}) \cdot (\mathbf{A}^T \mathbf{t})\right] = \exp\left[\frac{1}{2}\mathbf{t}^T \mathbf{A} \mathbf{A}^T \mathbf{t}\right] = \exp\left[\frac{1}{2}\mathbf{t} \cdot \mathbf{V} \mathbf{t}\right] \quad (5.7.40)$$

Of course, the moment generating function completely determines the distribution. Thus, if a random vector \mathbf{X} in \mathbb{R}^n has a moment generating function of the form given above, for some $\boldsymbol{\mu} \in \mathbb{R}^n$ and symmetric, positive definite $\mathbf{V} \in \mathbb{R}^{n \times n}$, then \mathbf{X} has the n -dimensional normal distribution with mean $\boldsymbol{\mu}$ and variance-covariance matrix \mathbf{V} .

Note again that in the representation $\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$, the distribution of \mathbf{X} is uniquely determined by the expected value vector $\boldsymbol{\mu}$ and the variance-covariance matrix $\mathbf{V} = \mathbf{A}\mathbf{A}^T$, but not by $\boldsymbol{\mu}$ and \mathbf{A} . In general, for a given positive definite matrix \mathbf{V} , there are many invertible matrices \mathbf{A} such that $\mathbf{V} = \mathbf{A}\mathbf{A}^T$ (the matrix \mathbf{A} is a bit like a square root of \mathbf{V}). A theorem in matrix theory states that there is a unique lower triangular matrix \mathbf{L} with this property. The representation $\mathbf{X} = \boldsymbol{\mu} + \mathbf{L}\mathbf{Z}$ is known as the *canonical representation* of \mathbf{X} .

If $\mathbf{X} = (X, Y)$ has bivariate normal distribution with parameters $(\mu, \nu, \sigma, \tau, \rho)$, then the lower triangular matrix \mathbf{L} such that $\mathbf{L}\mathbf{L}^T = \text{vc}(\mathbf{X})$ is

$$\mathbf{L} = \begin{bmatrix} \sigma & 0 \\ \tau\rho & \tau\sqrt{1-\rho^2} \end{bmatrix} \quad (5.7.41)$$

Proof

Note that

$$\mathbf{L}\mathbf{L}^T = \begin{bmatrix} \sigma^2 & \sigma\tau\rho \\ \sigma\tau\rho & \tau^2 \end{bmatrix} = \text{vc}(X, Y) \quad (5.7.42)$$

Note that the matrix \mathbf{L} above gives the canonical representation of (X, Y) in terms of the standard normal vector (Z, W) in the original [definition](#), namely $X = \mu + \sigma Z$, $Y = \nu + \tau\rho Z + \tau\sqrt{1-\rho^2} W$.

If the matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ in the [definition](#) is not invertible, then the variance-covariance matrix $\mathbf{V} = \mathbf{A}\mathbf{A}^T$ is symmetric, but only positive semi-definite. The random vector $\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$ takes values in a lower dimensional affine subspace of \mathbb{R}^n that has measure 0 relative to n -dimensional Lebesgue measure λ_n . Thus, \mathbf{X} does not have a probability density function relative to λ_n , and so the distribution is degenerate. However, the formula for the [moment generating function](#) still holds. Degenerate normal distributions are discussed in more detail [below](#).

Transformations

The multivariate normal distribution is invariant under two basic types of transformations on the underlying random vectors: affine transformations (with linearly independent rows), and concatenation of independent vectors. As simple corollaries of these two results, the normal distribution is also invariant with respect to subsequences of the random vector, re-orderings of the terms in the random vector, and sums of independent random vectors. The main tool that we will use is the moment generating function. We start with the first main result on affine transformations.

Suppose that \mathbf{X} has the n -dimensional normal distribution with mean vector $\boldsymbol{\mu}$ and variance-covariance matrix \mathbf{V} . Suppose also that $\mathbf{a} \in \mathbb{R}^m$ and that $\mathbf{A} \in \mathbb{R}^{m \times n}$ has linearly independent rows (thus, $m \leq n$). Then $\mathbf{Y} = \mathbf{a} + \mathbf{A}\mathbf{X}$ has an m -dimensional normal distribution, with

1. $\mathbb{E}(\mathbf{Y}) = \mathbf{a} + \mathbf{A}\boldsymbol{\mu}$
2. $\text{vc}(\mathbf{Y}) = \mathbf{A}\mathbf{V}\mathbf{A}^T$

Proof

For $\mathbf{t} \in \mathbb{R}^m$, $\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{Y})] = \exp(\mathbf{t} \cdot \mathbf{a}) \mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{A}\mathbf{X})]$. but $\mathbf{t} \cdot \mathbf{A}\mathbf{X} = (\mathbf{A}^T \mathbf{t}) \cdot \mathbf{X}$, so using the [MGF](#) of \mathbf{X} we have

$$\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{A}\mathbf{X})] = \exp\left[(\mathbf{A}^T \mathbf{t}) \cdot \boldsymbol{\mu} + \frac{1}{2}(\mathbf{A}^T \mathbf{t}) \cdot \mathbf{V}(\mathbf{A}^T \mathbf{t})\right] \quad (5.7.43)$$

But $(\mathbf{A}^T \mathbf{t}) \cdot \boldsymbol{\mu} = \mathbf{t} \cdot \mathbf{A}\boldsymbol{\mu}$ and $(\mathbf{A}^T \mathbf{t}) \cdot \mathbf{V}(\mathbf{A}^T \mathbf{t}) = \mathbf{t} \cdot (\mathbf{A}\mathbf{V}\mathbf{A}^T) \mathbf{t}$, so letting $\mathbf{b} = \mathbf{a} + \mathbf{A}\boldsymbol{\mu}$ and $\mathbf{U} = \mathbf{A}\mathbf{V}\mathbf{A}^T$ and putting the pieces together, we have $\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{Y})] = \exp[\mathbf{b} \cdot \mathbf{t} + \frac{1}{2}\mathbf{t} \cdot \mathbf{U} \mathbf{t}]$.

A clearly important special case is $m = n$, which generalizes the [definition](#). Thus, if $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{A} \in \mathbb{R}^{n \times n}$ is invertible, then $\mathbf{Y} = \mathbf{a} + \mathbf{A}\mathbf{X}$ has an n -dimensional normal distribution. Here are some other corollaries:

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ has an n -dimensional normal distribution. If $\{i_1, i_2, \dots, i_m\}$ is a set of distinct indices, then $\mathbf{Y} = (X_{i_1}, X_{i_2}, \dots, X_{i_m})$ has an m -dimensional normal distribution.

Proof

Let $A \in \mathbb{R}^{m \times n}$ be the matrix defined by the condition that for $j \in \{1, 2, \dots, m\}$, row j has 1 in position i_j and has 0 in all other positions. Then A has linearly independent rows (since the i_j are distinct in j) and $Y = AX$. Thus the result follows from the general theorem on [affine transformations](#).

In the context of the previous result, if X has mean vector μ and variance-covariance matrix V , then Y has mean vector $A\mu$ and variance-covariance matrix $AVAT$, where A is the 0-1 matrix defined in the proof. As simple corollaries, note that if $X = (X_1, X_2, \dots, X_n)$ has an n -dimensional normal distribution, then any permutation of the coordinates of X also has an n -dimensional normal distribution, and (X_1, X_2, \dots, X_m) has an m -dimensional normal distribution for any $m \leq n$. Here is a slight extension of the last statement.

Suppose that X is a random vector in \mathbb{R}^m , Y is a random vector in \mathbb{R}^n , and that (X, Y) has an $(m+n)$ -dimensional normal distribution. Then

1. X has an m -dimensional normal distribution.
2. Y has an n -dimensional normal distribution.
3. X and Y are independent if and only if $\text{cov}(X, Y) = \mathbf{0}$ (the $m \times n$ zero matrix).

Proof

As we already noted, parts (a) and (b) are a simple consequence of the [previous theorem](#). Thus, we just need to verify (c). In block form, note that

$$\text{vc}(X, Y) = \begin{bmatrix} \text{vc}(X) & \text{cov}(X, Y) \\ \text{cov}(Y, X) & \text{vc}(Y) \end{bmatrix} \quad (5.7.44)$$

Now let M denote the moment generating function of (X, Y) , M_1 the MGF of X , and M_2 the MGF of Y . From the form of the [MGF](#), note that $M(s, t) = M_1(s)M_2(t)$ for all $s \in \mathbb{R}^m$, $t \in \mathbb{R}^n$ if and only if $\text{cov}(X, Y) = \mathbf{0}$, the $m \times n$ zero matrix.

Next is the converse to part (c) of the previous result: concatenating independent normally distributed vectors produces another normally distributed vector.

Suppose that X has the m -dimensional normal distribution with mean vector μ and variance-covariance matrix U , Y has the n -dimensional normal distribution with mean vector ν and variance-covariance matrix V , and that X and Y are independent. Then $Z = (X, Y)$ has the $m+n$ -dimensional normal distribution with

1. $\mathbb{E}(X, Y) = (\mu, \nu)$
2. $\text{vc}(X, Y) = \begin{bmatrix} \text{vc}(X) & \mathbf{0} \\ \mathbf{0}^T & \text{vc}(Y) \end{bmatrix}$ where $\mathbf{0}$ is the $m \times n$ zero matrix.

Proof

For $t \in \mathbb{R}^{m+n}$, write t in block form as $t = (r, s)$ where $r \in \mathbb{R}^m$ and $s \in \mathbb{R}^n$. By independence, the MGF of (X, Y) is

$$\mathbb{E}(\exp[t \cdot (X, Y)]) = \mathbb{E}[r \cdot X + s \cdot Y] = \mathbb{E}[\exp(r \cdot X)] \mathbb{E}[\exp(s \cdot Y)] \quad (5.7.45)$$

Using the formula for the [normal MGF](#) we have

$$\mathbb{E}(\exp[t \cdot (X, Y)]) = \exp\left(r \cdot \mu + \frac{1}{2} r \cdot U r\right) \exp\left(s \cdot \nu + \frac{1}{2} s \cdot V s\right) = \exp\left[\left(r \cdot \mu + s \cdot \nu\right) + \frac{1}{2} (r \cdot U r + s \cdot V s)\right] \quad (5.7.46)$$

But $r \cdot \mu + s \cdot \nu = t \cdot (\mu, \nu)$ and $r \cdot U r + s \cdot V s = t \cdot \begin{bmatrix} \text{vc}(X) & \mathbf{0} \\ \mathbf{0}^T & \text{vc}(Y) \end{bmatrix} t$ so the proof is complete

Just as in the univariate case, the normal family of distributions is closed with respect to sums of independent variables. The proof follows easily from the previous result.

Suppose that X has the n -dimensional normal distribution with mean vector μ and variance-covariance matrix U , Y has the n -dimensional normal distribution with mean vector ν and variance-covariance matrix V , and that X and Y are independent. Then $X + Y$ has the n -dimensional normal distribution with

1. $\mathbb{E}(X + Y) = \mu + \nu$
2. $\text{vc}(X + Y) = U + V$

Proof

From the previous result (X, Y) has a $2n$ -dimensional normal distribution. Moreover, $X + Y = A(X, Y)$ where A is the $n \times 2n$ matrix defined by the condition that for $i \in \{1, 2, \dots, n\}$, row i has 1 in positions i and $n+i$ and 0 in all other positions. The matrix A has linearly

independent rows and thus the result follows from the general theorem on [affine transformations](#). Parts (a) and (b) are standard results for sums of independent random vectors.

We close with a trivial corollary to the general result on [affine transformation](#), but this corollary points the way to a further generalization of the multivariate normal distribution that includes the degenerate distributions.

Suppose that \mathbf{X} has an n -dimensional normal distribution with mean vector $\boldsymbol{\mu}$ and variance-covariance matrix \mathbf{V} , and that $\mathbf{a} \in \mathbb{R}^n$ with $\mathbf{a} \neq \mathbf{0}$. Then $Y = \mathbf{a} \cdot \mathbf{X}$ has a (univariate) normal distribution with

1. $\mathbb{E}(Y) = \mathbf{a} \cdot \boldsymbol{\mu}$
2. $\text{var}(Y) = \mathbf{a} \cdot \mathbf{V} \mathbf{a}$

Proof

Note again that $\mathbf{a} \cdot \mathbf{X} = \mathbf{a}^T \mathbf{X}$. Since $\mathbf{a} \neq \mathbf{0}$, the single row of \mathbf{a}^T is linearly independent and hence the result follows from the general theorem on [affine transformations](#).

A Further Generalization

The last result can be used to give a simple, elegant definition of the multivariate normal distribution that includes the degenerate distributions as well as the ones we have considered so far. First we will adopt our general definition of the univariate normal distribution that includes constant random variables.

A random variable \mathbf{X} that takes values in \mathbb{R}^n has an n -dimensional normal distribution if and only if $\mathbf{a} \cdot \mathbf{X}$ has a univariate normal distribution for every $\mathbf{a} \in \mathbb{R}^n$.

Although an n -dimensional normal distribution may not have a probability density function with respect to n -dimensional Lebesgue measure λ_n , the form of the moment generating function is unchanged.

Suppose that \mathbf{X} has mean vector $\boldsymbol{\mu}$ and variance-covariance matrix \mathbf{V} , and that \mathbf{X} has an n -dimensional normal distribution. The moment generating function of \mathbf{X} is given by

$$\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{X})] = \exp\left[\mathbb{E}(\mathbf{t} \cdot \mathbf{X}) + \frac{1}{2}\text{var}(\mathbf{t} \cdot \mathbf{X})\right] = \exp\left(\mathbf{t} \cdot \boldsymbol{\mu} + \frac{1}{2}\mathbf{t} \cdot \mathbf{V} \mathbf{t}\right), \quad \mathbf{t} \in \mathbb{R}^n \quad (5.7.47)$$

Proof

If $\mathbf{t} \in \mathbb{R}^n$, then by definition, $\mathbf{t} \cdot \mathbf{X}$ has a univariate normal distribution. Thus $\mathbb{E}[\exp(\mathbf{t} \cdot \mathbf{X})]$ is simply the moment generating function of $\mathbf{t} \cdot \mathbf{X}$, evaluated at the argument 1. The results then follow from the univariate MGF.

Our new general definition really is a generalization.

Suppose that \mathbf{X} has an n -dimensional normal distribution in the sense of the [general definition](#), and that the distribution of \mathbf{X} has a probability density function on \mathbb{R}^n with respect to Lebesgue measure λ_n . Then \mathbf{X} has an n -dimensional normal distribution in the sense of our [original definition](#).

Proof

This follows from our previous results, since both the MGF and the PDF completely determine the distribution of \mathbf{X} .

This page titled [5.7: The Multivariate Normal Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.8: The Gamma Distribution

In this section we will study a family of distributions that has special importance in probability and statistics. In particular, the arrival times in the Poisson process have gamma distributions, and the chi-square distribution in statistics is a special case of the gamma distribution. Also, the gamma distribution is widely used to model physical quantities that take positive values.

The Gamma Function

Before we can study the gamma *distribution*, we need to introduce the gamma *function*, a special function whose values will play the role of the normalizing constants.

Definition

The *gamma function* Γ is defined as follows

$$\Gamma(k) = \int_0^{\infty} x^{k-1} e^{-x} dx, \quad k \in (0, \infty) \quad (5.8.1)$$

The function is well defined, that is, the integral converges for any $k > 0$. On the other hand, the integral diverges to ∞ for $k \leq 0$.

Proof

Note that

$$\int_0^{\infty} x^{k-1} e^{-x} dx = \int_0^1 x^{k-1} e^{-x} dx + \int_1^{\infty} x^{k-1} e^{-x} dx \quad (5.8.2)$$

For the first integral on the right,

$$\int_0^1 x^{k-1} e^{-x} dx \leq \int_0^1 x^{k-1} dx = \frac{1}{k} \quad (5.8.3)$$

For the second integral, let $n = \lceil k \rceil$. Then

$$\int_1^{\infty} x^{k-1} e^{-x} dx \leq \int_1^{\infty} x^{n-1} e^{-x} dx \quad (5.8.4)$$

The last integral can be evaluated explicitly by integrating by parts, and is finite for every $n \in \mathbb{N}_+$.

Finally, if $k \leq 0$, note that

$$\int_0^1 x^{k-1} e^{-x} dx \geq e^{-1} \int_0^1 x^{k-1} dx = \infty \quad (5.8.5)$$

The gamma function was first introduced by Leonhard Euler.




Figure 5.8.1: The graph of the gamma function on the interval $(0, 5)$

The (lower) *incomplete gamma function* is defined by

$$\Gamma(k, x) = \int_0^x t^{k-1} e^{-t} dt, \quad k, x \in (0, \infty) \quad (5.8.6)$$

Properties

Here are a few of the essential properties of the gamma function. The first is the fundamental identity.

$$\Gamma(k+1) = k \Gamma(k) \quad \text{for } k \in (0, \infty).$$

Proof

This follows from integrating by parts, with $u = x^k$ and $dv = e^{-x} dx$:

$$\Gamma(k+1) = \int_0^\infty x^k e^{-x} dx = (-x^k e^{-x})_0^\infty + \int_0^\infty kx^{k-1} e^{-x} dx = 0 + k \Gamma(k) \quad (5.8.7)$$

Applying this result repeatedly gives

$$\Gamma(k+n) = k(k+1) \cdots (k+n-1) \Gamma(k), \quad n \in \mathbb{N}_+ \quad (5.8.8)$$

It's clear that the gamma function is a continuous extension of the factorial function.

$$\Gamma(k+1) = k! \quad \text{for } k \in \mathbb{N}.$$

Proof

This follows from the [fundamental identity](#) and the fact that $\Gamma(1) = 1$.

The values of the gamma function for non-integer arguments generally cannot be expressed in simple, closed forms. However, there are exceptions.

$$\Gamma\left(\frac{1}{2}\right) = \sqrt{\pi}.$$

Proof

By definition,

$$\Gamma\left(\frac{1}{2}\right) = \int_0^\infty x^{-1/2} e^{-x} dx \quad (5.8.9)$$

Substituting $x = z^2/2$ gives

$$\Gamma\left(\frac{1}{2}\right) = \int_0^\infty \sqrt{2} e^{-z^2/2} dz = 2\sqrt{\pi} \int_0^\infty \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz \quad (5.8.10)$$

But the last integrand is the PDF of the standard normal distribution, and so the integral evaluates to $\frac{1}{2}$

We can generalize the last result to odd multiples of $\frac{1}{2}$.

For $n \in \mathbb{N}$,

$$\Gamma\left(\frac{2n+1}{2}\right) = \frac{1 \cdot 3 \cdots (2n-1)}{2^n} \sqrt{\pi} = \frac{(2n)!}{4^n n!} \sqrt{\pi} \quad (5.8.11)$$

Proof

This follows from the [previous result](#) and the [fundamental identity](#).

Stirling's Approximation

One of the most famous asymptotic formulas for the gamma function is *Stirling's formula*, named for James Stirling. First we need to recall a definition.

Suppose that $f, g: D \rightarrow (0, \infty)$ where $D = (0, \infty)$ or $D = \mathbb{N}_+$. Then $f(x) \approx g(x)$ as $x \rightarrow \infty$ means that

$$\frac{f(x)}{g(x)} \rightarrow 1 \text{ as } x \rightarrow \infty \quad (5.8.12)$$

Stirling's formula

$$\Gamma(x+1) \approx \left(\frac{x}{e}\right)^x \sqrt{2\pi x} \text{ as } x \rightarrow \infty \quad (5.8.13)$$

As a special case, Stirling's result gives an asymptotic formula for the factorial function:

$$n! \approx \left(\frac{n}{e}\right)^n \sqrt{2\pi n} \text{ as } n \rightarrow \infty \quad (5.8.14)$$

The Standard Gamma Distribution

Distribution Functions

The *standard gamma distribution* with *shape parameter* $k \in (0, \infty)$ is a continuous distribution on $(0, \infty)$ with probability density function f given by

$$f(x) = \frac{1}{\Gamma(k)} x^{k-1} e^{-x}, \quad x \in (0, \infty) \quad (5.8.15)$$

Clearly f is a valid probability density function, since $f(x) > 0$ for $x > 0$, and by definition, $\Gamma(k)$ is the normalizing constant for the function $x \mapsto x^{k-1} e^{-x}$ on $(0, \infty)$. The following theorem shows that the gamma density has a rich variety of shapes, and shows why k is called the shape parameter.

The gamma probability density function f with shape parameter $k \in (0, \infty)$ satisfies the following properties:

1. If $0 < k < 1$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
2. If $k = 1$, f is decreasing with $f(0) = 1$.
3. If $k > 1$, f increases and then decreases, with mode at $k - 1$.
4. If $0 < k \leq 1$, f is concave upward.
5. If $1 < k \leq 2$, f is concave downward and then upward, with inflection point at $k - 1 + \sqrt{k - 1}$.
6. If $k > 2$, f is concave upward, then downward, then upward again, with inflection points at $k - 1 \pm \sqrt{k - 1}$.

Proof

These results follow from standard calculus. For $x > 0$,

$$\begin{aligned} f'(x) &= \frac{1}{\Gamma(k)} x^{k-2} e^{-x} [(k-1) - x] \\ f''(x) &= \frac{1}{\Gamma(k)} x^{k-3} e^{-x} [(k-1)(k-2) - 2(k-1)x + x^2] \end{aligned}$$

The special case $k = 1$ gives the *standard exponential distribution*. When $k \geq 1$, the distribution is unimodal.

In the simulation of the special distribution simulator, select the gamma distribution. Vary the shape parameter and note the shape of the density function. For various values of k , run the simulation 1000 times and compare the empirical density function to the true probability density function.

The distribution function and the quantile function do not have simple, closed representations for most values of the shape parameter. However, the distribution function has a trivial representation in terms of the incomplete and complete gamma functions.

The distribution function F of the standard gamma distribution with shape parameter $k \in (0, \infty)$ is given by

$$F(x) = \frac{\Gamma(k, x)}{\Gamma(k)}, \quad x \in (0, \infty) \quad (5.8.16)$$

Approximate values of the distribution and quantile functions can be obtained from special distribution calculator, and from most mathematical and statistical software packages.

Using the special distribution calculator, find the median, the first and third quartiles, and the interquartile range in each of the following cases:

1. $k = 1$
2. $k = 2$
3. $k = 3$

Moments

Suppose that X has the standard gamma distribution with shape parameter $k \in (0, \infty)$. The mean and variance are both simply the shape parameter.

The mean and variance of X are

1. $\mathbb{E}(X) = k$
2. $\text{var}(X) = k$

Proof

1. From the [fundamental identity](#),

$$\mathbb{E}(X) = \int_0^{\infty} x \frac{1}{\Gamma(k)} x^{k-1} e^{-x} dx = \frac{\Gamma(k+1)}{\Gamma(k)} = k \quad (5.8.17)$$

2. From the [fundamental identity](#) again

$$\mathbb{E}(X^2) = \int_0^{\infty} x^2 \frac{1}{\Gamma(k)} x^{k-1} e^{-x} dx = \frac{\Gamma(k+2)}{\Gamma(k)} = (k+1)k \quad (5.8.18)$$

and hence $\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2 = k$

In the simulation of the special distribution simulator, select the gamma distribution. Vary the shape parameter and note the size and location of the mean \pm standard deviation bar. For selected values of k , run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

More generally, the moments can be expressed easily in terms of the gamma function:

The moments of X are

1. $\mathbb{E}(X^a) = \Gamma(a+k)/\Gamma(k)$ if $a > -k$
2. $\mathbb{E}(X^n) = k^{[n]} = k(k+1) \cdots (k+n-1)$ if $n \in \mathbb{N}$

Proof

1. For $a > -k$,

$$\mathbb{E}(X^a) = \int_0^{\infty} x^a \frac{1}{\Gamma(k)} x^{k-1} e^{-x} dx = \frac{1}{\Gamma(k)} \int_0^{\infty} x^{a+k} e^{-x} dx = \frac{\Gamma(a+k)}{\Gamma(k)} \quad (5.8.19)$$

2. If $n \in \mathbb{N}$, then by the [fundamental identity](#), $\Gamma(k+n) = k(k+1) \cdots (k+n-1)\Gamma(k)$, so the result follows from (a).

Note also that $\mathbb{E}(X^a) = \infty$ if $a \leq -k$. We can now also compute the skewness and the kurtosis.

The skewness and kurtosis of X are

1. $\text{skew}(X) = \frac{2}{\sqrt{k}}$
2. $\text{kurt}(X) = 3 + \frac{6}{k}$

Proof

These results follows from the previous [moment results](#) and the computational formulas for skewness and kurtosis.

In particular, note that $\text{skew}(X) \rightarrow 0$ and $\text{kurt}(X) \rightarrow 3$ as $k \rightarrow \infty$. Note also that the *excess kurtosis* $\text{kurt}(X) - 3 \rightarrow 0$ as $k \rightarrow \infty$.

In the simulation of the special distribution simulator, select the gamma distribution. Increase the shape parameter and note the shape of the density function in light of the previous results on skewness and kurtosis. For various values of k , run the simulation 1000 times and compare the empirical density function to the true probability density function.

The following theorem gives the moment generating function.

The moment generating function of X is given by

$$\mathbb{E}(e^{tX}) = \frac{1}{(1-t)^k}, \quad t < 1 \quad (5.8.20)$$

Proof

For $t < 1$,

$$\mathbb{E}(e^{tX}) = \int_0^\infty e^{tx} \frac{1}{\Gamma(k)} x^{k-1} e^{-x} dx = \int_0^\infty \frac{1}{\Gamma(k)} x^{k-1} e^{-x(1-t)} dx \quad (5.8.21)$$

Substituting $u = x(1-t)$ so that $x = u/(1-t)$ and $dx = du/(1-t)$ gives

$$\mathbb{E}(e^{tX}) = \frac{1}{(1-t)^k} \int_0^\infty \frac{1}{\Gamma(k)} u^{k-1} e^{-u} du = \frac{1}{(1-t)^k} \quad (5.8.22)$$

The General Gamma Distribution

The gamma distribution is usually generalized by adding a scale parameter.

If Z has the standard gamma distribution with shape parameter $k \in (0, \infty)$ and if $b \in (0, \infty)$, then $X = bZ$ has the *gamma distribution* with *shape parameter* k and *scale parameter* b .

The reciprocal of the scale parameter, $r = 1/b$ is known as the *rate parameter*, particularly in the context of the Poisson process. The gamma distribution with parameters $k = 1$ and b is called the exponential distribution with scale parameter b (or rate parameter $r = 1/b$). More generally, when the shape parameter k is a positive integer, the gamma distribution is known as the *Erlang distribution*, named for the Danish mathematician Agner Erlang. The exponential distribution governs the time between arrivals in the Poisson model, while the Erlang distribution governs the actual arrival times.

Basic properties of the general gamma distribution follow easily from corresponding properties of the standard distribution and basic results for scale transformations.

Distribution Functions

Suppose that X has the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

X has probability density function f given by

$$f(x) = \frac{1}{\Gamma(k)b^k} x^{k-1} e^{-x/b}, \quad x \in (0, \infty) \quad (5.8.23)$$

Proof

Recall that if g is the [PDF of the standard gamma distribution](#) with shape parameter k then $f(x) = \frac{1}{b} g\left(\frac{x}{b}\right)$ for $x > 0$.

Recall that the inclusion of a scale parameter does not change the shape of the density function, but simply scales the graph horizontally and vertically. In particular, we have the same basic shapes as for the [standard gamma density function](#).

The probability density function f of X satisfies the following properties:

1. If $0 < k < 1$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
2. If $k = 1$, f is decreasing with $f(0) = 1$.
3. If $k > 1$, f increases and then decreases, with mode at $(k-1)b$.
4. If $0 < k \leq 1$, f is concave upward.
5. If $1 < k \leq 2$, f is concave downward and then upward, with inflection point at $b(k-1+\sqrt{k-1})$.
6. If $k > 2$, f is concave upward, then downward, then upward again, with inflection points at $b(k-1 \pm \sqrt{k-1})$.

In the simulation of the special distribution simulator, select the gamma distribution. Vary the shape and scale parameters and note the shape and location of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the true probability density function.

Once again, the distribution function and the quantile function do not have simple, closed representations for most values of the shape parameter. However, the distribution function has a simple representation in terms of the incomplete and complete gamma functions.

The distribution function F of X is given by

$$F(x) = \frac{\Gamma(k, x/b)}{\Gamma(k)}, \quad x \in (0, \infty) \quad (5.8.24)$$

Proof

From the [definition](#) we can take $X = bZ$ where Z has the standard gamma distribution with shape parameter k . Then $\mathbb{P}(X \leq x) = \mathbb{P}(Z \leq x/b)$ for $x \in (0, \infty)$, so the result follows from the [distribution function of \$Z\$](#) .

Approximate values of the distribution and quantile functions can be obtained from special distribution calculator, and from most mathematical and statistical software packages.

Open the special distribution calculator. Vary the shape and scale parameters and note the shape and location of the distribution and quantile functions. For selected values of the parameters, find the median and the first and third quartiles.

Moments

Suppose again that X has the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

The mean and variance of X are

1. $\mathbb{E}(X) = bk$
2. $\text{var}(X) = b^2k$

Proof

From the [definition](#), we can take $X = bZ$ where Z has the standard gamma distribution with shape parameter k . Then using the [mean and variance of \$Z\$](#) ,

1. $\mathbb{E}(X) = b\mathbb{E}(Z) = bk$
2. $\text{var}(X) = b^2\text{var}(Z) = b^2k$

In the special distribution simulator, select the gamma distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The moments of X are

1. $\mathbb{E}(X^a) = b^a \Gamma(a+k) / \Gamma(k)$ for $a > -k$
2. $\mathbb{E}(X^n) = b^n k^{[n]} = b^n k(k+1) \cdots (k+n-1)$ if $n \in \mathbb{N}$

Proof

Again, from the [definition](#), we can take $X = bZ$ where Z has the standard gamma distribution with shape parameter k . The results follow from the [moment results for \$Z\$](#) , since $E(X^a) = b^a E(Z^a)$.

Note also that $E(X^a) = \infty$ if $a \leq -k$. Recall that skewness and kurtosis are defined in terms of the standard score, and hence are unchanged by the addition of a scale parameter.

The skewness and kurtosis of X are

1. $\text{skew}(X) = \frac{2}{\sqrt{k}}$
2. $\text{kurt}(X) = 3 + \frac{6}{k}$

The moment generating function of X is given by

$$\mathbb{E}(e^{tX}) = \frac{1}{(1-bt)^k}, \quad t < \frac{1}{b} \quad (5.8.25)$$

Proof

From the [definition](#), we can take $X = bZ$ where Z has the standard gamma distribution with shape parameter k . Then $\mathbb{E}(e^{tX}) = \mathbb{E}[e^{(tb)Z}]$, so the result follows from the [moment generating function of \$Z\$](#) .

Relations

Our first result is simply a restatement of the meaning of the term *scale parameter*.

Suppose that X has the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. If $c \in (0, \infty)$, then cX has the gamma distribution with shape parameter k and scale parameter bc .

Proof

From the [definition](#), we can take $X = bZ$ where Z has the standard gamma distribution with shape parameter k . Then $cX = cbZ$.

More importantly, if the scale parameter is fixed, the gamma family is closed with respect to sums of independent variables.

Suppose that X_1 and X_2 are independent random variables, and that X_i has the gamma distribution with shape parameter $k_i \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ for $i \in \{1, 2\}$. Then $X_1 + X_2$ has the gamma distribution with shape parameter $k_1 + k_2$ and scale parameter b .

Proof

Recall that the MGF of $X = X_1 + X_2$ is the product of the MGFs of X_1 and X_2 , so

$$\mathbb{E}(e^{tX}) = \frac{1}{(1-bt)^{k_1}} \frac{1}{(1-bt)^{k_2}} = \frac{1}{(1-bt)^{k_1+k_2}}, \quad t < \frac{1}{b} \quad (5.8.26)$$

From the previous result, it follows that the gamma distribution is infinitely divisible:

Suppose that X has the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. For $n \in \mathbb{N}_+$, X has the same distribution as $\sum_{i=1}^n X_i$, where (X_1, X_2, \dots, X_n) is a sequence of independent random variables, each with the gamma distribution with shape parameter k/n and scale parameter b .

From the sum result and the central limit theorem, it follows that if k is large, the gamma distribution with shape parameter k and scale parameter b can be approximated by the normal distribution with mean kb and variance kb^2 . Here is the precise statement:

Suppose that X_k has the gamma distribution with shape parameter $k \in (0, \infty)$ and fixed scale parameter $b \in (0, \infty)$. Then the distribution of the standardized variable below converges to the standard normal distribution as $k \rightarrow \infty$:

$$Z_k = \frac{X_k - kb}{\sqrt{kb}} \quad (5.8.27)$$

In the special distribution simulator, select the gamma distribution. For various values of the scale parameter, increase the shape parameter and note the increasingly “normal” shape of the density function. For selected values of the parameters, run the experiment 1000 times and compare the empirical density function to the true probability density function.

The gamma distribution is a member of the general exponential family of distributions:

The gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ is a two-parameter exponential family with natural parameters $(k-1, -1/b)$, and natural statistics $(\ln X, X)$.

Proof

This follows from the definition of the general exponential family. The gamma PDF can be written as

$$f(x) = \frac{1}{b^k \Gamma(k)} \exp \left[(k-1) \ln x - \frac{1}{b} x \right], \quad x \in (0, \infty) \quad (5.8.28)$$

For $n \in (0, \infty)$, the gamma distribution with shape parameter $n/2$ and scale parameter 2 is known as the *chi-square distribution* with n degrees of freedom. The chi-square distribution is important enough to deserve a separate section.

Computational Exercise

Suppose that the lifetime of a device (in hours) has the gamma distribution with shape parameter $k = 4$ and scale parameter $b = 100$.

1. Find the probability that the device will last more than 300 hours.
2. Find the mean and standard deviation of the lifetime.

Answer

Let X denote the lifetime in hours.

1. $\mathbb{P}(X > 300) = 13e^{-3} \approx 0.6472$
2. $\mathbb{E}(X) = 400, \text{sd}(X) = 200$

Suppose that Y has the gamma distribution with parameters $k = 10$ and $b = 2$. For each of the following, compute the true value using the special distribution calculator and then compute the normal approximation. Compare the results.

1. $\mathbb{P}(18 < X < 25)$
2. The 80th percentile of Y

Answer

1. $\mathbb{P}(18 < X < 25) = 0.3860, \mathbb{P}(18 < X < 25) \approx 0.4095$
2. $y_{0.8} = 25.038, y_{0.8} \approx 25.325$

This page titled [5.8: The Gamma Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.9: Chi-Square and Related Distribution

In this section we will study a distribution, and some relatives, that have special importance in statistics. In particular, the chi-square distribution will arise in the study of the sample variance when the underlying distribution is normal and in goodness of fit tests.

The Chi-Square Distribution

Distribution Functions

For $n \in (0, \infty)$, the gamma distribution with shape parameter $n/2$ and scale parameter 2 is called the *chi-square distribution* with n degrees of freedom. The probability density function f is given by

$$f(x) = \frac{1}{2^{n/2}\Gamma(n/2)} x^{n/2-1} e^{-x/2}, \quad x \in (0, \infty) \quad (5.9.1)$$

So the chi-square distribution is a continuous distribution on $(0, \infty)$. For reasons that will be clear later, n is usually a positive integer, although technically this is not a mathematical requirement. When n is a positive integer, the gamma function in the normalizing constant can be given explicitly.

If $n \in \mathbb{N}_+$ then

1. $\Gamma(n/2) = (n/2 - 1)!$ if n is even.
2. $\Gamma(n/2) = \frac{(n-1)!}{2^{n-1}(n/2-1/2)!} \sqrt{\pi}$ if n is odd.

The chi-square distribution has a rich collection of shapes.

The chi-square probability density function with $n \in (0, \infty)$ degrees of freedom satisfies the following properties:

1. If $0 < n < 2$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
2. If $n = 2$, f is decreasing with $f(0) = \frac{1}{2}$.
3. If $n > 2$, f increases and then decreases with mode at $n - 2$.
4. If $0 < n \leq 2$, f is concave downward.
5. If $2 < n \leq 4$, f is concave downward and then upward, with inflection point at $n - 2 + \sqrt{2n - 4}$
6. If $n > 4$ then f is concave upward then downward and then upward again, with inflection points at $n - 2 \pm \sqrt{2n - 4}$

In the special distribution simulator, select the chi-square distribution. Vary n with the scroll bar and note the shape of the probability density function. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.

The distribution function and the quantile function do not have simple, closed-form representations for most values of the parameter. However, the distribution function can be given in terms of the complete and incomplete gamma functions.

Suppose that X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom. The distribution function F of X is given by

$$F(x) = \frac{\Gamma(n/2, x/2)}{\Gamma(n/2)}, \quad x \in (0, \infty) \quad (5.9.2)$$

Approximate values of the distribution and quantile functions can be obtained from the special distribution calculator, and from most mathematical and statistical software packages.

In the special distribution calculator, select the chi-square distribution. Vary the parameter and note the shape of the probability density, distribution, and quantile functions. In each of the following cases, find the median, the first and third quartiles, and the interquartile range.

1. $n = 1$
2. $n = 2$
3. $n = 5$
4. $n = 10$

Moments

The mean, variance, moments, and moment generating function of the chi-square distribution can be obtained easily from general results for the gamma distribution.

If X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom then

1. $\mathbb{E}(X) = n$
2. $\text{var}(X) = 2n$

In the simulation of the special distribution simulator, select the chi-square distribution. Vary n with the scroll bar and note the size and location of the mean \pm standard deviation bar. For selected values of n , run the simulation 1000 times and compare the empirical moments to the distribution moments.

The skewness and kurtosis of the chi-square distribution are given next.

If X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom, then

1. $\text{skew}(X) = 2\sqrt{2/n}$
2. $\text{kurt}(X) = 3 + 12/n$

Note that $\text{skew}(X) \rightarrow 0$ and $\text{kurt}(X) \rightarrow 3$ as $n \rightarrow \infty$. In particular, the *excess kurtosis* $\text{kurt}(X) - 3 \rightarrow 0$ as $n \rightarrow \infty$.

In the simulation of the special distribution simulator, select the chi-square distribution. Increase n with the scroll bar and note the shape of the probability density function in light of the previous results on skewness and kurtosis. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.

The next result gives the general moments of the chi-square distribution.

If X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom, then for $k > -n/2$,

$$\mathbb{E}(X^k) = 2^k \frac{\Gamma(n/2 + k)}{\Gamma(n/2)} \quad (5.9.3)$$

In particular, if $k \in \mathbb{N}_+$ then

$$\mathbb{E}(X^k) = 2^k \left(\frac{n}{2}\right) \left(\frac{n}{2} + 1\right) \cdots \left(\frac{n}{2} + k - 1\right) \quad (5.9.4)$$

Note also $\mathbb{E}(X^k) = \infty$ if $k \leq -n/2$.

If X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom, then X has moment generating function

$$\mathbb{E}(e^{tX}) = \frac{1}{(1 - 2t)^{n/2}}, \quad t < \frac{1}{2} \quad (5.9.5)$$

Relations

The chi-square distribution is connected to a number of other special distributions. Of course, the most important relationship is the definition—the chi-square distribution with n degrees of freedom is a special case of the gamma distribution, corresponding to shape parameter $n/2$ and scale parameter 2. On the other hand, any gamma distributed variable can be re-scaled into a variable with a chi-square distribution.

If X has the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ then $Y = \frac{2}{b}X$ has the chi-square distribution with $2k$ degrees of freedom.

Proof

Since the gamma distribution is a scale family, Y has a gamma distribution with shape parameter k and scale parameter $b\frac{2}{b} = 2$. Hence Y has the chi-square distribution with $2k$ degrees of freedom.

The chi-square distribution with 2 degrees of freedom is the exponential distribution with scale parameter 2.

Proof

The chi-square distribution with 2 degrees of freedom is the gamma distribution with shape parameter 1 and scale parameter 2, which we already know is the exponential distribution with scale parameter 2.

If Z has the standard normal distribution then $X = Z^2$ has the chi-square distribution with 1 degree of freedom.

Proof

As usual, let ϕ and Φ denote the PDF and CDF of the standard normal distribution, respectively. Then for $x > 0$,

$$\mathbb{P}(X \leq x) = \mathbb{P}(-\sqrt{x} \leq Z \leq \sqrt{x}) = 2\Phi(\sqrt{x}) - 1 \quad (5.9.6)$$

Differentiating with respect to x gives the density function f of X :

$$f(x) = \phi(\sqrt{x})x^{-1/2} = \frac{1}{\sqrt{2\pi}}x^{-1/2}e^{-x/2}, \quad x \in (0, \infty) \quad (5.9.7)$$

which we recognize as the chi-square PDF with 1 degree of freedom.

Recall that if we add independent gamma variables with a common scale parameter, the resulting random variable also has a gamma distribution, with the common scale parameter and with shape parameter that is the sum of the shape parameters of the terms. Specializing to the chi-square distribution, we have the following important result:

If X has the chi-square distribution with $m \in (0, \infty)$ degrees of freedom, Y has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom, and X and Y are independent, then $X + Y$ has the chi-square distribution with $m + n$ degrees of freedom.

The last two results lead to the following theorem, which is fundamentally important in statistics.

Suppose that $n \in \mathbb{N}_+$ and that (Z_1, Z_2, \dots, Z_n) is a sequence of independent standard normal variables. Then the sum of the squares

$$V = \sum_{i=1}^n Z_i^2 \quad (5.9.8)$$

has the chi-square distribution with n degrees of freedom:

This theorem is the reason that the chi-square distribution deserves a name of its own, and the reason that the degrees of freedom parameter is usually a positive integer. Sums of squares of independent normal variables occur frequently in statistics.

From the central limit theorem, and previous results for the gamma distribution, it follows that if n is large, the chi-square distribution with n degrees of freedom can be approximated by the normal distribution with mean n and variance $2n$. Here is the precise statement:

If X_n has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom, then the distribution of the standard score

$$Z_n = \frac{X_n - n}{\sqrt{2n}} \quad (5.9.9)$$

converges to the standard normal distribution as $n \rightarrow \infty$.

In the simulation of the special distribution simulator, select the chi-square distribution. Start with $n = 1$ and increase n . Note the shape of the probability density function in light of the previous theorem. For selected values of n , run the experiment 1000 times and compare the empirical density function to the true density function.

Like the gamma distribution, the chi-square distribution is infinitely divisible:

Suppose that X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom. For $k \in \mathbb{N}_+$, X has the same distribution as $\sum_{i=1}^k X_i$, where (X_1, X_2, \dots, X_k) is a sequence of independent random variables, each with the chi-square distribution with n/k degrees of freedom.

Also like the gamma distribution, the chi-square distribution is a member of the general exponential family of distributions:

The chi-square distribution with $n \in (0, \infty)$ degrees of freedom is a one-parameter exponential family with natural parameter $n/2 - 1$, and natural statistic $\ln X$.

Proof

This follows from the definition of the general exponential family. The PDF can be written as

$$f(x) = \frac{e^{-x/2}}{2^{n/2}\Gamma(n/2)} \exp[(n/2 - 1) \ln x], \quad x \in (0, \infty) \quad (5.9.10)$$

The Chi Distribution

The chi distribution, appropriately enough, is the distribution of the square root of a variable with the chi-square distribution

Suppose that X has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom. Then $U = \sqrt{X}$ has the *chi distribution* with n degrees of freedom.

So like the chi-square distribution, the chi distribution is a continuous distribution on $(0, \infty)$.

Distribution Functions

The distribution function G of the chi distribution with $n \in (0, \infty)$ degrees of freedom is given by

$$G(u) = \frac{\Gamma(n/2, u^2/2)}{\Gamma(n/2)}, \quad u \in (0, \infty) \quad (5.9.11)$$

Proof

Suppose that U has the chi distribution with n degrees of freedom so that $X = U^2$ has the chi-square distribution with n degrees of freedom. For $u \in (0, \infty)$,

$$G(u) = \mathbb{P}(U \leq u) = \mathbb{P}(U^2 \leq u^2) = \mathbb{P}(X \leq u^2) = F(u^2) \quad (5.9.12)$$

where F is the [chi-square distribution function](#) with n degrees of freedom.

The probability density function g of the chi distribution with $n \in (0, \infty)$ degrees of freedom is given by

$$g(u) = \frac{1}{2^{n/2-1}\Gamma(n/2)} u^{n-1} e^{-u^2/2}, \quad u \in (0, \infty) \quad (5.9.13)$$

Proof

Suppose again that U has the chi distribution with n degrees of freedom so that $X = U^2$ has the chi-square distribution with n degrees of freedom. The transformation $u = \sqrt{x}$ maps $(0, \infty)$ one-to-one onto $(0, \infty)$. The inverse transformation is $x = u^2$ with $dx/du = 2u$. Hence by the standard change of variables formula,

$$g(u) = f(x) \frac{dx}{du} = f(u^2) 2u \quad (5.9.14)$$

where f is the [chi-square PDF](#).

The chi probability density function also has a variety of shapes.

The chi probability density function with $n \in (0, \infty)$ degrees of freedom satisfies the following properties:

1. If $0 < n < 1$, g is decreasing with $g(u) \rightarrow \infty$ as $u \downarrow 0$.
2. If $n = 1$, g is decreasing with $g(0) = \sqrt{2/\pi}$ as $u \downarrow 0$.
3. If $n > 1$, g increases and then decreases with mode $u = \sqrt{n-1}$
4. If $0 < n < 1$, g is concave upward.
5. If $1 \leq n \leq 2$, g is concave downward and then upward with inflection point at $u = \sqrt{\frac{1}{2}[2n-1 + \sqrt{8n-7}]}$
6. If $n > 2$, g is concave upward then downward then upward again with inflection points at $u = \sqrt{\frac{1}{2}[2n-1 \pm \sqrt{8n-7}]}$

Moments

The raw moments of the chi distribution are easy to compute in terms of the gamma function.

Suppose that U has the chi distribution with $n \in (0, \infty)$ degrees of freedom. Then

$$\mathbb{E}(U^k) = 2^{k/2} \frac{\Gamma[(n+k)/2]}{\Gamma(n/2)}, \quad k \in (0, \infty) \quad (5.9.15)$$

Proof

By definition

$$E(U^k) = \int_0^\infty u^k g(u) du = \frac{1}{2^{n/2-1} \Gamma(n/2)} \int_0^\infty u^{n+k-1} e^{-u^2/2} du \quad (5.9.16)$$

The change of variables $v = u^2/2$, so that $u = 2^{1/2}v^{1/2}$ and $du = 2^{-1/2}v^{-1/2} dv$ gives (after simplification)

$$E(U^k) = \frac{2^{k/2}}{\Gamma(n/2)} \int_0^\infty v^{(n+k)/2-1} e^{-v} dv \quad (5.9.17)$$

The last integral is $\Gamma[(n+k)/2]$.

Curiously, the second moment is simply the degrees of freedom parameter.

Suppose again that U has the chi distribution with $n \in (0, \infty)$ degrees of freedom. Then

1. $\mathbb{E}(U) = 2^{1/2} \frac{\Gamma[(n+1)/2]}{\Gamma(n/2)}$
2. $\mathbb{E}(U^2) = n$
3. $\text{var}(U) = n - 2 \frac{\Gamma^2[(n+1)/2]}{\Gamma^2(n/2)}$

Proof

For part (b), using the fundamental identity of the gamma function we have

$$\mathbb{E}(U^2) = 2 \frac{\Gamma(n/2 + 1)}{\Gamma(n/2)} = 2 \frac{(n/2) \Gamma(n/2)}{\Gamma(n/2)} = n \quad (5.9.18)$$

The other parts follow from direct substitution.

Relations

The fundamental relationship of course is the one between the chi distribution and the chi-square distribution given in the [definition](#). In turn, this leads to a fundamental relationship between the chi distribution and the normal distribution.

Suppose that $n \in \mathbb{N}_+$ and that (Z_1, Z_2, \dots, Z_n) is a sequence of independent variables, each with the standard normal distribution. Then

$$U = \sqrt{Z_1^2 + Z_2^2 + \cdots + Z_n^2} \quad (5.9.19)$$

has the chi distribution with n degrees of freedom.

Note that the random variable U in the last result is the standard Euclidean norm of (Z_1, Z_2, \dots, Z_n) , thought of as a vector in \mathbb{R}^n . Note also that the chi distribution with 1 degree of freedom is the distribution of $|Z|$, the absolute value of a standard normal variable, which is known as the standard half-normal distribution.

The Non-Central Chi-Square Distribution

Much of the importance of the chi-square distribution stems from the fact that it is the distribution that governs the sum of squares of independent, standard normal variables. A natural generalization, and one that is important in statistical applications, is to consider the distribution of a sum of squares of independent normal variables, each with variance 1 but with different means.

Suppose that $n \in \mathbb{N}_+$ and that (X_1, X_2, \dots, X_n) is a sequence of independent variables, where X_k has the normal distribution with mean $\mu_k \in \mathbb{R}$ and variance 1 for $k \in \{1, 2, \dots, n\}$. The distribution of $Y = \sum_{k=1}^n X_k^2$ is the *non-central chi-square distribution* with n degrees of freedom and non-centrality parameter $\lambda = \sum_{k=1}^n \mu_k^2$.

Note that the degrees of freedom is a positive integer while the non-centrality parameter $\lambda \in [0, \infty)$, but we will soon generalize the degrees of freedom.

Distribution Functions

Like the chi-square and chi distributions, the non-central chi-square distribution is a continuous distribution on $(0, \infty)$. The probability density function and distribution function do not have simple, closed expressions, but there is a fascinating connection to the Poisson distribution. To set up the notation, let f_k and F_k denote the probability density and distribution functions of the chi-square distribution with $k \in (0, \infty)$ degrees of freedom. Suppose that Y has the non-central chi-square distribution with $n \in \mathbb{N}_+$ degrees of freedom and non-centrality parameter $\lambda \in [0, \infty)$. The following fundamental theorem gives the probability density function of Y as an infinite series, and shows that the distribution does in fact depend only on n and λ .

The probability density function g of Y is given by

$$g(y) = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} f_{n+2k}(y), \quad y \in (0, \infty) \quad (5.9.20)$$

Proof

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent random variables, where X_i has the normal distribution with mean μ_i and variance 1, and where $\lambda = \sum_{i=1}^n \mu_i^2$. So by definition, $Y = \sum_{i=1}^n X_i^2$ has the non-central chi-square distribution with n degrees of freedom and non-centrality parameter λ . The random vector \mathbf{X} has a multivariate normal distribution with mean vector $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_n)$ and variance-covariance matrix I (the $n \times n$ identity matrix). The (joint) PDF h of \mathbf{X} is symmetric about $\boldsymbol{\mu}$: $h(\boldsymbol{\mu} - \mathbf{x}) = h(\boldsymbol{\mu} + \mathbf{x})$ for $\mathbf{x} \in \mathbb{R}^n$. Because of this symmetry, the distribution of Y depends on $\boldsymbol{\mu}$ only through the parameter λ . It follows that Y has the same distribution as $\sum_{i=1}^n U_i^2$ where (U_1, U_2, \dots, U_n) are independent, U_1 has the normal distribution with mean $\sqrt{\lambda}$ and variance 1, and (U_2, U_3, \dots, U_n) are standard normal.

The distribution of U_1^2 is found by the usual change of variables methods. Let ϕ and Φ denote the standard normal PDF and CDF, respectively, so that U_1 has CDF given by $\mathbb{P}(U_1 \leq x) = \Phi(x - \sqrt{\lambda})$ for $x \in \mathbb{R}$. Thus,

$$\mathbb{P}(U_1^2 \leq x) = \mathbb{P}(-\sqrt{x} \leq U_1 \leq \sqrt{x}) = \Phi(\sqrt{x} - \sqrt{\lambda}) - \Phi(-\sqrt{x} - \sqrt{\lambda}), \quad x \in (0, \infty) \quad (5.9.21)$$

Taking derivatives, the PDF g of U_1^2 is given by

$$g(x) = \frac{1}{2\sqrt{x}} [\phi(\sqrt{x} - \sqrt{\lambda}) + \phi(-\sqrt{x} - \sqrt{\lambda})] \frac{1}{2\sqrt{x}} [\phi(\sqrt{x} - \sqrt{\lambda}) + \phi(\sqrt{x} + \sqrt{\lambda})], \quad x \in (0, \infty) \quad (5.9.22)$$

But $\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ for $z \in \mathbb{R}$, so substituting and simplifying gives

$$g(x) = \frac{1}{\sqrt{2\pi x}} e^{-\frac{1}{2}(x+\lambda)} \frac{1}{2} (e^{\sqrt{\lambda x}} + e^{-\sqrt{\lambda x}}) = \frac{1}{\sqrt{2\pi x}} e^{-\frac{1}{2}(x+\lambda)} \cosh(\sqrt{\lambda x}), \quad x \in (0, \infty) \quad (5.9.23)$$

Next, recall that the Taylor series for the hyperbolic cosine function is

$$\cosh(x) = \sum_{k=0}^{\infty} \frac{x^{2k}}{(2k)!}, \quad x \in \mathbb{R} \quad (5.9.24)$$

which leads to

$$g(x) = \sum_{k=0}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(x+\lambda)} \frac{\lambda^k x^{k-1/2}}{(2k)!}, \quad x \in (0, \infty) \quad (5.9.25)$$

After a bit more algebra, we get the representation in the theorem, with $n = 1$. That is,

$$g(x) = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} \frac{1}{2^{(2k+1)/2} \Gamma[(2k+1)/2]} x^{(2k+1)/2-1} e^{-x/2}, \quad x \in (0, \infty) \quad (5.9.26)$$

Or in functional form, $g = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} f_{2k+1}$.

To complete the proof, we know that $\sum_{j=2}^n U_j^2$ has the chi-square distribution with $n-1$ degrees of freedom, and hence has PDF f_{n-1} , and is independent of U_1 . Therefore the distribution of $\sum_{j=1}^n U_j^2$ is

$$g * f_{n-1} = \left(\sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} f_{2k+1} \right) * f_{n-1} = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} f_{2k+n} \quad (5.9.27)$$

where $*$ denotes convolution as usual, and where we have used the fundamental result above on the [sum of independent chi-square variables](#).

The function $k \mapsto e^{-\lambda/2} \frac{(\lambda/2)^k}{k!}$ on \mathbb{N} is the probability density function of the Poisson distribution with parameter $\lambda/2$. So it follows that if N has the Poisson distribution with parameter $\lambda/2$ and the conditional distribution of Y given N is chi-square with parameter $n+2N$, then Y has the distribution discussed here—non-central chi-square with n degrees of freedom and non-centrality parameter λ . Moreover, it's clear that g is a valid probability density function for any $n \in (0, \infty)$, so we can generalize our definition a bit.

For $n \in (0, \infty)$ and $\lambda \in [0, \infty)$, the distribution with probability density function g above is the *non-central chi-square distribution* with n degrees of freedom and non-centrality parameter λ .

The distribution function G is given by

$$G(y) = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} F_{n+2k}(y), \quad y \in (0, \infty) \quad (5.9.28)$$

Proof

This follows immediately from the result for the [PDF](#), since $G(0) = 0$ and $G' = g$.

Moments

In this discussion, we assume again that Y has the non-central chi-square distribution with $n \in (0, \infty)$ degrees of freedom and non-centrality parameter $\lambda \in [0, \infty)$.

The moment generating function M of Y is given by

$$M(t) = \mathbb{E}(e^{tY}) = \frac{1}{(1-2t)^{n/2}} \exp\left(\frac{\lambda t}{1-2t}\right), \quad t \in (-\infty, 1/2) \quad (5.9.29)$$

Proof

We will use the fundamental relationship mentioned above. Thus, suppose that N has the Poisson distribution with parameter $\lambda/2$, and that given N , Y has the chi-square distribution with $n+2N$ degrees of freedom. Conditioning and using the [MGF of the chi-square distribution](#) above gives

$$E(e^{tY}) = E[E(e^{tY} | N)] = E\left(\frac{1}{(1-2t)^{(n+2N)/2}}\right) = \frac{1}{(1-2t)^{n/2}} E\left[\left(\frac{1}{1-2t}\right)^N\right] \quad (5.9.30)$$

The last expected value is the probability generating function of N , evaluated at $\frac{1}{1-2t}$. Hence

$$E(e^{tY}) = \frac{1}{1-2t} \exp\left[\frac{\lambda}{2} \left(\frac{1}{1-2t} - 1\right)\right] = \frac{1}{(1-2t)^{n/2}} \exp\left(\frac{\lambda t}{1-2t}\right) \quad (5.9.31)$$

The mean and variance of Y are

1. $E(Y) = n + \lambda$
2. $\text{var}(Y) = 2(n + 2\lambda)$

Proof

These results can be obtained by taking derivatives of the MGF, but the derivation using the connection with the Poisson distribution is more interesting. So suppose again that N has the Poisson distribution with parameter $\lambda/2$ and that the conditional distribution of Y given N is chi-square with $n + 2N$ degrees of freedom. Conditioning and using the means and variances of the chi-square and Poisson distributions, we have

1. $E(Y) = E[E(Y | N)] = E(n + 2N) = n + 2(\lambda/2) = n + \lambda$
2. $\text{var}(Y) = E[\text{var}(Y | N)] + \text{var}[E(Y | N)] = E[2(n + 2N)] + \text{var}(n + 2N) = 2n + 4(\lambda/2) + 4\lambda/2 = 2n + 4\lambda$

The skewness and kurtosis of Y are

1. $\text{skew}(Y) = 2^{3/2} \frac{n+3\lambda}{(n+2\lambda)^{3/2}}$
2. $\text{kurt}(Y) = 3 + 12 \frac{n+4\lambda}{(n+2\lambda)^2}$

Note that $\text{skew}(Y) \rightarrow 0$ as $n \rightarrow \infty$ or as $\lambda \rightarrow \infty$. Note also that the excess kurtosis is $\text{kurt}(Y) - 3 = 12 \frac{n+4\lambda}{(n+2\lambda)^2}$. So $\text{kurt}(Y) \rightarrow 3$ (the kurtosis of the normal distribution) as $n \rightarrow \infty$ or as $\lambda \rightarrow \infty$.

Relations

Trivially of course, the ordinary chi-square distribution is a special case of the non-central chi-square distribution, with non-centrality parameter 0. The most important relation is the original [definition above](#). The non-central chi-square distribution with $n \in \mathbb{N}_+$ degrees of freedom and non-centrality parameter $\lambda \in [0, \infty)$ is the distribution of the sum of the squares of n independent normal variables with variance 1 and whose means satisfy $\sum_{k=1}^n \mu_k^2 = \lambda$. The next most important relation is the one that arose in the probability density function and was so useful for computing moments. We state this one again for emphasis.

Suppose that N has the Poisson distribution with parameter $\lambda/2$, where $\lambda \in (0, \infty)$, and that the conditional distribution of Y given N is chi-square with $n + 2N$ degrees of freedom, where $n \in (0, \infty)$. Then the (unconditional) distribution of Y is non-central chi-square with n degree of freedom and non-centrality parameter λ .

Proof

For $j \in \mathbb{N}_+$, let f_j denote the chi-square PDF with j degrees of freedom. Then from the assumptions, the PDF g of Y is given by

$$g(y) = \sum_{n=0}^{\infty} \mathbb{P}(N = k) f_{n+2k}(y) = \sum_{n=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} f_{n+2k}(y), \quad y \in (0, \infty) \quad (5.9.32)$$

which is the PDF of the non-central chi-square distribution with n degrees of freedom and non-centrality parameter λ , derived above.

As the asymptotic results for the skewness and kurtosis suggest, there is also a central limit theorem.

Suppose that Y has the non-central chi-square distribution with $n \in (0, \infty)$ degrees of freedom and non-centrality parameter $\lambda \in (0, \infty)$. Then the distribution of the standard score

$$\frac{Y - (n + \lambda)}{\sqrt{2(n + 2\lambda)}} \quad (5.9.33)$$

converges to the standard normal distribution as $n \rightarrow \infty$ or as $\lambda \rightarrow \infty$.

Computational Exercises

Suppose that a missile is fired at a target at the origin of a plane coordinate system, with units in meters. The missile lands at (X, Y) where X and Y are independent and each has the normal distribution with mean 0 and variance 100. The missile will destroy the target if it lands within 20 meters of the target. Find the probability of this event.

Answer

Let Z denote the distance from the missile to the target. $\mathbb{P}(Z < 20) = 1 - e^{-2} \approx 0.8647$

Suppose that X has the chi-square distribution with $n = 18$ degrees of freedom. For each of the following, compute the true value using the special distribution calculator and then compute the normal approximation. Compare the results.

1. $\mathbb{P}(15 < X < 20)$
2. The 75th percentile of X .

Answer

1. $\mathbb{P}(15 < X < 20) = 0.3252, \mathbb{P}(15 < X < 20) \approx 0.3221$
2. $x_{0.75} = 21.605, x_{0.75} \approx 22.044$

This page titled [5.9: Chi-Square and Related Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.10: The Student t Distribution

In this section we will study a distribution that has special importance in statistics. In particular, this distribution will arise in the study of a standardized version of the sample mean when the underlying distribution is normal.

Basic Theory

Definition

Suppose that Z has the standard normal distribution, V has the chi-squared distribution with $n \in (0, \infty)$ degrees of freedom, and that Z and V are independent. Random variable

$$T = \frac{Z}{\sqrt{V/n}} \quad (5.10.1)$$

has the *student t distribution* with n degrees of freedom.

The student t distribution is well defined for any $n > 0$, but in practice, only positive integer values of n are of interest. This distribution was first studied by William Gosset, who published under the pseudonym *Student*.

Distribution Functions

Suppose that T has the t distribution with $n \in (0, \infty)$ degrees of freedom. Then T has a continuous distribution on \mathbb{R} with probability density function f given by

$$f(t) = \frac{\Gamma[(n+1)/2]}{\sqrt{n\pi} \Gamma(n/2)} \left(1 + \frac{t^2}{n}\right)^{-(n+1)/2}, \quad t \in \mathbb{R} \quad (5.10.2)$$

Proof

For $v > 0$, the conditional distribution of T given $V = v$ is normal with mean 0 and variance n/v . By definition, V has the chi-square distribution with n degrees of freedom. Hence, the joint PDF of (T, V) is

$$g(t, v) = \sqrt{\frac{v}{2\pi n}} e^{-vz^2/2n} \frac{1}{2^{n/2} \Gamma(n/2)} v^{n/2-1} e^{-v/2} = \frac{1}{2^{(n+1)/2} \sqrt{n\pi} \Gamma(n/2)} v^{(n+1)/2-1} e^{-v(1+t^2/n)/2}, \quad t \in \mathbb{R}, v \in (0, \infty) \quad (5.10.3)$$

The PDF of T is

$$f(t) = \int_0^\infty g(t, v) dv = \frac{1}{2^{(n+1)/2} \sqrt{n\pi} \Gamma(n/2)} \int_0^\infty v^{(n+1)/2-1} e^{-v(1+t^2/n)/2} dv, \quad t \in \mathbb{R} \quad (5.10.4)$$

Except for the missing normalizing constant, the integrand is the gamma PDF with shape parameter $(n+1)/2$ and scale parameter $2/(1+t^2/n)$. Hence

$$f(t) = \frac{1}{2^{(n+1)/2} \sqrt{n\pi} \Gamma(n/2)} \Gamma[(n+1)/2] \left(\frac{2}{1+t^2/n}\right)^{(n+1)/2}, \quad t \in \mathbb{R} \quad (5.10.5)$$

Simplifying gives the result.

The proof of this theorem provides a good way of thinking of the t distribution: the distribution arises when the variance of a mean 0 normal distribution is randomized in a certain way.

In the special distribution simulator, select the student t distribution. Vary n and note the shape of the probability density function. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.

The Student **probability density function** f with $n \in (0, \infty)$ degrees of freedom has the following properties:

1. f is symmetric about $t = 0$.
2. f is increasing and then decreasing with mode $t = 0$.
3. f is concave upward, then downward, then upward again with inflection points at $\pm\sqrt{n/(n+1)}$.
4. $f(t) \rightarrow 0$ as $t \rightarrow \infty$ and as $t \rightarrow -\infty$.

In particular, the distribution is unimodal with mode and median at $t = 0$. Note also that the inflection points converge to ± 1 as $n \rightarrow \infty$.

The distribution function and the quantile function of the general t distribution do not have simple, closed-form representations. Approximate values of these functions can be obtained from the special distribution calculator, and from most mathematical and statistical software packages.

In the special distribution calculator, select the student distribution. Vary the parameter and note the shape of the probability density, distribution, and quantile functions. In each of the following cases, find the first and third quartiles:

1. $n = 2$
2. $n = 5$
3. $n = 10$
4. $n = 20$

Moments

Suppose that T has a t distribution. The representation in the [definition](#) can be used to find the mean, variance and other moments of T . The main point to remember in the proofs that follow is that since V has the chi-square distribution with n degrees of freedom, $E(V^k) = \infty$ if $k \leq -\frac{n}{2}$, while if $k > -\frac{n}{2}$,

$$\mathbb{E}(V^k) = 2^k \frac{\Gamma(k + n/2)}{\Gamma(n/2)} \quad (5.10.6)$$

Suppose that T has the t distribution with $n \in (0, \infty)$ degrees of freedom. Then

1. $\mathbb{E}(T)$ is undefined if $0 < n \leq 1$
2. $\mathbb{E}(T) = 0$ if $1 < n < \infty$

Proof

By independence, $\mathbb{E}(T) = \sqrt{n} \mathbb{E}(V^{-1/2}) \mathbb{E}(Z)$. Of course $\mathbb{E}(Z) = 0$. On the other hand, $\mathbb{E}(V^{-1/2}) = \infty$ if $n \leq 1$ and $\mathbb{E}(V^{-1/2}) < \infty$ if $n > 1$.

Suppose again that T has the t distribution with $n \in (0, \infty)$ degrees of freedom then

1. $\text{var}(T)$ is undefined if $0 < n \leq 1$
2. $\text{var}(T) = \infty$ if $1 < n \leq 2$
3. $\text{var}(T) = \frac{n}{n-2}$ if $2 < n < \infty$

Proof

By independence, $\mathbb{E}(T^2) = n \mathbb{E}(Z^2) \mathbb{E}(V^{-1})$. Of course $\mathbb{E}(Z^2) = 1$. On the other hand, $\mathbb{E}(V^{-1}) = \infty$ if $n \leq 2$ and $\mathbb{E}(V^{-1}) = 1/(n-2)$ if $n > 2$. The results now follow from the previous result on the [mean](#).

Note that $\text{var}(T) \rightarrow 1$ as $n \rightarrow \infty$.

In the simulation of the special distribution simulator, select the student t distribution. Vary n and note the location and shape of the mean \pm standard deviation bar. For selected values of n , run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Next we give the general moments of the t distribution.

Suppose again that T has the t distribution with $n \in (0, \infty)$ degrees of freedom and $k \in \mathbb{N}$. Then

1. $\mathbb{E}(T^k)$ is undefined if k is odd and $k \geq n$
2. $\mathbb{E}(T^k) = \infty$ if k is even and $k \geq n$
3. $\mathbb{E}(T^k) = 0$ if k is odd and $k < n$
4. If k is even and $k < n$ then

$$\mathbb{E}(T^k) = \frac{n^{k/2} 1 \cdot 3 \cdots (k-1) \Gamma((n-k)/2)}{2^{k/2} \Gamma(n/2)} = \frac{n^{k/2} k! \Gamma((n-k)/2)}{2^k (k/2)! \Gamma(n/2)} \quad (5.10.7)$$

Proof

By independence, $\mathbb{E}(T^k) = n^{k/2} \mathbb{E}(Z^k) \mathbb{E}(V^{-k/2})$. Recall that $\mathbb{E}(Z^k) = 0$ if k is odd, while

$$\mathbb{E}(Z^k) = 1 \cdot 3 \cdots (k-1) = \frac{k!}{(k/2)! 2^{k/2}} \quad (5.10.8)$$

if k is even. Also, $\mathbb{E}(V^{-k/2}) = \infty$ if $k \geq n$, while

$$\mathbb{E}(V^{-k/2}) = \frac{2^{-k/2} \Gamma((n-k)/2)}{\Gamma(n/2)} \quad (5.10.9)$$

if $k < n$. The results now follow by considering the various cases.

From the general moments, we can compute the skewness and kurtosis of T .

Suppose again that T has the t distribution with $n \in (0, \infty)$ degrees of freedom. Then

1. $\text{skew}(T) = 0$ if $n > 3$
2. $\text{kurt}(T) = 3 + \frac{6}{n-4}$ if $n > 4$

Proof

1. This follows from the symmetry of the distribution of T , although $\text{skew}(T)$ only exists if $\mathbb{E}(T^3)$ exists.
2. For $n > 4$,

$$\text{kurt}(T) = \frac{\mathbb{E}(T^4)}{[\mathbb{E}(T^2)]^2} = \frac{3n^2 \Gamma[(n-4)/2] / 4\Gamma(n/2)}{(n/(n-2))^2} = \frac{3(n-2)^2 \Gamma[(n-4)/2]}{4\Gamma(n/2)} \quad (5.10.10)$$

But $\Gamma(n/2) = (n/2 - 1)(n/2 - 2)\Gamma(n/2 - 2)$. Simplifying gives the result.

Note that $\text{kurt}(T) \rightarrow 3$ as $n \rightarrow \infty$ and hence the *excess kurtosis* $\text{kurt}(T) - 3 \rightarrow 0$ as $n \rightarrow \infty$.

In the special distribution simulator, select the student t distribution. Vary n and note the shape of the probability density function in light of the previous results on skewness and kurtosis. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.

Since T does not have moments of all orders, there is no interval about 0 on which the moment generating function of T is finite. The characteristic function exists, of course, but has no simple representation, except in terms of special functions.

Relations

The t distribution with 1 degree of freedom is known as the *Cauchy distribution*. The probability density function is

$$f(t) = \frac{1}{\pi(1+t^2)}, \quad t \in \mathbb{R} \quad (5.10.11)$$

The Cauchy distribution is named after Augustin Cauchy and is studied in more detail in a separate section.

You probably noticed that, qualitatively at least, the t probability density function is very similar to the standard normal probability density function. The similarity is quantitative as well:

Let f_n denote the t probability density function with $n \in (0, \infty)$ degrees of freedom. Then for fixed $t \in \mathbb{R}$,

$$f_n(t) \rightarrow \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2} \text{ as } n \rightarrow \infty \quad (5.10.12)$$

Proof

From a basic limit theorem in calculus,

$$\left(1 + \frac{t^2}{n}\right)^{-(n+1)/2} \rightarrow e^{-t^2/2} \text{ as } n \rightarrow \infty \quad (5.10.13)$$

An application of Stirling's approximation shows that

$$\frac{\Gamma[(n+1)/2]}{\sqrt{n\pi} \Gamma(n/2)} \rightarrow \frac{1}{\sqrt{2\pi}} \text{ as } n \rightarrow \infty \quad (5.10.14)$$

Note that the function on the right is the probability density function of the standard normal distribution. We can also get convergence of the t distribution to the standard normal distribution from the basic random variable representation in the [definition](#).

Suppose that T_n has the t distribution with $n \in \mathbb{N}_+$ degrees of freedom, so that we can represent T_n as

$$T_n = \frac{Z}{\sqrt{V_n/n}} \quad (5.10.15)$$

where Z has the standard normal distribution, V_n has the chi-square distribution with n degrees of freedom, and Z and V_n are independent. Then $T_n \rightarrow Z$ as $n \rightarrow \infty$ with probability 1.

Proof

We can represent V_n as $V_n = Z_1^2 + Z_2^2 + \cdots + Z_n^2$ where (Z_1, Z_2, \dots, Z_n) are independent, standard normal variables, independent of Z . Note that $V_n/n \rightarrow 1$ as $n \rightarrow \infty$ with probability 1 by the strong law of large numbers.

The t distribution has more probability in the tails, and consequently less probability near 0, compared to the standard normal distribution.

The Non-Central t Distribution

One natural way to generalize the student t distribution is to replace the standard normal variable Z in the [definition above](#) with a normal variable having an arbitrary mean (but still unit variance). The reason this particular generalization is important is because it arises in hypothesis tests about the mean based on a random sample from the normal distribution, when the null hypothesis is false. For details see the sections on tests in the normal model and tests in the bivariate normal model in the chapter on Hypothesis Testing.

Suppose that Z has the standard normal distribution, $\mu \in \mathbb{R}$, V has the chi-squared distribution with $n \in (0, \infty)$ degrees of freedom, and that Z and V are independent. Random variable

$$T = \frac{Z + \mu}{\sqrt{V/n}} \quad (5.10.16)$$

has the *non-central student t distribution* with n degrees of freedom and non-centrality parameter μ .

The standard functions that characterize a distribution—the probability density function, distribution function, and quantile function—do not have simple representations for the non-central t distribution, but can only be expressed in terms of other special functions. Similarly, the moments do not have simple, closed form expressions either. For the beginning student of statistics, the most important fact is that the probability density function of the non-central t distribution is similar (but not exactly the same) as that of the standard t distribution (with the same degrees of freedom), but shifted and scaled. The density function is shifted to the right or left, depending on whether $\mu > 0$ or $\mu < 0$.

Computational Exercises

Suppose that T has the t distribution with $n = 10$ degrees of freedom. For each of the following, compute the true value using the special distribution calculator and then compute the normal approximation. Compare the results.

1. $\mathbb{P}(-0.8 < T < 1.2)$
2. The 90th percentile of T .

Answer

1. $\mathbb{P}(-0.8 < T < 1.2) = 0.650$, $\mathbb{P}(-0.8 < T < 1.2) \approx 0.673$
2. $x_{0.90} = 1.372$, $x_{0.90} \approx 1.281$

This page titled [5.10: The Student \$t\$ Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.11: The F Distribution

In this section we will study a distribution that has special importance in statistics. In particular, this distribution arises from ratios of sums of squares when sampling from a normal distribution, and so is important in estimation and in the two-sample normal model and in hypothesis testing in the two-sample normal model.

Basic Theory

Definition

Suppose that U has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom, V has the chi-square distribution with $d \in (0, \infty)$ degrees of freedom, and that U and V are independent. The distribution of

$$X = \frac{U/n}{V/d} \quad (5.11.1)$$

is the F distribution with n degrees of freedom in the numerator and d degrees of freedom in the denominator.

The F distribution was first derived by George Snedecor, and is named in honor of Sir Ronald Fisher. In practice, the parameters n and d are usually positive integers, but this is not a mathematical requirement.

Distribution Functions

Suppose that X has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator. Then X has a continuous distribution on $(0, \infty)$ with probability density function f given by

$$f(x) = \frac{\Gamma(n/2 + d/2)}{\Gamma(n/2)\Gamma(d/2)} \frac{n}{d} \frac{[(n/d)x]^{n/2-1}}{[1 + (n/d)x]^{n/2+d/2}}, \quad x \in (0, \infty) \quad (5.11.2)$$

where Γ is the gamma function.

Proof

The trick, once again, is conditioning. The conditional distribution of X given $V = v \in (0, \infty)$ is gamma with shape parameter $n/2$ and scale parameter $2d/nv$. Hence the conditional PDF is

$$x \mapsto \frac{1}{\Gamma(n/2)(2d/nv)^{n/2}} x^{n/2-1} e^{-x(nv/2d)} \quad (5.11.3)$$

By definition, V has the chi-square distribution with d degrees of freedom, and so has PDF

$$v \mapsto \frac{1}{\Gamma(d/2)2^{d/2}} v^{d/2-1} e^{-v/2} \quad (5.11.4)$$

The joint PDF of (X, V) is the product of these functions:

$$g(x, v) = \frac{1}{\Gamma(n/2)\Gamma(d/2)2^{(n+d)/2}} \left(\frac{n}{d}\right)^{n/2} x^{n/2-1} v^{(n+d)/2-1} e^{-v(n/d+1)/2}; \quad x, v \in (0, \infty) \quad (5.11.5)$$

The PDF of X is therefore

$$f(x) = \int_0^\infty g(x, v) dv = \frac{1}{\Gamma(n/2)\Gamma(d/2)2^{(n+d)/2}} \left(\frac{n}{d}\right)^{n/2} x^{n/2-1} \int_0^\infty v^{(n+d)/2-1} e^{-v(n/d+1)/2} dv \quad (5.11.6)$$

Except for the normalizing constant, the integrand in the last integral is the gamma PDF with shape parameter $(n+d)/2$ and scale parameter $2d/(nx+d)$. Hence the integral evaluates to

$$\Gamma\left(\frac{n+d}{2}\right) \left(\frac{2d}{nx+d}\right)^{(n+d)/2} \quad (5.11.7)$$

Simplifying gives the result.

Recall that the beta function B can be written in terms of the gamma function by

$$B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}, \quad a, b \in (0, \infty) \quad (5.11.8)$$

Hence the probability density function of the F distribution above can also be written as

$$f(x) = \frac{1}{B(n/2, d/2)} \frac{n}{d} \frac{[(n/d)x]^{n/2-1}}{[1 + (n/d)x]^{n/2+d/2}}, \quad x \in (0, \infty) \quad (5.11.9)$$

When $n \geq 2$, the probability density function is defined at $x = 0$, so the support interval is $[0, \infty)$ in this case.

In the special distribution simulator, select the F distribution. Vary the parameters with the scroll bars and note the shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Both parameters influence the shape of the F probability density function, but some of the basic *qualitative* features depend only on the numerator degrees of freedom. For the remainder of this discussion, let f denote the F probability density function with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom

in the denominator.

Probability density function f satisfies the following properties:

1. If $0 < n < 2$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
2. If $n = 2$, f is decreasing with mode at $x = 0$.
3. If $n > 2$, f increases and then decreases, with mode at $x = \frac{(n-2)d}{n(d+2)}$.

Proof

These properties follow from standard calculus. The first derivative of f is

$$f'(x) = \frac{1}{B(n/2, d/2)} \left(\frac{n}{d}\right)^2 \frac{[(n/d)x]^{n/2-2}}{[1 + (n/2)x]^{n/2+d/2+1}} [(n/2-1) - (n/d)(d/2+1)x], \quad x \in (0, \infty) \quad (5.11.10)$$

Qualitatively, the second order properties of f also depend only on n , with transitions at $n = 2$ and $n = 4$.

For $n > 2$, define

$$x_1 = \frac{d}{n} \frac{(n-2)(d+4) - \sqrt{2(n-2)(d+4)(n+d)}}{(d+2)(d+4)} \quad (5.11.11)$$

$$x_2 = \frac{d}{n} \frac{(n-2)(d+4) + \sqrt{2(n-2)(d+4)(n+d)}}{(d+1)(d+4)} \quad (5.11.12)$$

The probability density function f satisfies the following properties:

1. If $0 < n \leq 2$, f is concave upward.
2. If $2 < n \leq 4$, f is concave downward and then upward, with inflection point at x_2 .
3. If $n > 4$, f is concave upward, then downward, then upward again, with inflection points at x_1 and x_2 .

Proof

These results follow from standard calculus. The second derivative of f is

$$f''(x) = \frac{1}{B(n/2, d/2)} \left(\frac{n}{d}\right)^3 \frac{[(n/d)x]^{n/2-3}}{[1 + (n/d)x]^{n/2+d/2+2}} [(n/2-1)(n/2-2) - 2(n/2-1)(d/2+2)(n/d)x + (d/2+1)(d/2+2)(n/d)^2 x^2], \quad (5.11.13)$$

$$x \in (0, \infty)$$

The distribution function and the quantile function do not have simple, closed-form representations. Approximate values of these functions can be obtained from the special distribution calculator and from most mathematical and statistical software packages.

In the special distribution calculator, select the F distribution. Vary the parameters and note the shape of the probability density function and the distribution function. In each of the following cases, find the median, the first and third quartiles, and the interquartile range.

1. $n = 5$, $d = 5$
2. $n = 5$, $d = 10$
3. $n = 10$, $d = 5$
4. $n = 10$, $d = 10$

The general probability density function of the F distribution is a bit complicated, but it simplifies in a couple of special cases.

Special cases.

1. If $n = 2$,

$$f(x) = \frac{1}{(1+2x/d)^{1+d/2}}, \quad x \in (0, \infty) \quad (5.11.14)$$

2. If $n = d \in (0, \infty)$,

$$f(x) = \frac{\Gamma(n)}{\Gamma^2(n/2)} \frac{x^{n/2-1}}{(1+x)^n}, \quad x \in (0, \infty) \quad (5.11.15)$$

3. If $n = d = 2$,

$$f(x) = \frac{1}{(1+x)^2}, \quad x \in (0, \infty) \quad (5.11.16)$$

4. If $n = d = 1$,

$$f(x) = \frac{1}{\pi\sqrt{x(1+x)}}, \quad x \in (0, \infty) \quad (5.11.17)$$

Moments

The random variable representation in the [definition](#), along with the moments of the chi-square distribution can be used to find the mean, variance, and other moments of the F distribution. For the remainder of this discussion, suppose that X has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator.

Mean

1. $\mathbb{E}(X) = \infty$ if $0 < d \leq 2$
2. $\mathbb{E}(X) = \frac{d}{d-2}$ if $d > 2$

Proof

By independence, $\mathbb{E}(X) = \frac{d}{n} \mathbb{E}(U) \mathbb{E}(V^{-1})$. Recall that $\mathbb{E}(U) = n$. Similarly if $d \leq 2$, $\mathbb{E}(V^{-1}) = \infty$ while if $d > 2$,

$$\mathbb{E}(V^{-1}) = \frac{\Gamma(d/2 - 1)}{2\Gamma(d/2)} = \frac{1}{d-2} \quad (5.11.18)$$

Thus, the mean depends only on the degrees of freedom in the denominator.

Variance

1. $\text{var}(X)$ is undefined if $0 < d \leq 2$
2. $\text{var}(X) = \infty$ if $2 < d \leq 4$
3. If $d > 4$ then

$$\text{var}(X) = 2 \left(\frac{d}{d-2} \right)^2 \frac{n+d-2}{n(d-4)} \quad (5.11.19)$$

Proof

By independence, $\mathbb{E}(X^2) = \frac{d^2}{n^2} \mathbb{E}(U^2) \mathbb{E}(V^{-2})$. Recall that

$$\mathbb{E}(U^2) = 4 \frac{\Gamma(n/2 + 2)}{\Gamma(n/2)} = (n+2)n \quad (5.11.20)$$

Similarly if $d \leq 4$, $\mathbb{E}(V^{-2}) = \infty$ while if $d > 4$,

$$\mathbb{E}(V^{-2}) = \frac{\Gamma(d/2 - 2)}{4\Gamma(d/2)} = \frac{1}{(d-2)(d-4)} \quad (5.11.21)$$

Hence $\mathbb{E}(X^2) = \infty$ if $d \leq 4$ while if $d > 4$,

$$\mathbb{E}(X^2) = \frac{(n+2)d^2}{n(d-2)(d-4)} \quad (5.11.22)$$

The results now follow from the previous result on the [mean](#) and the computational formula $\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2$.

In the simulation of the special distribution simulator, select the F distribution. Vary the parameters with the scroll bar and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

General moments. For $k > 0$,

1. $\mathbb{E}(X^k) = \infty$ if $0 < d \leq 2k$
2. If $d > 2k$ then

$$\mathbb{E}(X^k) = \left(\frac{d}{n} \right)^k \frac{\Gamma(n/2 + k) \Gamma(d/2 - k)}{\Gamma(n/2) \Gamma(d/2)} \quad (5.11.23)$$

Proof

By independence, $\mathbb{E}(X^k) = \left(\frac{d}{n} \right)^k \mathbb{E}(U^k) \mathbb{E}(V^{-k})$. Recall that

$$\mathbb{E}(U^k) = \frac{2^k \Gamma(n/2 + k)}{\Gamma(n/2)} \quad (5.11.24)$$

On the other hand, $\mathbb{E}(V^{-k}) = \infty$ if $d/2 \leq k$ while if $d/2 > k$,

$$\mathbb{E}(V^{-k}) = \frac{2^{-k} \Gamma(d/2 - k)}{\Gamma(d/2)} \quad (5.11.25)$$

If $k \in \mathbb{N}$, then using the fundamental identity of the gamma distribution and some algebra,

$$\mathbb{E}(X^k) = \left(\frac{d}{n} \right)^k \frac{n(n+2) \cdots [n+2(k-1)]}{(d-2)(d-4) \cdots (d-2k)} \quad (5.11.26)$$

From the general moment formula, we can compute the skewness and kurtosis of the F distribution.

Skewness and kurtosis

1. If $d > 6$,

$$\text{skew}(X) = \frac{(2n+d-2)\sqrt{8(d-4)}}{(d-6)\sqrt{n(n+d-2)}} \quad (5.11.27)$$

2. If $d > 8$,

$$\text{kurt}(X) = 3 + 12 \frac{n(5d-22)(n+d-2) + (d-4)(d-2)^2}{n(d-6)(d-8)(n+d-2)} \quad (5.11.28)$$

Proof

These results follow from the formulas for $\mathbb{E}(X^k)$ for $k \in \{1, 2, 3, 4\}$ and the standard computational formulas for skewness and kurtosis.

Not surprisingly, the F distribution is positively skewed. Recall that the *excess kurtosis* is

$$\text{kurt}(X) - 3 = 12 \frac{n(5d-22)(n+d-2) + (d-4)(d-2)^2}{n(d-6)(d-8)(n+d-2)} \quad (5.11.29)$$

In the simulation of the special distribution simulator, select the F distribution. Vary the parameters with the scroll bar and note the shape of the probability density function in light of the previous results on skewness and kurtosis. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Relations

The most important relationship is the one in the [definition](#), between the F distribution and the chi-square distribution. In addition, the F distribution is related to several other special distributions.

Suppose that X has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator. Then $1/X$ has the F distribution with d degrees of freedom in the numerator and n degrees of freedom in the denominator.

Proof

This follows easily from the random variable interpretation in the [definition](#). We can write

$$X = \frac{U/n}{V/d} \quad (5.11.30)$$

where U and V are independent and have chi-square distributions with n and d degrees of freedom, respectively. Hence

$$\frac{1}{X} = \frac{V/d}{U/n} \quad (5.11.31)$$

Suppose that T has the t distribution with $n \in (0, \infty)$ degrees of freedom. Then $X = T^2$ has the F distribution with 1 degree of freedom in the numerator and n degrees of freedom in the denominator.

Proof

This follows easily from the random variable representations of the t and F distributions. We can write

$$T = \frac{Z}{\sqrt{V/n}} \quad (5.11.32)$$

where Z has the standard normal distribution, V has the chi-square distribution with n degrees of freedom, and Z and V are independent. Hence

$$T^2 = \frac{Z^2}{V/n} \quad (5.11.33)$$

Recall that Z^2 has the chi-square distribution with 1 degree of freedom.

Our next relationship is between the F distribution and the exponential distribution.

Suppose that X and Y are independent random variables, each with the exponential distribution with rate parameter $r \in (0, \infty)$. Then $Z = X/Y$ has the F distribution with 2 degrees of freedom in both the numerator and denominator.

Proof

We first find the distribution function F of Z by conditioning on X :

$$F(z) = \mathbb{P}(Z \leq z) = \mathbb{P}(Y \geq X/z) = \mathbb{E}[\mathbb{P}(Y \geq X/z | X)] \quad (5.11.34)$$

But $\mathbb{P}(Y \geq y) = e^{-ry}$ for $y \geq 0$ so $F(z) = \mathbb{E}(e^{-rX/z})$. Also, X has PDF $g(x) = re^{-rx}$ for $x \geq 0$ so

$$F(z) = \int_0^\infty e^{-rx/z} re^{-rx} dx = \int_0^\infty re^{-rx(1+1/z)} dx = \frac{1}{1+1/z} = \frac{z}{1+z}, \quad z \in (0, \infty) \quad (5.11.35)$$

Differentiating gives the PDF of Z

$$f(z) = \frac{1}{(1+z)^2}, \quad z \in (0, \infty) \quad (5.11.36)$$

which we recognize as the PDF of the F distribution with 2 degrees of freedom in the numerator and the denominator.

A simple transformation can change a variable with the F distribution into a variable with the beta distribution, and conversely.

Connections between the F distribution and the beta distribution.

1. If X has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator, then

$$Y = \frac{(n/d)X}{1 + (n/d)X} \quad (5.11.37)$$

has the beta distribution with left parameter $n/2$ and right parameter $d/2$.

2. If Y has the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$ then

$$X = \frac{bY}{a(1-Y)} \quad (5.11.38)$$

has the F distribution with $2a$ degrees of freedom in the numerator and $2b$ degrees of freedom in the denominator.

Proof

The two statements are equivalent and follow from the standard change of variables formula. The function

$$y = \frac{(n/d)x}{1 + (n/d)x} \quad (5.11.39)$$

maps $(0, \infty)$ one-to-one onto $(0, 1)$, with inverse

$$x = \frac{d}{n} \frac{y}{1-y} \quad (5.11.40)$$

Let f denote the PDF of the F distribution with n degrees of freedom in the numerator and d degrees of freedom in the denominator, and let g denote the PDF of the beta distribution with left parameter $n/2$ and right parameter $d/2$. Then f and g are related by

1. $g(y) = f(x) \frac{dx}{dy}$
2. $f(x) = g(y) \frac{dy}{dx}$

The F distribution is closely related to the beta prime distribution by a simple scale transformation.

Connections with the beta prime distributions.

1. If X has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator, then $Y = \frac{a}{d}X$ has the beta prime distribution with parameters $n/2$ and $d/2$.
2. If Y has the beta prime distribution with parameters $a \in (0, \infty)$ and $b \in (0, \infty)$ then $X = \frac{b}{a}Y$ has the F distribution with $2a$ degrees of the freedom in the numerator and $2b$ degrees of freedom in the denominator.

Proof

Let f denote the PDF of X and g the PDF of Y .

1. By the change of variables formula,

$$g(y) = \frac{d}{n} f\left(\frac{d}{n}y\right), \quad y \in (0, \infty) \quad (5.11.41)$$

Substituting into the beta F PDF shows that Y has the appropriate beta prime distribution.

2. Again using the change of variables formula,

$$f(x) = \frac{a}{b} g\left(\frac{a}{b}x\right), \quad x \in (0, \infty) \quad (5.11.42)$$

Substituting into the beta prime PDF shows that X has the appropriate F PDF.

The Non-Central F Distribution

The F distribution can be generalized in a natural way by replacing the ordinary chi-square variable in the numerator in the [definition above](#) with a variable having a non-central chi-square distribution. This generalization is important in analysis of variance.

Suppose that U has the non-central chi-square distribution with $n \in (0, \infty)$ degrees of freedom and non-centrality parameter $\lambda \in [0, \infty)$, V has the chi-square distribution with $d \in (0, \infty)$ degrees of freedom, and that U and V are independent. The distribution of

$$X = \frac{U/n}{V/d} \quad (5.11.43)$$

is the *non-central F distribution* with n degrees of freedom in the numerator, d degrees of freedom in the denominator, and *non-centrality parameter* λ .

One of the most interesting and important results for the non-central chi-square distribution is that it is a Poisson mixture of ordinary chi-square distributions. This leads to a similar result for the non-central F distribution.

Suppose that N has the Poisson distribution with parameter $\lambda/2$, and that the conditional distribution of X given N is the F distribution with $N + 2n$ degrees of freedom in the numerator and d degrees of freedom in the denominator, where $\lambda \in [0, \infty)$ and $n, d \in (0, \infty)$. Then X has the non-central F distribution with n degrees of freedom in the numerator, d degrees of freedom in the denominator, and non-centrality parameter λ .

Proof

As in the theorem, let N have the Poisson distribution with parameter $\lambda/2$, and suppose also that the conditional distribution of U given N is chi-square with $n+2N$ degrees of freedom, and that V has the chi-square distribution with d degrees of freedom and is independent of (N, U) . Let $X = (U/n)/(V/d)$. Since V is independent of (N, U) , the variable X satisfies the condition in the theorem; that is, the conditional distribution of X given N is the F distribution with $n+2N$ degrees of freedom in the numerator and d degrees of freedom in the denominator. But then also, (unconditionally) U has the non-central chi-square distribution with n degrees of freedom in the numerator and non-centrality parameter λ , V has the chi-square distribution with d degrees of freedom, and U and V are independent. So by definition X has the F distribution with n degrees of freedom in the numerator, d degrees of freedom in the denominator, and non-centrality parameter λ .

From the last result, we can express the probability density function and distribution function of the non-central F distribution as a series in terms of ordinary F density and distribution functions. To set up the notation, for $j, k \in (0, \infty)$ let f_{jk} be the probability density function and F_{jk} the distribution function of the F distribution with j degrees of freedom in the numerator and k degrees of freedom in the denominator. For the rest of this discussion, $\lambda \in [0, \infty)$ and $n, d \in (0, \infty)$ as usual.

The probability density function g of the non-central F distribution with n degrees of freedom in the numerator, d degrees of freedom in the denominator, and non-centrality parameter λ is given by

$$g(x) = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} f_{n+2k,d}(x), \quad x \in (0, \infty) \quad (5.11.44)$$

The distribution function G of the non-central F distribution with n degrees of freedom in the numerator, d degrees of freedom in the denominator, and non-centrality parameter λ is given by

$$G(x) = \sum_{k=0}^{\infty} e^{-\lambda/2} \frac{(\lambda/2)^k}{k!} F_{n+2k,d}(x), \quad x \in (0, \infty) \quad (5.11.45)$$

This page titled 5.11: The F Distribution is shared under a CC BY 2.0 license and was authored, remixed, and/or curated by Kyle Siegrist (Random Services) via source content that was edited to the style and standards of the LibreTexts platform.

5.12: The Lognormal Distribution

Basic Theory

Definition

Suppose that Y has the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$. Then $X = e^Y$ has the *lognormal distribution* with parameters μ and σ .

1. The parameter σ is the *shape parameter* of the distribution.
2. The parameter e^μ is the *scale parameter* of the distribution.

If Z has the standard normal distribution then $W = e^Z$ has the *standard lognormal distribution*.

So equivalently, if X has a lognormal distribution then $\ln X$ has a normal distribution, hence the name. The lognormal distribution is a continuous distribution on $(0, \infty)$ and is used to model random quantities when the distribution is believed to be skewed, such as certain income and lifetime variables. It's easy to write a general lognormal variable in terms of a standard lognormal variable. Suppose that Z has the standard normal distribution and let $W = e^Z$ so that W has the standard lognormal distribution. If $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ then $Y = \mu + \sigma Z$ has the normal distribution with mean μ and standard deviation σ and hence $X = e^Y$ has the lognormal distribution with parameters μ and σ . But

$$X = e^Y = e^{\mu + \sigma Z} = e^\mu (e^Z)^\sigma = e^\mu W^\sigma \quad (5.12.1)$$

Distribution Functions

Suppose that X has the lognormal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$.

The probability density function f of X is given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma x} \exp\left[-\frac{(\ln x - \mu)^2}{2\sigma^2}\right], \quad x \in (0, \infty) \quad (5.12.2)$$

1. f increases and then decreases with mode at $x = \exp(\mu - \sigma^2)$.
2. f is concave upward then downward then upward again, with inflection points at $x = \exp\left(\mu - \frac{3}{2}\sigma^2 \pm \frac{1}{2}\sigma\sqrt{\sigma^2 + 4}\right)$
3. $f(x) \rightarrow 0$ as $x \downarrow 0$ and as $x \rightarrow \infty$.

Proof

The form of the PDF follows from the change of variables theorem. Let g denote the PDF of the normal distribution with mean μ and standard deviation σ , so that

$$g(y) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{y - \mu}{\sigma}\right)^2\right], \quad y \in \mathbb{R} \quad (5.12.3)$$

The mapping $x = e^y$ maps \mathbb{R} one-to-one onto $(0, \infty)$ with inverse $y = \ln x$. Hence the PDF f of $X = e^Y$ is

$$f(x) = g(y) \frac{dy}{dx} = g(\ln x) \frac{1}{x} \quad (5.12.4)$$

Substituting gives the result. Parts (a)–(d) follow from standard calculus.

In the special distribution simulator, select the lognormal distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the true probability density function.

Let Φ denote the standard normal distribution function, so that Φ^{-1} is the standard normal quantile function. Recall that values of Φ and Φ^{-1} can be obtained from the special distribution calculator, as well as standard mathematical and statistical software packages, and in fact these functions are considered to be special functions in mathematics. The following two results show how to compute the lognormal distribution function and quantiles in terms of the standard normal distribution function and quantiles.

The distribution function F of X is given by

$$F(x) = \Phi\left(\frac{\ln x - \mu}{\sigma}\right), \quad x \in (0, \infty) \quad (5.12.5)$$

Proof

Once again, write $X = e^{\mu + \sigma Z}$ where Z has the standard normal distribution. For $x > 0$,

$$F(x) = \mathbb{P}(X \leq x) = \mathbb{P}\left(Z \leq \frac{\ln x - \mu}{\sigma}\right) = \Phi\left(\frac{\ln x - \mu}{\sigma}\right) \quad (5.12.6)$$

The quantile function of X is given by

$$F^{-1}(p) = \exp[\mu + \sigma \Phi^{-1}(p)], \quad p \in (0, 1) \quad (5.12.7)$$

Proof

This follows by solving $p = F(x)$ for x in terms of p .

In the special distribution calculator, select the lognormal distribution. Vary the parameters and note the shape and location of the probability density function and the distribution function. With $\mu = 0$ and $\sigma = 1$, find the median and the first and third quartiles.

Moments

The moments of the lognormal distribution can be computed from the moment generating function of the normal distribution. Once again, we assume that X has the lognormal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$.

For $t \in \mathbb{R}$,

$$\mathbb{E}(X^t) = \exp\left(\mu t + \frac{1}{2}\sigma^2 t^2\right) \quad (5.12.8)$$

Proof

Recall that if Y has the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$, then Y has moment generating function given by

$$\mathbb{E}(e^{tY}) = \exp\left(\mu t + \frac{1}{2}\sigma^2 t^2\right), \quad t \in \mathbb{R} \quad (5.12.9)$$

Hence the result follows immediately since $\mathbb{E}(X^t) = \mathbb{E}(e^{tY})$.

In particular, the mean and variance of X are

1. $\mathbb{E}(X) = \exp\left(\mu + \frac{1}{2}\sigma^2\right)$
2. $\text{var}(X) = \exp[2(\mu + \sigma^2)] - \exp(2\mu + \sigma^2)$

In the simulation of the special distribution simulator, select the lognormal distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical moments to the true moments.

From the general formula for the moments, we can also compute the skewness and kurtosis of the lognormal distribution.

The skewness and kurtosis of X are

1. $\text{skew}(X) = (e^{\sigma^2} + 2) \sqrt{e^{\sigma^2} - 1}$
2. $\text{kurt}(X) = e^{4\sigma^2} + 2e^{3\sigma^2} + 3e^{2\sigma^2} - 3$

Proof

These result follow from the first 4 [moments of the lognormal distribution](#) and the standard computational formulas for skewness and kurtosis.

The fact that the skewness and kurtosis do not depend on μ is due to the fact that μ is a scale parameter. Recall that skewness and kurtosis are defined in terms of the standard score and so are independent of location and scale parameters. Naturally, the lognormal distribution is positively skewed. Finally, note that the *excess kurtosis* is

$$\text{kurt}(X) - 3 = e^{4\sigma^2} + 2e^{3\sigma^2} + 3e^{2\sigma^2} - 6 \quad (5.12.10)$$

Even though the lognormal distribution has finite moments of all orders, the moment generating function is infinite at any positive number. This property is one of the reasons for the fame of the lognormal distribution.

$\mathbb{E}(e^{tX}) = \infty$ for every $t > 0$.

Proof

By definition, $X = e^Y$ where Y has the normal distribution with mean μ and standard deviation σ . Using the change of variables formula for expected value we have

$$\mathbb{E}(e^{tX}) = \mathbb{E}(e^{te^Y}) = \int_{-\infty}^{\infty} \exp(te^y) \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2\right] dy = \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \exp\left[te^y - \frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2\right] dy \quad (5.12.11)$$

If $t > 0$ the integrand in the last integral diverges to ∞ as $y \rightarrow \infty$, so there is no hope that the integral converges.

Related Distributions

The most important relations are the ones between the lognormal and normal distributions in the definition: if X has a lognormal distribution then $\ln X$ has a normal distribution; conversely if Y has a normal distribution then e^Y has a lognormal distribution. The lognormal distribution is also a scale family.

Suppose that X has the lognormal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ and that $c \in (0, \infty)$. Then cX has the lognormal distribution with parameters $\mu + \ln c$ and σ .

Proof

From the [definition](#), we can write $X = e^Y$ where Y has the normal distribution with mean μ and standard deviation σ . Hence

$$cX = ce^Y = e^{\ln c} e^Y = e^{\ln c + Y} \quad (5.12.12)$$

But $\ln c + Y$ has the normal distribution with mean $\ln c + \mu$ and standard deviation σ .

The reciprocal of a lognormal variable is also lognormal.

If X has the lognormal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ then $1/X$ has the lognormal distribution with parameters $-\mu$ and σ .

Proof

Again from the [definition](#), we can write $X = e^Y$ where Y has the normal distribution with mean μ and standard deviation σ . Hence $1/X = e^{-Y}$. But $-Y$ has the normal distribution with mean $-\mu$ and standard deviation σ .

The lognormal distribution is closed under non-zero powers of the underlying variable. In particular, this generalizes the previous result.

Suppose that X has the lognormal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ and that $a \in \mathbb{R} \setminus \{0\}$. Then X^a has the lognormal distribution with parameters $a\mu$ and $|a|\sigma$.

Proof

Again from the [definition](#), we can write $X = e^Y$ where Y has the normal distribution with mean μ and standard deviation σ . Hence $X^a = e^{aY}$. But aY has the normal distribution with mean $a\mu$ and standard deviation $|a|\sigma$.

Since the normal distribution is closed under sums of independent variables, it's not surprising that the lognormal distribution is closed under products of independent variables.

Suppose that $n \in \mathbb{N}_+$ and that (X_1, X_2, \dots, X_n) is a sequence of independent variables, where X_i has the lognormal distribution with parameters $\mu_i \in \mathbb{R}$ and $\sigma_i \in (0, \infty)$ for $i \in \{1, 2, \dots, n\}$. Then $\prod_{i=1}^n X_i$ has the lognormal distribution with parameters μ and σ where $\mu = \sum_{i=1}^n \mu_i$ and $\sigma^2 = \sum_{i=1}^n \sigma_i^2$.

Proof

Again from the [definition](#), we can write $X_i = e^{Y_i}$ where Y_i has the normal distribution with mean μ_i and standard deviation σ_i for $i \in \{1, 2, \dots, n\}$ and where (Y_1, Y_2, \dots, Y_n) is an independent sequence. Hence $\prod_{i=1}^n X_i = \exp(\sum_{i=1}^n Y_i)$. But $\sum_{i=1}^n Y_i$ has the normal distribution with mean $\sum_{i=1}^n \mu_i$ and variance $\sum_{i=1}^n \sigma_i^2$.

Finally, the lognormal distribution belongs to the family of general exponential distributions.

Suppose that X has the lognormal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$. The distribution of X is a 2-parameter exponential family with natural parameters and natural statistics, respectively, given by

1. $(-1/2\sigma^2, \mu/\sigma^2)$
2. $(\ln^2(X), \ln X)$

Proof

This follows from the definition of the general exponential family, since we can write the [lognormal PDF](#) in the form

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \frac{1}{x} \exp\left[-\frac{1}{2\sigma^2} \ln^2(x) + \frac{\mu}{\sigma^2} \ln x\right], \quad x \in (0, \infty) \quad (5.12.13)$$

Computational Exercises

Suppose that the income X of a randomly chosen person in a certain population (in \$1000 units) has the lognormal distribution with parameters $\mu = 2$ and $\sigma = 1$. Find $\mathbb{P}(X > 20)$.

Answer

$$\mathbb{P}(X > 20) = 0.1497$$

Suppose that the income X of a randomly chosen person in a certain population (in \$1000 units) has the lognormal distribution with parameters $\mu = 2$ and $\sigma = 1$. Find each of the following:

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

Answer

1. $\mathbb{E}(X) = e^{5/2} \approx 12.1825$
2. $\text{sd}(X) = \sqrt{e^6 - e^5} \approx 15.9629$

This page titled [5.12: The Lognormal Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.13: The Folded Normal Distribution

The General Folded Normal Distribution

Introduction

The *folded normal distribution* is the distribution of the absolute value of a random variable with a normal distribution. As has been emphasized before, the normal distribution is perhaps the most important in probability and is used to model an incredible variety of random phenomena. Since one may only be interested in the magnitude of a normally distributed variable, the folded normal arises in a very natural way. The name stems from the fact that the probability measure of the normal distribution on $(-\infty, 0]$ is “folded over” to $[0, \infty)$. Here is the formal definition:

Suppose that Y has a normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$. Then $X = |Y|$ has the *folded normal distribution* with parameters μ and σ .

So in particular, the folded normal distribution is a continuous distribution on $[0, \infty)$.

Distribution Functions

Suppose that Z has the standard normal distribution. Recall that Z has probability density function ϕ and distribution function Φ given by

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad z \in \mathbb{R} \quad (5.13.1)$$

$$\Phi(z) = \int_{-\infty}^z \phi(x) dx = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx, \quad z \in \mathbb{R} \quad (5.13.2)$$

The standard normal distribution is so important that Φ is considered a special function and can be computed using most mathematical and statistical software. If $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, then $Y = \mu + \sigma Z$ has the normal distribution with mean μ and standard deviation σ , and therefore $X = |Y| = |\mu + \sigma Z|$ has the folded normal distribution with parameters μ and σ . For the remainder of this discussion we assume that X has this folded normal distribution.

X has distribution function F given by

$$F(x) = \Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(\frac{-x-\mu}{\sigma}\right) = \Phi\left(\frac{x-\mu}{\sigma}\right) + \Phi\left(\frac{x+\mu}{\sigma}\right) - 1 \quad (5.13.3)$$

$$= \int_0^x \frac{1}{\sigma\sqrt{2\pi}} \left\{ \exp\left[-\frac{1}{2}\left(\frac{y+\mu}{\sigma}\right)^2\right] + \exp\left[-\frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2\right] \right\} dy, \quad x \in [0, \infty) \quad (5.13.4)$$

Proof

For $x \in [0, \infty)$,

$$F(x) = \mathbb{P}(X \leq x) = \mathbb{P}(|Y| \leq x) = \mathbb{P}(|\mu + \sigma Z| \leq x) = \mathbb{P}(-x \leq \mu + \sigma Z \leq x) \quad (5.13.5)$$

$$= \mathbb{P}\left(\frac{-x-\mu}{\sigma} \leq Z \leq \frac{x-\mu}{\sigma}\right) = \Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(\frac{-x-\mu}{\sigma}\right) \quad (5.13.6)$$

which gives the first expression. The second expression follows since $\Phi(-z) = 1 - \Phi(z)$ for $z \in \mathbb{R}$. Finally, the integral formula follows from the form of Φ given above and simple substitution.

We cannot compute the quantile function F^{-1} in closed form, but values of this function can be approximated.

Open the special distribution calculator and select the folded normal distribution, and set the view to CDF. Vary the parameters and note the shape of the distribution function. For selected values of the parameters, compute the median and the first and third quartiles.

X has probability density function f given by

$$f(x) = \frac{1}{\sigma} \left[\phi\left(\frac{x-\mu}{\sigma}\right) + \phi\left(\frac{x+\mu}{\sigma}\right) \right] \quad (5.13.7)$$

$$= \frac{1}{\sigma\sqrt{2\pi}} \left\{ \exp\left[-\frac{1}{2}\left(\frac{x+\mu}{\sigma}\right)^2\right] + \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right] \right\}, \quad x \in [0, \infty) \quad (5.13.8)$$

Proof

This follows from differentiating the CDF with respect to x , since $F'(x) = f(x)$ and $\Phi'(z) = \phi(z)$.

Open the special distribution simulator and select the folded normal distribution. Vary the parameters μ and σ and note the shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the true probability density function.

Note that the folded normal distribution is unimodal for some values of the parameters and decreasing for other values. Note also that μ is not a location parameter nor is σ a scale parameter; both influence the shape of the probability density function.

Moments

We cannot compute the mean of the folded normal distribution in closed form, but the mean can at least be given in terms of Φ . Once again, we assume that X has the folded normal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$.

The first two moments of x are

1. $\mathbb{E}(X) = \mu[1 - 2\Phi(-\mu/\sigma)] + \sigma\sqrt{2/\pi}\exp(-\mu^2/2\sigma^2)$
2. $\mathbb{E}(X^2) = \mu^2 + \sigma^2$

Proof

From the definition, we can assume $X = |\mu + \sigma Z|$ where Z has the standard normal distribution. Then

$$\mathbb{E}(X) = \mathbb{E}(|\mu + \sigma Z|) = \mathbb{E}(\mu + \sigma Z; Z \geq -\mu/\sigma) - \mathbb{E}(\mu + \sigma Z; Z \leq -\mu/\sigma) \quad (5.13.9)$$

$$= \mathbb{E}(\mu + \sigma Z) - 2\mathbb{E}(\mu + \sigma Z; Z \leq -\mu/\sigma) = \mu - 2\mu\Phi(-\mu/\sigma) - 2\sigma\mathbb{E}(Z; Z \leq -\mu/\sigma) \quad (5.13.10)$$

So we just need to compute the last expected value. Using the change of variables $u = z^2/2$ we get

$$\mathbb{E}(Z; Z \leq -\mu/\sigma) = \int_{-\infty}^{-\mu/\sigma} z \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = - \int_{\mu^2/2\sigma^2}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-u} du = - \frac{1}{\sqrt{2\pi}} e^{-\mu^2/2\sigma^2} \quad (5.13.11)$$

Substituting gives the result in (a). For (b), let Y have the normal distribution with mean μ and standard deviation σ so that we can take $X = |Y|$. Then $\mathbb{E}(X^2) = \mathbb{E}(Y^2) = \text{var}(Y) + [\mathbb{E}(Y)]^2 = \sigma^2 + \mu^2$.

In particular, the variance of X is

$$\text{var}(X) = \mu^2 + \sigma^2 - \left\{ \mu \left[1 - 2\Phi\left(-\frac{\mu}{\sigma}\right) \right] + \sigma \sqrt{\frac{2}{\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \right\}^2 \quad (5.13.12)$$

Open the special distribution simulator and select the folded normal distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

Related Distributions

The most important relation is the one between the folded normal distribution and the normal distribution in the definition: If Y has a normal distribution then $X = |Y|$ has a folded normal distribution. The folded normal distribution is also related to itself through a symmetry property that is perhaps not completely obvious from the initial definition:

For $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$, the folded normal distribution with parameters $-\mu$ and σ is the same as the folded normal distribution with parameters μ and σ .

Proof 1

The PDF is unchanged if μ is replaced with $-\mu$.

Proof 2

Suppose that Y has the normal distribution with mean μ and standard deviation σ so that $|Y|$ has the folded normal distribution with parameters μ and σ . Then $-Y$ has the normal distribution with mean $-\mu$ and standard deviation σ so that $|-Y|$ has the folded normal distribution with parameters $-\mu$ and σ . But $|Y| = |-Y|$.

The folded normal distribution is also closed under scale transformations.

Suppose that X has the folded normal distribution with parameters $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ and that $b \in (0, \infty)$. Then bX has the folded normal distribution with parameters $b\mu$ and $b\sigma$.

Proof

Once again from the definition, we can assume $X = |Y|$ where Y has the normal distribution with mean μ and standard deviation σ . But then $bX = b|Y| = |bY|$, and bY has the normal distribution with mean $b\mu$ and standard deviation $b\sigma$.

The Half-Normal Distribution

When $\mu = 0$, results for the folded normal distribution are much simpler, and fortunately this special case is the most important one. We are more likely to be interested in the magnitude of a normally distributed variable when the mean is 0, and moreover, this distribution arises in the study of Brownian motion.

Suppose that Z has the standard normal distribution and that $\sigma \in (0, \infty)$. Then $X = \sigma|Z|$ has the *half-normal distribution* with scale parameter σ . If $\sigma = 1$ so that $X = |Z|$, then X has the *standard half-normal distribution*.

Distribution Functions

For our next discussion, suppose that X has the half-normal distribution with parameter $\sigma \in (0, \infty)$. Once again, Φ and Φ^{-1} denote the distribution function and quantile function, respectively, of the standard normal distribution.

The distribution function F and quantile function F^{-1} of X are

$$F(x) = 2\Phi\left(\frac{x}{\sigma}\right) - 1 = \int_0^x \frac{1}{\sigma} \sqrt{\frac{2}{\pi}} \exp\left(-\frac{y^2}{2\sigma^2}\right) dy, \quad x \in [0, \infty) \quad (5.13.13)$$

$$F^{-1}(p) = \sigma \Phi^{-1}\left(\frac{1+p}{2}\right), \quad p \in [0, 1) \quad (5.13.14)$$

Proof

The result for the CDF follows from the CDF of the folded normal distribution with $\mu = 0$. Recall that $\Phi(-z) = 1 - \Phi(z)$ for $z \in \mathbb{R}$. The result for the quantile function follows from the result for the CDF and simple algebra.

Open the special distribution calculator and select the folded normal distribution. Select CDF view and keep $\mu = 0$. Vary σ and note the shape of the CDF. For various values of σ , compute the median and the first and third quartiles.

The probability density function f of X is given by

$$f(x) = \frac{2}{\sigma} \phi\left(\frac{x}{\sigma}\right) = \frac{1}{\sigma} \sqrt{\frac{2}{\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad x \in [0, \infty) \quad (5.13.15)$$

1. f is decreasing with mode at $x = 0$.
2. f is concave downward and then upward, with inflection point at $x = \sigma$.

Proof

The formula for f follows from differentiating the CDF above. Properties (a) and (b) follow from standard calculus.

Open the special distribution simulator and select the folded normal distribution. Keep $\mu = 0$ and vary σ , and note the shape of the probability density function. For selected values of σ , run the simulation 1000 times and compare the empirical density

function to the true probability density function.

Moments

The moments of the half-normal distribution can be computed explicitly. Once again we assume that X has the half-normal distribution with parameter $\sigma \in (0, \infty)$.

For $n \in \mathbb{N}$

$$\mathbb{E}(X^{2n}) = \sigma^{2n} \frac{(2n)!}{n!2^n} \quad (5.13.16)$$

$$\mathbb{E}(X^{2n+1}) = \sigma^{2n+1} 2^n \sqrt{\frac{2}{\pi}} n! \quad (5.13.17)$$

Proof

As in the [definition](#), we can take $X = \sigma|Z|$ where Z has the standard normal distribution. The even order moments of X are the same as the even order moments of σZ . These were computed in the section on the normal distribution. For the odd order moments we again use the simple substitution $z = x^2/2$ to get

$$\mathbb{E}(X^{2n+1}) = \sigma^{2n+1} \int_0^\infty x^{2n+1} \sqrt{\frac{2}{\pi}} e^{-x^2/2} dx = \sigma^{2n+1} 2^n \sqrt{\frac{2}{\pi}} \int_0^\infty z^n e^{-z} dz = \sigma^{2n+1} 2^n \sqrt{\frac{2}{\pi}} n! \quad (5.13.18)$$

In particular, we have $\mathbb{E}(X) = \sigma\sqrt{2/\pi}$ and $\text{var}(X) = \sigma^2(1 - 2/\pi)$

Open the special distribution simulator and select the folded normal distribution. Keep $\mu = 0$ and vary σ , and note the size and location of the mean \pm standard deviation bar. For selected values of σ , run the simulation 1000 times and compare the mean and standard deviation to the true mean and standard deviation.

Next are the skewness and kurtosis of the half-normal distribution.

Skewness and kurtosis

1. The skewness of X is

$$\text{skew}(X) = \frac{\sqrt{2/\pi}(4/\pi - 1)}{(1 - 2/\pi)^{3/2}} \approx 0.99527 \quad (5.13.19)$$

2. The kurtosis of X is

$$\text{kurt}(X) = \frac{3 - 4/\pi - 12/\pi^2}{(1 - 2/\pi)^2} \approx 3.8692 \quad (5.13.20)$$

Proof

Skewness and kurtosis are functions of the standard score and so do not depend on the scale parameter σ . The results then follow by letting $\sigma = 1$ and using the standard computational formulas for skewness and kurtosis in terms of the [moments of the half-normal distribution](#).

Related Distributions

Once again, the most important relation is the one in the [definition](#): If Y has a normal distribution with mean 0 then $X = |Y|$ has a half-normal distribution. Since the half normal distribution is a scale family, it is trivially closed under scale transformations.

Suppose that X has the half-normal distribution with parameter σ and that $b \in (0, \infty)$. Then bX has the half-normal distribution with parameter $b\sigma$.

Proof

As in the [definition](#), let $X = \sigma|Z|$ where Z is standard normal. Then $bX = b\sigma|Z|$.

The standard half-normal distribution is also a special case of the chi distribution.

The standard half-normal distribution is the chi distribution with 1 degree of freedom.

Proof

If Z is a standard normal variable, then Z^2 has the chi-square distribution with 1 degree of freedom, and hence $|Z| = \sqrt{Z^2}$ has the chi distribution with 1 degree of freedom.

This page titled [5.13: The Folded Normal Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.14: The Rayleigh Distribution

The *Rayleigh distribution*, named for William Strutt, Lord Rayleigh, is the distribution of the magnitude of a two-dimensional random vector whose coordinates are independent, identically distributed, mean 0 normal variables. The distribution has a number of applications in settings where magnitudes of normal variables are important.

The Standard Rayleigh Distribution

Definition

Suppose that Z_1 and Z_2 are independent random variables with standard normal distributions. The magnitude $R = \sqrt{Z_1^2 + Z_2^2}$ of the vector (Z_1, Z_2) has the *standard Rayleigh distribution*.

So in this definition, (Z_1, Z_2) has the standard bivariate normal distribution

Distribution Functions

We give five functions that completely characterize the standard Rayleigh distribution: the distribution function, the probability density function, the quantile function, the reliability function, and the failure rate function. For the remainder of this discussion, we assume that R has the standard Rayleigh distribution.

R has distribution function G given by $G(x) = 1 - e^{-x^2/2}$ for $x \in [0, \infty)$.

Proof

(Z_1, Z_2) has joint PDF $(z_1, z_2) \mapsto \frac{1}{2\pi} e^{-(z_1^2 + z_2^2)/2}$ on \mathbb{R}^2 . Hence

$$\mathbb{P}(R \leq x) = \int_{C_x} \frac{1}{2\pi} e^{-(z_1^2 + z_2^2)/2} d(z_1, z_2) \quad (5.14.1)$$

where $C_x = \{(z_1, z_2) \in \mathbb{R}^2 : z_1^2 + z_2^2 \leq x^2\}$. Convert to polar coordinates with $z_1 = r \cos \theta$, $z_2 = r \sin \theta$ to get

$$\mathbb{P}(R \leq x) = \int_0^{2\pi} \int_0^x \frac{1}{2\pi} e^{-r^2/2} r dr d\theta \quad (5.14.2)$$

The result now follows by simple integration.

R has probability density function g given by $g(x) = x e^{-x^2/2}$ for $x \in [0, \infty)$.

1. g increases and then decreases with mode at $x = 1$.
2. g is concave downward and then upward with inflection point at $x = \sqrt{3}$.

Proof

The formula for the PDF follows immediately from the [distribution function](#) since $g(x) = G'(x)$.

1. $g'(x) = e^{-x^2/2}(1 - x^2)$
2. $g''(x) = x e^{-x^2/2}(x^2 - 3)$.

Open the Special Distribution Simulator and select the Rayleigh distribution. Keep the default parameter value and note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

R has quantile function G^{-1} given by $G^{-1}(p) = \sqrt{-2 \ln(1-p)}$ for $p \in [0, 1)$. In particular, the quartiles of R are

1. $q_1 = \sqrt{4 \ln 2 - 2 \ln 3} \approx 0.7585$, the first quartile
2. $q_2 = \sqrt{2 \ln 2} \approx 1.1774$, the median
3. $q_3 = \sqrt{4 \ln 2} \approx 1.6651$, the third quartile

Proof

The formula for the quantile function follows immediately from the [distribution function](#) by solving $p = G(x)$ for x in terms of $p \in [0, 1)$.

Open the Special Distribution Calculator and select the Rayleigh distribution. Keep the default parameter value. Note the shape and location of the distribution function. Compute selected values of the distribution function and the quantile function.

R has reliability function G^c given by $G^c(x) = e^{-x^2/2}$ for $x \in [0, \infty)$.

Proof

Recall that the reliability function is simply the right-tail distribution function, so $G^c(x) = 1 - G(x)$.

R has failure rate function h given by $h(x) = x$ for $x \in [0, \infty)$. In particular, R has increasing failure rate.

Proof

Recall that the failure rate function is $h(x) = g(x)/G^c(x)$.

Moments

Once again we assume that R has the standard Rayleigh distribution. We can express the moment generating function of R in terms of the standard normal distribution function Φ . Recall that Φ is so commonly used that it is a special function of mathematics.

R has moment generating function m given by

$$m(t) = \mathbb{E}(e^{tR}) = 1 + \sqrt{2\pi}te^{t^2/2}\Phi(t), \quad t \in \mathbb{R} \quad (5.14.3)$$

Proof

By definition $m(t) = \int_0^\infty e^{tx} x e^{-x^2/2} dx$. Combining the exponential and completing the square in x gives

$$m(t) = e^{t^2/2} \int_0^\infty x e^{-(x-t)^2/2} dx = \sqrt{2\pi} \int_0^\infty \frac{1}{\sqrt{2\pi}} x e^{-(x-t)^2/2} dx \quad (5.14.4)$$

But $x \mapsto \frac{1}{\sqrt{2\pi}} e^{-(x-t)^2/2}$ is the PDF of the normal distribution with mean t and variance 1. The rest of the derivation follows from basic calculus.

The mean, variance of R are

1. $\mathbb{E}(R) = \sqrt{\pi/2} \approx 1.2533$
2. $\text{var}(R) = 2 - \pi/2$

Proof

1. Note that

$$\mathbb{E}(R) = \int_0^\infty x^2 e^{-x^2/2} dx = \sqrt{2\pi} \int_0^\infty x^2 \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \quad (5.14.5)$$

But $x \mapsto \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ is the PDF of the standard normal distribution. Hence the second integral is $\frac{1}{2}$ (since the variance of the standard normal distribution is 1).

2. An integration by parts gives

$$\mathbb{E}(R^2) = \int_0^\infty x^3 e^{-x^2/2} dx = 0 + 2 \int_0^\infty x e^{-x^2/2} dx = 2 \quad (5.14.6)$$

Numerically, $\mathbb{E}(R) \approx 1.2533$ and $\text{sd}(R) \approx 0.6551$.

Open the Special Distribution Simulator and select the Rayleigh distribution. Keep the default parameter value. Note the size and location of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation.

deviation to the true mean and standard deviation.

The general moments of R can be expressed in terms of the gamma function Γ .

$$\mathbb{E}(R^n) = 2^{n/2} \Gamma(1 + n/2) \text{ for } n \in \mathbb{N}.$$

Proof

The substitution $u = x^2/2$ gives

$$\mathbb{E}(R^n) = \int_0^\infty x^n x e^{-x^2/2} dx = \int_0^\infty (2u)^{n/2} e^{-u} du = 2^{n/2} \int_0^\infty u^{n/2} e^{-u} du \quad (5.14.7)$$

The last integral is $\Gamma(1 + n/2)$ by definition.

Of course, the formula for the general moments gives an alternate derivation of the [mean and variance above](#), since $\Gamma(3/2) = \sqrt{\pi}/2$ and $\Gamma(2) = 1$. On the other hand, the moment generating function can be also be used to derive the formula for the general moments.

The skewness and kurtosis of R are

1. $\text{skew}(R) = 2\sqrt{\pi}(\pi - 3)/(4 - \pi)^{3/2} \approx 0.6311$
2. $\text{kurt}(R) = (32 - 3\pi^2)/(4 - \pi)^2 \approx 3.2451$

Proof

These results follow from the standard formulas for the skewness and kurtosis in terms of the moments, since $\mathbb{E}(R) = \sqrt{\pi/2}$, $\mathbb{E}(R^2) = 2$, $\mathbb{E}(R^3) = 3\sqrt{2\pi}$, and $\mathbb{E}(R^4) = 8$.

Related Distributions

The fundamental connection between the standard Rayleigh distribution and the standard normal distribution is given in the very [definition](#) of the standard Rayleigh, as the distribution of the magnitude of a point with independent, standard normal coordinates.

Connections to the chi-square distribution.

1. If R has the standard Rayleigh distribution then R^2 has the chi-square distribution with 2 degrees of freedom.
2. If V has the chi-square distribution with 2 degrees of freedom then \sqrt{V} has the standard Rayleigh distribution.

Proof

This follows directly from the definition of the standard Rayleigh variable $R = \sqrt{Z_1^2 + Z_2^2}$, where Z_1 and Z_2 are independent standard normal variables.

Recall also that the chi-square distribution with 2 degrees of freedom is the same as the exponential distribution with scale parameter 2.

Since the [quantile function](#) is in closed form, the standard Rayleigh distribution can be simulated by the random quantile method.

Connections between the standard Rayleigh distribution and the standard uniform distribution.

1. If U has the standard uniform distribution (a random number) then $R = G^{-1}(U) = \sqrt{-2 \ln(1 - U)}$ has the standard Rayleigh distribution.
2. If R has the standard Rayleigh distribution then $U = G(R) = 1 - \exp(-R^2/2)$ has the standard uniform distribution

In part (a), note that $1 - U$ has the same distribution as U (the standard uniform). Hence $R = \sqrt{-2 \ln U}$ also has the standard Rayleigh distribution.

Open the random quantile simulator and select the Rayleigh distribution with the default parameter value (standard). Run the simulation 1000 times and compare the empirical density function to the true density function.

There is another connection with the uniform distribution that leads to the most common method of simulating a pair of independent standard normal variables. We have seen this before, but it's worth repeating. The result is closely related to the [definition](#) of the standard Rayleigh variable as the magnitude of a standard bivariate normal pair, but with the addition of the polar coordinate angle.

Suppose that R has the standard Rayleigh distribution, Θ is uniformly distributed on $[0, 2\pi)$ and that R and Θ are independent. Let $Z = R \cos \Theta$, $W = R \sin \Theta$. Then (Z, W) have the standard bivariate normal distribution.

Proof

By independence, the joint PDF f of (R, Θ) is given by

$$f(r, \theta) = r e^{-r^2/2} \frac{1}{2\pi}, \quad r \in [0, \infty), \theta \in [0, 2\pi) \quad (5.14.8)$$

As we recall from calculus, the Jacobian of the transformation $z = r \cos \theta$, $w = r \sin \theta$ is r , and hence the Jacobian of the inverse transformation that takes (z, w) into (r, θ) is $1/r$. Moreover, $r = \sqrt{z^2 + w^2}$. From the change of variables theorem, the PDF g of (Z, W) is given by $g(z, w) = f(r, \theta) \frac{1}{r}$. This leads to

$$g(z, w) = \frac{1}{2\pi} e^{-(z^2 + w^2)/2} = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} \frac{1}{\sqrt{2\pi}} e^{-w^2/2}, \quad z \in \mathbb{R}, w \in \mathbb{R} \quad (5.14.9)$$

Hence (Z, W) has the standard bivariate normal distribution.

The General Rayleigh Distribution

Definition

The standard Rayleigh distribution is generalized by adding a scale parameter.

If R has the standard Rayleigh distribution and $b \in (0, \infty)$ then $X = bR$ has the *Rayleigh distribution* with scale parameter b .

Equivalently, the Rayleigh distribution is the distribution of the magnitude of a two-dimensional vector whose components have independent, identically distributed mean 0 normal variables.

If U_1 and U_2 are independent normal variables with mean 0 and standard deviation $\sigma \in (0, \infty)$ then $X = \sqrt{U_1^2 + U_2^2}$ has the Rayleigh distribution with scale parameter σ .

Proof

We can take $U_1 = \sigma Z_1$ and $U_2 = \sigma Z_2$ where Z_1 and Z_2 are independent standard normal variables. Then $X = \sigma \sqrt{Z_1^2 + Z_2^2} = \sigma R$ where R has the standard Rayleigh distribution.

Distribution Functions

In this section, we assume that X has the Rayleigh distribution with scale parameter $b \in (0, \infty)$.

X has cumulative distribution function F given by $F(x) = 1 - \exp\left(-\frac{x^2}{2b^2}\right)$ for $x \in [0, \infty)$.

Proof

Recall that $F(x) = G(x/b)$ where G is the [standard Rayleigh CDF](#).

X has probability density function f given by $f(x) = \frac{x}{b^2} \exp\left(-\frac{x^2}{2b^2}\right)$ for $x \in [0, \infty)$.

1. f increases and then decreases with mode at $x = b$.
2. f is concave downward and then upward with inflection point at $x = \sqrt{3}b$.

Proof

Recall that $f(x) = \frac{1}{b} g\left(\frac{x}{b}\right)$ where g is the [standard Rayleigh PDF](#).

Open the Special Distribution Simulator and select the Rayleigh distribution. Vary the scale parameter and note the shape and location of the probability density function. For various values of the scale parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has quantile function F^{-1} given by $F^{-1}(p) = b\sqrt{-2\ln(1-p)}$ for $p \in [0, 1)$. In particular, the quartiles of X are

1. $q_1 = b\sqrt{4\ln 2 - 2\ln 3}$, the first quartile
2. $q_2 = b\sqrt{2\ln 2}$, the median
3. $q_3 = b\sqrt{4\ln 2}$, the third quartile

Proof

Recall that $F^{-1}(p) = bG^{-1}(p)$ where G^{-1} is the [standard Rayleigh quantile function](#).

Open the Special Distribution Calculator and select the Rayleigh distribution. Vary the scale parameter and note the location and shape of the distribution function. For various values of the scale parameter, compute selected values of the distribution function and the quantile function.

X has reliability function F^c given by $F^c(x) = \exp\left(-\frac{x^2}{2b^2}\right)$ for $x \in [0, \infty)$.

Proof

Recall that $F^c(x) = 1 - F(x)$.

X has failure rate function h given by $h(x) = x/b^2$ for $x \in [0, \infty)$. In particular, X has increasing failure rate.

Proof

Recall that $h(x) = f(x)/H(x)$.

Moments

Again, we assume that X has the Rayleigh distribution with scale parameter b , and recall that Φ denotes the standard normal distribution function.

X has moment generating function M given by

$$M(t) = \mathbb{E}(e^{tX}) = 1 + \sqrt{2\pi}bt \exp\left(\frac{b^2 t^2}{2}\right) \Phi(t), \quad t \in \mathbb{R} \quad (5.14.10)$$

Proof

Recall that $M(t) = m(bt)$ where m is the [standard Rayleigh MGF](#).

The mean and variance of R are

1. $\mathbb{E}(X) = b\sqrt{\pi/2}$
2. $\text{var}(X) = b^2(2 - \pi/2)$

Proof

These result follow from [standard mean and variance](#) and basic properties of expected value and variance.

Open the Special Distribution Simulator and select the Rayleigh distribution. Vary the scale parameter and note the size and location of the mean±standard deviation bar. For various values of the scale parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

Again, the general moments can be expressed in terms of the gamma function Γ .

$\mathbb{E}(X^n) = b^n 2^{n/2} \Gamma(1 + n/2)$ for $n \in \mathbb{N}$.

Proof

This follows from the [standard moments](#) and basic properties of expected value.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 2\sqrt{\pi}(\pi - 3)/(4 - \pi)^{3/2} \approx 0.6311$
2. $\text{kurt}(X) = (32 - 3\pi^2)/(4 - \pi)^2 \approx 3.2451$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are unchanged by a scale transformation. Thus the results follow from the [standard skewness and kurtosis](#).

Related Distributions

The fundamental connection between the Rayleigh distribution and the normal distribution is the [definition](#), and of course, is the primary reason that the Rayleigh distribution is *special* in the first place. By construction, the Rayleigh distribution is a scale family, and so is closed under scale transformations.

If X has the Rayleigh distribution with scale parameter $b \in (0, \infty)$ and if $c \in (0, \infty)$ then cX has the Rayleigh distribution with scale parameter bc .

The Rayleigh distribution is a special case of the Weibull distribution.

The Rayleigh distribution with scale parameter $b \in (0, \infty)$ is the Weibull distribution with shape parameter 2 and scale parameter $\sqrt{2}b$.

The following result generalizes the connection between the [standard Rayleigh and chi-square distributions](#).

If X has the Rayleigh distribution with scale parameter $b \in (0, \infty)$ then X^2 has the exponential distribution with scale parameter $2b^2$.

Proof

We can take $X = bR$ where R has the standard Rayleigh distribution. Then $X^2 = b^2 R^2$, and R^2 has the exponential distribution with scale parameter 2. Hence X^2 has the exponential distribution with scale parameter $2b^2$.

Since the [quantile function](#) is in closed form, the Rayleigh distribution can be simulated by the random quantile method.

Suppose that $b \in (0, \infty)$.

1. If U has the standard uniform distribution (a random number) then $X = F^{-1}(U) = b\sqrt{-2\ln(1-U)}$ has the Rayleigh distribution with scale parameter b .
2. If X has the Rayleigh distribution with scale parameter b then $U = F(X) = 1 - \exp(-X^2/2b^2)$ has the standard uniform distribution

In part (a), note that $1 - U$ has the same distribution as U (the standard uniform). Hence $X = b\sqrt{-2\ln U}$ also has the Rayleigh distribution with scale parameter b .

Open the random quantile simulator and select the Rayleigh distribution. For selected values of the scale parameter, run the simulation 1000 times and compare the empirical density function to the true density function.

Finally, the Rayleigh distribution is a member of the general exponential family.

If X has the Rayleigh distribution with scale parameter $b \in (0, \infty)$ then X has a one-parameter exponential distribution with natural parameter $-1/b^2$ and natural statistic $X^2/2$.

Proof

This follows directly from the definition of the general exponential distribution.

This page titled [5.14: The Rayleigh Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.15: The Maxwell Distribution

The *Maxwell distribution*, named for James Clerk Maxwell, is the distribution of the magnitude of a three-dimensional random vector whose coordinates are independent, identically distributed, mean 0 normal variables. The distribution has a number of applications in settings where magnitudes of normal variables are important, particularly in physics. It is also called the *Maxwell-Boltzmann* distribution in honor also of Ludwig Boltzmann. The Maxwell distribution is closely related to the Rayleigh distribution, which governs the magnitude of a two-dimensional random vector whose coordinates are independent, identically distributed, mean 0 normal variables.

The Standard Maxwell Distribution

Definition

Suppose that Z_1 , Z_2 , and Z_3 are independent random variables with standard normal distributions. The magnitude $R = \sqrt{Z_1^2 + Z_2^2 + Z_3^2}$ of the vector (Z_1, Z_2, Z_3) has the *standard Maxwell distribution*.

So in the context of the definition, (Z_1, Z_2, Z_3) has the standard trivariate normal distribution. The Maxwell distribution is a continuous distribution on $[0, \infty)$.

Distribution Functions

In this discussion, we assume that R has the standard Maxwell distribution. The distribution function of R can be expressed in terms of the standard normal distribution function Φ . Recall that Φ occurs so frequently that it is considered a special function in mathematics.

R has distribution function G given by

$$G(x) = 2\Phi(x) - \sqrt{\frac{2}{\pi}} x e^{-x^2/2} - 1, \quad x \in [0, \infty) \quad (5.15.1)$$

Proof

(Z_1, Z_2, Z_3) has joint PDF $(z_1, z_2, z_3) \mapsto \frac{1}{(2\pi)^{3/2}} e^{-(z_1^2 + z_2^2 + z_3^2)/2}$ on \mathbb{R}^3 . Hence

$$\mathbb{P}(R \leq x) = \int_{B_x} \frac{1}{(2\pi)^{3/2}} e^{-(z_1^2 + z_2^2 + z_3^2)/2} d(z_1, z_2, z_3), \quad x \in [0, \infty) \quad (5.15.2)$$

where $B_x = \{(z_1, z_2, z_3) \in \mathbb{R}^3 : z_1^2 + z_2^2 + z_3^2 \leq x^2\}$, the spherical region of radius x centered at the origin. Convert to spherical coordinates with $z_1 = \rho \sin \phi \cos \theta$, $z_2 = \rho \sin \phi \sin \theta$, $z_3 = \rho \cos \phi$ to get

$$\mathbb{P}(R \leq x) = \int_0^\pi \int_0^{2\pi} \int_0^x \frac{1}{(2\pi)^{3/2}} e^{-\rho^2/2} \rho^2 \sin \phi d\rho d\theta d\phi, \quad x \in [0, \infty) \quad (5.15.3)$$

The result now follows by simple integration.

R has probability density function g given by

$$g(x) = \sqrt{\frac{2}{\pi}} x^2 e^{-x^2/2}, \quad x \in [0, \infty) \quad (5.15.4)$$

1. g increases and then decreases with mode at $x = \sqrt{2}$.
2. g is concave upward, then downward, then upward again, with inflection points at $x_1 = \sqrt{(5 - \sqrt{17})/2} \approx 0.6622$ and $x_2 = \sqrt{(5 + \sqrt{17})/2} \approx 2.1358$

Proof

The formula for the PDF follows immediately from the [distribution function](#) since $g(x) = G'(x)$.

1. $g'(x) = \sqrt{2/\pi} x e^{-x^2/2} (2 - x^2)$
2. $g''(x) = \sqrt{2/\pi} e^{-x^2/2} (x^4 - 5x^2 + 2)$

Open the Special Distribution Simulator and select the Maxwell distribution. Keep the default parameter value and note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function has no simple closed-form expression.

Open the Special Distribution Calculator and select the Maxwell distribution. Keep the default parameter value. Find approximate values of the median and the first and third quartiles.

Moments

Suppose again that R has the standard Maxwell distribution. The moment generating function of R , like the distribution function, can be expressed in terms of the standard normal distribution function Φ .

R has moment generating function m given by

$$m(t) = \mathbb{E}(e^{tR}) = \sqrt{\frac{2}{\pi}} t + 2(1+t^2)e^{t^2/2}\Phi(t), \quad t \in \mathbb{R} \quad (5.15.5)$$

Proof

Completing the square in x gives

$$m(t) = \int_0^\infty \sqrt{\frac{2}{\pi}} x^2 e^{-x^2/2} e^{tx} dx = \sqrt{\frac{2}{\pi}} e^{t^2/2} \int_0^\infty x^2 e^{-(x-t)^2/2} dx \quad (5.15.6)$$

The substitution $z = x - t$ gives

$$m(t) = \sqrt{\frac{2}{\pi}} e^{t^2/2} \int_{-t}^\infty (z+t)^2 e^{-z^2/2} dz = \sqrt{\frac{2}{\pi}} e^{t^2/2} \int_{-t}^\infty (z^2 + 2tz + t^2) e^{-z^2/2} dz \quad (5.15.7)$$

Integrating by parts or by simple substitution, using the fact that $z \mapsto \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ is the standard normal PDF, and that $1 - \Phi(-t) = \Phi(t)$ we have

$$\int_{-t}^\infty z^2 e^{-z^2/2} dz = -te^{-t^2/2} + \sqrt{2\pi}\Phi(t) \quad (5.15.8)$$

$$\int_{-t}^\infty 2tz e^{-z^2/2} dz = 2te^{-t^2/2} \quad (5.15.9)$$

$$\int_{-t}^\infty t^2 e^{-z^2/2} dz = t^2 \sqrt{2\pi}\Phi(t) \quad (5.15.10)$$

Simplifying gives the result.

The mean and variance of R can be found from the moment generating function, but direct computations are also easy.

The mean and variance of R are

1. $\mathbb{E}(R) = 2\sqrt{2/\pi}$
2. $\text{var}(R) = 3 - 8/\pi$

Proof

The integration methods are by parts and by simple substitution.

$$\mathbb{E}(R) = \int_0^\infty \sqrt{\frac{2}{\pi}} x^3 e^{-x^2/2} dx = 2\sqrt{\frac{2}{\pi}} \int_0^\infty x e^{-x^2/2} dx = 2\sqrt{\frac{2}{\pi}} \quad (5.15.11)$$

$$\mathbb{E}(R^2) = \int_0^\infty \sqrt{\frac{2}{\pi}} x^4 e^{-x^2/2} dx = 3 \int_0^\infty \sqrt{\frac{2}{\pi}} x^2 e^{-x^2/2} dx = 3 \quad (5.15.12)$$

Numerically, $\mathbb{E}(R) \approx 1.5958$ and $\text{sd}(R) \approx 0.6734$

Open the Special Distribution Simulator and select the Maxwell distribution. Keep the default parameter value. Note the size and location of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The general moments of R can be expressed in terms of the gamma function Γ

For $n \in \mathbb{N}_+$,

$$\mathbb{E}(R^n) = \frac{2^{n/2+1}}{\sqrt{\pi}} \Gamma\left(\frac{n+3}{2}\right) \quad (5.15.13)$$

Proof

The substitution $u = x^2/2$ gives

$$\mathbb{E}(R^n) = \int_0^\infty \sqrt{\frac{2}{\pi}} x^{n+1} x e^{-x^2/2} dx = \int_0^\infty \sqrt{\frac{2}{\pi}} (2u)^{(n+1)/2} e^{-u} du = \frac{2^{n/2+1}}{\sqrt{\pi}} \int_0^\infty u^{(n+1)/2} e^{-u} du \quad (5.15.14)$$

The last integral is $\Gamma[(n+3)/2]$ by definition.

Of course, the formula for the general moments gives an alternate derivation for the [mean and variance](#) above since $\Gamma(2) = 1$ and $\Gamma(5/2) = 3\sqrt{\pi}/4$. On the other hand, the moment generating function can be also be used to derive the formula for the general moments. Finally, we give the skewness and kurtosis of R .

The skewness and kurtosis of R are

1. $\text{skew}(R) = 2\sqrt{2}(16 - 5\pi)/(3\pi - 8)^{3/2} \approx 0.4857$
2. $\text{kurt}(R) = (15\pi^2 + 16\pi - 192)/(3\pi - 8)^2 \approx 3.1082$

Proof

These results follow from the standard formulas for the skewness and kurtosis in terms of the moments, since $\mathbb{E}(R) = 2\sqrt{2/\pi}$, $\mathbb{E}(R^2) = 3$, $\mathbb{E}(R^3) = 8\sqrt{2/\pi}$, and $\mathbb{E}(R^4) = 15$.

Related Distributions

The fundamental connection between the standard Maxwell distribution and the standard normal distribution is given in the very [definition](#) of the standard Maxwell, as the distribution of the magnitude of a vector in \mathbb{R}^3 with independent, standard normal coordinates.

Connections to the chi-square distribution.

1. If R has the standard Maxwell distribution then R^2 has the chi-square distribution with 3 degrees of freedom.
2. If V has the chi-square distribution with 3 degrees of freedom then \sqrt{V} has the standard Maxwell distribution.

Proof

This follows directly from the definition of the standard Maxwell variable $R = \sqrt{Z_1^2 + Z_2^2 + Z_3^2}$, where Z_1 , Z_2 , and Z_3 are independent standard normal variables.

Equivalently, the Maxwell distribution is simply the chi distribution with 3 degrees of freedom.

The General Maxwell Distribution

Definition

The standard Maxwell distribution is generalized by adding a scale parameter.

If R has the standard Maxwell distribution and $b \in (0, \infty)$ then $X = bR$ has the *Maxwell distribution* with scale parameter b .

Equivalently, the Maxwell distribution is the distribution of the magnitude of a three-dimensional vector whose components have independent, identically distributed, mean 0 normal variables.

If U_1, U_2 and U_3 are independent normal variables with mean 0 and standard deviation $\sigma \in (0, \infty)$ then $X = \sqrt{U_1^2 + U_2^2 + U_3^2}$ has the Maxwell distribution with scale parameter σ .

Proof

We can take $U_i = \sigma Z_i$ for $i \in \{1, 2, 3\}$ where Z_1, Z_2 , and Z_3 are independent standard normal variables. Then $X = \sigma \sqrt{Z_1^2 + Z_2^2 + Z_3^2} = \sigma R$ where R has the standard Maxwell distribution.

Distribution Functions

In this section, we assume that X has the Maxwell distribution with scale parameter $b \in (0, \infty)$. We can give the distribution function of X in terms of the standard normal distribution function Φ .

X has distribution function F given by

$$F(x) = 2\Phi\left(\frac{x}{b}\right) - \frac{1}{b}\sqrt{\frac{2}{\pi}}x \exp\left(-\frac{x^2}{2b^2}\right) - 1, \quad x \in [0, \infty) \quad (5.15.15)$$

Proof

Recall that $F(x) = G(x/b)$ where G is the [standard Maxwell CDF](#).

X has probability density function f given by

$$f(x) = \frac{1}{b^3} \sqrt{\frac{2}{\pi}} x^2 \exp\left(-\frac{x^2}{2b^2}\right), \quad x \in [0, \infty) \quad (5.15.16)$$

1. f increases and then decreases with mode at $x = b\sqrt{2}$.
2. f is concave upward, then downward, then upward again, with inflection points at $x = b\sqrt{(5 \pm \sqrt{17})/2}$.

Proof

Recall that $f(x) = \frac{1}{b}g\left(\frac{x}{b}\right)$ where g is the [standard Maxwell PDF](#).

Open the Special Distribution Simulator and select the Maxwell distribution. Vary the scale parameter and note the shape and location of the probability density function. For various values of the scale parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

Again, the quantile function does not have a simple, closed-form expression.

Open the Special Distribution Calculator and select the Maxwell distribution. For various values of the scale parameter, compute the median and the first and third quartiles.

Moments

Again, we assume that X has the Maxwell distribution with scale parameter $b \in (0, \infty)$. As before, the moment generating function of X can be written in terms of the standard normal distribution function Φ .

X has moment generating function M given by

$$M(t) = \mathbb{E}(e^{tX}) = \sqrt{\frac{2}{\pi}}bt + 2(1 + b^2t^2) \exp\left(\frac{b^2t^2}{2}\right) \Phi(bt), \quad t \in \mathbb{R} \quad (5.15.17)$$

Proof

Recall that $M(t) = m(bt)$ where m is the [standard Maxwell MGF](#).

The mean and variance of X are

1. $\mathbb{E}(X) = b2\sqrt{2/\pi}$
2. $\text{var}(X) = b^2(3 - 8/\pi)$

Proof

These result follow from the [standard mean and variance](#) and basic properties of expected value and variance.

Open the Special Distribution Simulator and select the Maxwell distribution. Vary the scale parameter and note the size and location of the mean \pm standard deviation bar. For various values of the scale parameter, run the simulation 1000 times compare the empirical mean and standard deviation to the true mean and standard deviation.

As before, the general moments can be expressed in terms of the gamma function Γ .

For $n \in \mathbb{N}$,

$$\mathbb{E}(X^n) = b^n \frac{2^{n/2+1}}{\sqrt{\pi}} \Gamma\left(\frac{n+3}{2}\right) \quad (5.15.18)$$

Proof

This follows from the [standard moments](#) and basic properties of expected value.

Finally, the skewness and kurtosis are unchanged.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 2\sqrt{2}(16 - 5\pi)/(3\pi - 8)^{3/2} \approx 0.4857$
2. $\text{kurt}(X) = (15\pi^2 + 16\pi - 192)/(3\pi - 8)^2 \approx 3.1082$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are unchanged by a scale transformation. Thus the results follow from the [standard skewness and kurtosis](#).

Related Distributions

The fundamental connection between the Maxwell distribution and the normal distribution is given in the [definition](#), and of course, is the primary reason that the Maxwell distribution is *special* in the first place.

By construction, the Maxwell distribution is a scale family, and so is closed under scale transformations.

If X has the Maxwell distribution with scale parameter $b \in (0, \infty)$ and if $c \in (0, \infty)$ then cX has the Maxwell distribution with scale parameter bc .

Proof

By [definition](#), we can assume that $X = bR$ where R has the standard Maxwell distribution. Hence $cX = (cb)R$ has the Maxwell distribution with scale parameter bc .

The Maxwell distribution is a generalized exponential distribution.

If X has the Maxwell distribution with scale parameter $b \in (0, \infty)$ then X is a one-parameter exponential family with natural parameter $-1/b^2$ and natural statistic $X^2/2$.

Proof

This follows directly from the definition of the general exponential distribution, and the form of the [PDF](#).

This page titled [5.15: The Maxwell Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.16: The Lévy Distribution

The *Lévy distribution*, named for the French mathematician Paul Lévy, is important in the study of [Brownian motion](#), and is one of only three stable distributions whose probability density function can be expressed in a simple, closed form.

The Standard Lévy Distribution

Definition

If Z has the standard normal distribution then $U = 1/Z^2$ has the *standard Lévy distribution*.

So the standard Lévy distribution is a continuous distribution on $(0, \infty)$.

Distribution Functions

We assume that U has the standard Lévy distribution. The distribution function of U has a simple expression in terms of the standard normal distribution function Φ , not surprising given the definition.

U has distribution function G given by

$$G(u) = 2 \left[1 - \Phi \left(\frac{1}{\sqrt{u}} \right) \right], \quad u \in (0, \infty) \quad (5.16.1)$$

Proof

For $u \in (0, \infty)$,

$$\mathbb{P} \left(\frac{1}{Z^2} \leq u \right) = \mathbb{P} \left(Z^2 \geq \frac{1}{u} \right) = \mathbb{P} \left(Z \geq \frac{1}{\sqrt{u}} \right) + \mathbb{P} \left(Z \leq -\frac{1}{\sqrt{u}} \right) = 2 \left[1 - \Phi \left(\frac{1}{\sqrt{u}} \right) \right] \quad (5.16.2)$$

Similarly, the quantile function of U has a simple expression in terms of the standard normal quantile function Φ^{-1} .

U has quantile function G^{-1} given by

$$G^{-1}(p) = \frac{1}{[\Phi^{-1}(1 - p/2)]^2}, \quad p \in [0, 1] \quad (5.16.3)$$

The quartiles of U are

1. $q_1 = [\Phi^{-1}(\frac{7}{8})]^{-2} \approx 0.7557$, the first quartile.
2. $q_2 = [\Phi^{-1}(\frac{3}{4})]^{-2} \approx 2.1980$, the median.
3. $q_3 = [\Phi^{-1}(\frac{5}{8})]^{-2} \approx 9.8516$, the third quartile.

Proof

The quantile function can be obtained from the distribution function by solving $p = G(u)$ for $u = G^{-1}(p)$.

Open the [Special Distribution Calculator](#) and select the Lévy distribution. Keep the default parameter values. Note the shape and location of the distribution function. Compute a few values of the distribution function and the quantile function.

Finally, the probability density function of U has a simple closed expression.

U has probability density function g given by

$$g(u) = \frac{1}{\sqrt{2\pi}} \frac{1}{u^{3/2}} \exp\left(-\frac{1}{2u}\right), \quad u \in (0, \infty) \quad (5.16.4)$$

1. g increases and then decreasing with mode at $x = \frac{1}{3}$.

2. g is concave upward, then downward, then upward again, with inflection points at $x = \frac{1}{3} - \frac{\sqrt{10}}{15} \approx 0.1225$ and at $x = \frac{1}{3} + \frac{\sqrt{10}}{15} \approx 0.5442$.

Proof

The formula for g follows from differentiating the CDF given above:

$$g(u) = -2\Phi'(u^{-1/2}) \left(-\frac{1}{2}u^{-3/2} \right), \quad u \in (0, \infty) \quad (5.16.5)$$

But $\Phi'(z) = \phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$, the standard normal PDF. Substitution and simplification then gives the results. Parts (a) and (b) also follow from standard calculus:

$$g'(u) = \frac{1}{2\sqrt{2\pi}} u^{-7/2} e^{-u^{-1/2}} (-3u + 1) \quad (5.16.6)$$

$$g''(u) = \frac{1}{4\sqrt{2\pi}} u^{-11/2} e^{-u^{-1/2}} (15u^2 - 10u + 1) \quad (5.16.7)$$

Open the Special Distribution Simulator and select the Lévy distribution. Keep the default parameter values. Note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

Moments

We assume again that U has the standard Lévy distribution. After exploring the graphs of the probability density function and distribution function above, you probably noticed that the Lévy distribution has a very heavy tail. The 99th percentile is about 6400, for example. The following result is not surprising.

$$\mathbb{E}(U) = \infty$$

Proof

Note that $u \mapsto e^{-1/2u}$ is increasing. Hence

$$\mathbb{E}(U) = \int_0^\infty u \frac{1}{\sqrt{2\pi}u^{3/2}} e^{-1/2u} du > \int_1^\infty \frac{1}{\sqrt{2\pi}} u^{-1/2} e^{-1/2u} du = \infty \quad (5.16.8)$$

Of course, the higher-order moments are infinite as well, and the variance, skewness, and kurtosis do not exist. The moment generating function is infinite at every positive value, and so is of no use. On the other hand, the characteristic function of the standard Lévy distribution is very useful. For the following result, recall that the *sign* function sgn is given by $\text{sgn}(t) = 1$ for $t > 0$, $\text{sgn}(t) = -1$ for $t < 0$, and $\text{sgn}(0) = 0$.

U has characteristic function χ_0 given by

$$\chi_0(t) = \mathbb{E}(e^{itU}) = \exp\left(-|t|^{1/2} [1 + i \text{sgn}(t)]\right), \quad t \in \mathbb{R} \quad (5.16.9)$$

Related Distributions

The most important relationship is the one in the definition: If Z has the standard normal distribution then $U = 1/Z^2$ has the standard Lévy distribution. The following result is basically the converse.

If U has the standard Lévy distribution, then $V = 1/\sqrt{U}$ has the standard half-normal distribution.

Proof

From the definition, we can take $U = 1/Z^2$ where Z has the standard normal distribution. Then $1/\sqrt{U} = |Z|$, and $|Z|$ has the standard half-normal distribution.

The General Lévy Distribution

Like so many other “standard distributions”, the standard Lévy distribution is generalized by adding location and scale parameters.

Definition

Suppose that U has the standard Lévy distribution, and $a \in \mathbb{R}$ and $b \in (0, \infty)$. Then $X = a + bU$ has the *Lévy distribution* with location parameter a and scale parameter b .

Note that X has a continuous distribution on the interval (a, ∞) .

Distribution Functions

Suppose that X has the Lévy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. As before, the distribution function of X has a simple expression in terms of the standard normal distribution function Φ .

X has distribution function G given by

$$F(x) = 2 \left[1 - \Phi \left(\sqrt{\frac{b}{x-a}} \right) \right], \quad x \in (a, \infty) \quad (5.16.10)$$

Proof

Recall that $F(x) = G\left(\frac{x-a}{b}\right)$ where G is the standard Lévy CDF.

Similarly, the quantile function of X has a simple expression in terms of the standard normal quantile function Φ^{-1} .

X has quantile function F^{-1} given by

$$F^{-1}(p) = a + \frac{b}{[\Phi^{-1}(1-p/2)]^2}, \quad p \in [0, 1) \quad (5.16.11)$$

The quartiles of X are

1. $q_1 = a + b[\Phi^{-1}(\frac{7}{8})]^{-2}$, the first quartile.
2. $q_2 = a + b[\Phi^{-1}(\frac{3}{4})]^{-2}$, the median.
3. $q_3 = a + b[\Phi^{-1}(\frac{5}{8})]^{-2}$, the third quartile.

Proof

Recall that $F^{-1}(p) = a + bG^{-1}(p)$, where G^{-1} is the standard Lévy quantile function.

Open the Special Distribution Calculator and select the Lévy distribution. Vary the parameter values and note the shape of the graph of the distribution function. For various values of the parameters, compute a few values of the distribution function and the quantile function.

Finally, the probability density function of X has a simple closed expression.

X has probability density function f given by

$$f(x) = \sqrt{\frac{b}{2\pi}} \frac{1}{(x-a)^{3/2}} \exp\left[-\frac{b}{2(x-a)}\right], \quad x \in (a, \infty) \quad (5.16.12)$$

1. f increases and then decreases with mode at $x = a + \frac{1}{3}b$.
2. f is concave upward, then downward, then upward again with inflection points at $x = a + \left(\frac{1}{3} \pm \frac{\sqrt{10}}{15}\right)b$.

Proof

Recall that $f(x) = \frac{1}{b}g\left(\frac{x-a}{b}\right)$ where g is the standard Lévy PDF, so the formula for f follow from the definition of g and simple algebra. Parts (a) and (b) follow from the corresponding results for g .

Open the Special Distribution Simulator and select the Lévy distribution. Vary the parameters and note the shape and location of the probability density function. For various parameter values, run the simulation 1000 times and compare the empirical density function to the probability density function.

Moments

Assume again that X has the Lévy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. Of course, since the standard Lévy distribution has infinite mean, so does the general Lévy distribution.

$$\mathbb{E}(X) = \infty$$

Also as before, the variance, skewness, and kurtosis of X are undefined. On the other hand, the characteristic function of X is very important.

X has characteristic function χ given by

$$\chi(t) = \mathbb{E}(e^{itX}) = \exp\left(ita - b^{1/2}|t|^{1/2}[1 + i \operatorname{sgn}(t)]\right), \quad t \in \mathbb{R} \quad (5.16.13)$$

Proof

This follows from the standard characteristic function since $\chi(t) = e^{ita}\chi_0(bt)$. Note that $\operatorname{sgn}(bt) = \operatorname{sgn}(t)$ since $b > 0$.

Related Distributions

Since the Lévy distribution is a location-scale family, it is trivially closed under location-scale transformations.

Suppose that X has the Lévy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, and that $c \in \mathbb{R}$ and $d \in (0, \infty)$. Then $Y = c + dX$ has the Lévy distribution with location parameter $c + ad$ and scale parameter bd .

Proof

From the definition, we can take $X = a + bU$ where U has the standard Lévy distribution. Hence $Y = c + dX = (c + ad) + (bd)U$ has the Lévy distribution with location parameter $c + ad$ and scale parameter bd .

Of more interest is the fact that the Lévy distribution is closed under convolution (corresponding to sums of independent variables).

Suppose that X_1 and X_2 are independent, and that, X_k has the Lévy distribution with location parameter $a_k \in \mathbb{R}$ and scale parameter $b_k \in (0, \infty)$ for $k \in \{1, 2\}$. Then $X_1 + X_2$ has the Lévy distribution with location parameter $a_1 + a_2$ and scale parameter $(b_1^{1/2} + b_2^{1/2})^2$.

Proof

The characteristic function of X_k is

$$\chi_k(t) = \exp\left(ita_k - b_k^{1/2}|t|^{1/2}[1 + i \operatorname{sgn}(t)]\right), \quad t \in \mathbb{R} \quad (5.16.14)$$

for $k \in \{1, 2\}$. Hence the characteristic function of $X_1 + X_2$ is

$$\begin{aligned} \chi(t) &= \chi_1(t)\chi_2(t) = \exp\left[it(a_1 + a_2) - (b_1^{1/2} + b_2^{1/2})|t|^{1/2}[1 + i \operatorname{sgn}(t)]\right] \\ &= \exp\left[itA - B^{1/2}|t|^{1/2}[1 + i \operatorname{sgn}(t)]\right], \quad t \in \mathbb{R} \end{aligned}$$

where $A = a_1 + a_2$ is the location parameter and $B = (b_1^{1/2} + b_2^{1/2})^2$ is the scale parameter.

As a corollary, the Lévy distribution is a stable distribution with index $\alpha = \frac{1}{2}$:

Suppose that $n \in \mathbb{N}_+$ and that (X_1, X_2, \dots, X_n) is a sequence of independent random variables, each having the Lévy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. Then $X_1 + X_2 + \dots + X_n$ has the Lévy distribution with location parameter na and scale parameter n^2b .

Stability is one of the reasons for the importance of the Lévy distribution. From the characteristic function, it follows that the skewness parameter is $\beta = 1$.

This page titled [5.16: The Lévy Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.17: The Beta Distribution

In this section, we will study the beta distribution, the most important distribution that has bounded support. But before we can study the beta *distribution* we must study the beta *function*.

The Beta Function

Definition

The beta function B is defined as follows:

$$B(a, b) = \int_0^1 u^{a-1} (1-u)^{b-1} du; \quad a, b \in (0, \infty) \quad (5.17.1)$$

Proof that B is well defined

We need to show that $B(a, b) < \infty$ for every $a, b \in (0, \infty)$. The integrand is positive on $(0, 1)$, so the integral exists, either as a real number or ∞ . If $a \geq 1$ and $b \geq 1$, the integrand is continuous on $[0, 1]$, so of course the integral is finite. Thus, the only cases of interest are when $0 < a < 1$ or $0 < b < 1$. Note that

$$\int_0^1 u^{a-1} (1-u)^{b-1} du = \int_0^{1/2} u^{a-1} (1-u)^{b-1} du + \int_{1/2}^1 u^{a-1} (1-u)^{b-1} du \quad (5.17.2)$$

If $0 < a < 1$, $(1-u)^{b-1}$ is bounded on $(0, \frac{1}{2}]$ and $\int_0^{1/2} u^{a-1} du = \frac{1}{a2^a}$. Hence the first integral on the right in the displayed equation is finite. Similarly, If $0 < b < 1$, u^{a-1} is bounded on $[\frac{1}{2}, 1)$ and $\int_{1/2}^1 (1-u)^{b-1} du = \frac{1}{b2^b}$. Hence the second integral on the right in the displayed equation is also finite.

The beta function was first introduced by Leonhard Euler.

Properties

The beta function satisfies the following properties:

1. $B(a, b) = B(b, a)$ for $a, b \in (0, \infty)$, so B is symmetric.
2. $B(a, 1) = \frac{1}{a}$ for $a \in (0, \infty)$
3. $B(1, b) = \frac{1}{b}$ for $b \in (0, \infty)$

Proof

1. Using the substitution $v = 1 - u$ we have

$$B(a, b) = \int_0^1 u^{a-1} (1-u)^{b-1} du = \int_0^1 (1-v)^{a-1} v^{b-1} dv = B(b, a) \quad (5.17.3)$$

2. $B(a, 1) = \int_0^1 u^{a-1} du = \frac{1}{a}$
3. This follows from (a) and (b).

The beta function has a simple expression in terms of the gamma function:

If $a, b \in (0, \infty)$ then

$$B(a, b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)} \quad (5.17.4)$$

Proof

From the definitions, we can express $\Gamma(a+b)B(a, b)$ as a double integral:

$$\Gamma(a+b)B(a, b) = \int_0^\infty x^{a+b-1} e^{-x} dx \int_0^1 y^{a-1} (1-y)^{b-1} dy = \int_0^\infty \int_0^1 (xy)^{a-1} [x(1-y)]^{b-1} x e^{-x} dx dy \quad (5.17.5)$$

Next we use the transformation $w = xy$, $z = x(1 - y)$ which maps $(0, \infty) \times (0, 1)$ one-to-one onto $(0, \infty) \times (0, \infty)$. The inverse transformation is $x = w + z$, $y = w/(w + z)$ and the absolute value of the Jacobian is

$$\left| \det \frac{\partial(x, y)}{\partial(w, z)} \right| = \frac{1}{(w + z)} \quad (5.17.6)$$

Thus, using the change of variables theorem for multiple integrals, the integral above becomes

$$\int_0^\infty \int_0^\infty w^{a-1} z^{b-1} (w + z) e^{-(w+z)} \frac{1}{w + z} dw dz \quad (5.17.7)$$

which after simplifying is $\Gamma(a)\Gamma(b)$.

Recall that the gamma function is a generalization of the factorial function. Here is the corresponding result for the beta function:

If $j, k \in \mathbb{N}_+$ then

$$B(j, k) = \frac{(j-1)!(k-1)!}{(j+k-1)!} \quad (5.17.8)$$

Proof

Recall that $\Gamma(n) = (n-1)!$ for $n \in \mathbb{N}_+$, so this result follows from the previous one.

Let's generalize this result. First, recall from our study of combinatorial structures that for $a \in \mathbb{R}$ and $j \in \mathbb{N}$, the *ascending power* of base a and order j is

$$a^{[j]} = a(a+1) \cdots [a+(j-1)] \quad (5.17.9)$$

If $a, b \in (0, \infty)$, and $j, k \in \mathbb{N}$, then

$$\frac{B(a+j, b+k)}{B(a, b)} = \frac{a^{[j]} b^{[k]}}{(a+b)^{[j+k]}} \quad (5.17.10)$$

Proof

Recall that $\Gamma(a+j) = a^{[j]} \Gamma(a)$, so the result follows from the representation [above](#) for the beta function in terms of the gamma function.

$$B\left(\frac{1}{2}, \frac{1}{2}\right) = \pi.$$

Proof




Figure 5.17.1: The graph of $B(a, b)$ on the square $0 < a < 5$, $0 < b < 5$

The Incomplete Beta Function

The integral that defines the beta function can be generalized by changing the interval of integration from $(0, 1)$ to $(0, x)$ where $x \in [0, 1]$.

The *incomplete beta function* is defined as follows

$$B(x; a, b) = \int_0^x u^{a-1} (1-u)^{b-1} du, \quad x \in (0, 1); a, b \in (0, \infty) \quad (5.17.11)$$

Of course, the ordinary (complete) beta function is $B(a, b) = B(1; a, b)$ for $a, b \in (0, \infty)$.

The Standard Beta Distribution

Distribution Functions

The beta distributions are a family of continuous distributions on the interval $(0, 1)$.

The (standard) *beta distribution* with *left parameter* $a \in (0, \infty)$ and *right parameter* $b \in (0, \infty)$ has probability density function f given by

$$f(x) = \frac{1}{B(a, b)} x^{a-1} (1-x)^{b-1}, \quad x \in (0, 1) \quad (5.17.12)$$

Of course, the beta function is simply the normalizing constant, so it's clear that f is a valid probability density function. If $a \geq 1$, f is defined at 0, and if $b \geq 1$, f is defined at 1. In these cases, it's customary to extend the domain of f to these endpoints. The beta distribution is useful for modeling random probabilities and proportions, particularly in the context of Bayesian analysis. The distribution has just two parameters and yet a rich variety of shapes (so in particular, both parameters are *shape parameters*). Qualitatively, the first order properties of f depend on whether each parameter is less than, equal to, or greater than 1.

For $a, b \in (0, \infty)$ with $a + b \neq 2$, define

$$x_0 = \frac{a-1}{a+b-2} \quad (5.17.13)$$

1. If $0 < a < 1$ and $0 < b < 1$, f decreases and then increases with minimum value at x_0 and with $f(x) \rightarrow \infty$ as $x \downarrow 0$ and as $x \uparrow 1$.
2. If $a = 1$ and $b = 1$, f is constant.
3. If $0 < a < 1$ and $b \geq 1$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
4. If $a \geq 1$ and $0 < b < 1$, f is increasing with $f(x) \rightarrow \infty$ as $x \uparrow 1$.
5. If $a = 1$ and $b > 1$, f is decreasing with mode at $x = 0$.
6. If $a > 1$ and $b = 1$, f is increasing with mode at $x = 1$.
7. If $a > 1$ and $b > 1$, f increases and then decreases with mode at x_0 .

Proof

These results follow from standard calculus. The first derivative is

$$f'(x) = \frac{1}{B(a, b)} x^{a-2} (1-x)^{b-2} [(a-1) - (a+b-2)x], \quad 0 < x < 1 \quad (5.17.14)$$

From part (b), note that the special case $a = 1$ and $b = 1$ gives the continuous uniform distribution on the interval $(0, 1)$ (the *standard uniform distribution*). Note also that when $a < 1$ or $b < 1$, the probability density function is unbounded, and hence the distribution has no mode. On the other hand, if $a \geq 1$, $b \geq 1$, and one of the inequalities is strict, the distribution has a unique mode at x_0 . The second order properties are more complicated.

For $a, b \in (0, \infty)$ with $a + b \notin \{2, 3\}$ and $(a-1)(b-1)(a+b-3) \geq 0$, define

$$x_1 = \frac{(a-1)(a+b-3) - \sqrt{(a-1)(b-1)(a+b-3)}}{(a+b-3)(a+b-2)} \quad (5.17.15)$$

$$x_2 = \frac{(a-1)(a+b-3) + \sqrt{(a-1)(b-1)(a+b-3)}}{(a+b-3)(a+b-2)} \quad (5.17.16)$$

For $a < 1$ and $a + b = 2$ or for $b < 1$ and $a + b = 2$, define $x_1 = x_2 = 1 - a/2$.

1. If $a \leq 1$ and $b \leq 1$, or if $a \leq 1$ and $b \geq 2$, or if $a \geq 2$ and $b \leq 1$, f is concave upward.
2. If $a \leq 1$ and $1 < b < 2$, f is concave upward and then downward with inflection point at x_1 .
3. If $1 < a < 2$ and $b \leq 1$, f is concave downward and then upward with inflection point at x_2 .
4. If $1 < a \leq 2$ and $1 < b \leq 2$, f is concave downward.
5. If $1 < a \leq 2$ and $b > 2$, f is concave downward and then upward with inflection point at x_2 .

6. If $a > 2$ and $1 < b \leq 2$, f is concave upward and then downward with inflection point at x_1 .
7. If $a > 2$ and $b > 2$, f is concave upward, then downward, then upward again, with inflection points at x_1 and x_2 .

Proof

These results follow from standard (but very tedious) calculus. The second derivative is

$$f''(x) = \frac{1}{B(a, b)} x^{a-3} (1-x)^{b-3} [(a+b-2)(a+b-3)x^2 - 2(a-1)(a+b-3)x + (a-1)(a-2)] \quad (5.17.17)$$

In the special distribution simulator, select the beta distribution. Vary the parameters and note the shape of the beta density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the true density function.

The special case $a = \frac{1}{2}$, $b = \frac{1}{2}$ is the *arcsine distribution*, with probability density function given by

$$f(x) = \frac{1}{\pi \sqrt{x(1-x)}}, \quad x \in (0, 1) \quad (5.17.18)$$

This distribution is important in a number of applications, and so the arcsine distribution is studied in a separate section.

The beta distribution function F can be easily expressed in terms of the [incomplete beta function](#). As usual a denotes the left parameter and b the right parameter.

The beta distribution function F with parameters $a, b \in (0, \infty)$ is given by

$$F(x) = \frac{B(x; a, b)}{B(a, b)}, \quad x \in (0, 1) \quad (5.17.19)$$

The distribution function F is sometimes known as the *regularized incomplete beta function*. In some special cases, the distribution function F and its inverse, the quantile function F^{-1} , can be computed in closed form, without resorting to special functions.

If $a \in (0, \infty)$ and $b = 1$ then

1. $F(x) = x^a$ for $x \in (0, 1)$
2. $F^{-1}(p) = p^{1/a}$ for $p \in (0, 1)$

If $a = 1$ and $b \in (0, \infty)$ then

1. $F(x) = 1 - (1-x)^b$ for $x \in (0, 1)$
2. $F^{-1}(p) = 1 - (1-p)^{1/b}$ for $p \in (0, 1)$

If $a = b = \frac{1}{2}$ (the arcsine distribution) then

1. $F(x) = \frac{2}{\pi} \arcsin(\sqrt{x})$ for $x \in (0, 1)$
2. $F^{-1}(p) = \sin^2\left(\frac{\pi}{2}p\right)$ for $p \in (0, 1)$

There is an interesting relationship between the distribution functions of the beta distribution and the binomial distribution, when the beta parameters are positive integers. To state the relationship we need to embellish our notation to indicate the dependence on the parameters. Thus, let $F_{a,b}$ denote the beta distribution function with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$, and let $G_{n,p}$ denote the binomial distribution function with trial parameter $n \in \mathbb{N}_+$ and success parameter $p \in (0, 1)$.

If $j, k \in \mathbb{N}_+$ and $x \in (0, 1)$ then

$$F_{j,k}(x) = G_{j+k-1, 1-x}(k-1) \quad (5.17.20)$$

Proof

By definition

$$F_{j,k}(x) = \frac{1}{B(j,k)} \int_0^x t^{j-1} (1-t)^{k-1} dt \quad (5.17.21)$$

Integrate by parts with $u = (1-t)^{k-1}$ and $dv = t^{j-1} dt$, so that $du = -(k-1)(1-t)^{k-2}$ and $v = t^j/j$. The result is

$$F_{j,k}(x) = \frac{1}{jB(j,k)} (1-x)^{k-1} x^j + \frac{k-1}{jB(j,k)} \int_0^x t^j (1-t)^{k-2} dt \quad (5.17.22)$$

But by the [property of the beta function](#) above, $B(j,k) = (j-1)!(k-1)!/(j+k-1)!$. Hence $1/jB(j,k) = \binom{j+k-1}{k-1}$ and $(k-1)/jB(j,k) = 1/B(j+1, k-1)$. Thus, the last displayed equation can be rewritten as

$$F_{j,k}(x) = \binom{j+k-1}{k-1} (1-x)^{k-1} x^j + F_{j+1,k-1}(x) \quad (5.17.23)$$

Recall from the [special case](#) above that $F_{j+k-1,1}(x) = x^{j+k-1}$. Iterating the last displayed equation gives the result.

In the special distribution calculator, select the beta distribution. Vary the parameters and note the shape of the density function and the distribution function. In each of the following cases, find the median, the first and third quartiles, and the interquartile range. Sketch the boxplot.

1. $a = 1, b = 1$
2. $a = 1, b = 3$
3. $a = 3, b = 1$
4. $a = 2, b = 4$
5. $a = 4, b = 2$
6. $a = 4, b = 4$

Moments

The moments of the beta distribution are easy to express in terms of the beta function. As before, suppose that X has the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$.

If $k \in [0, \infty)$ then

$$\mathbb{E}(X^k) = \frac{B(a+k, b)}{B(a, b)} \quad (5.17.24)$$

In particular, if $k \in \mathbb{N}$ then

$$\mathbb{E}(X^k) = \frac{a^{[k]}}{(a+b)^{[k]}} \quad (5.17.25)$$

Proof

Note that

$$\mathbb{E}(X^k) = \int_0^1 x^k \frac{1}{B(a, b)} x^{a-1} (1-x)^{b-1} dx = \frac{1}{B(a, b)} \int_0^1 x^{a+k-1} (1-x)^{b-1} dx = \frac{B(a+k, b)}{B(a, b)} \quad (5.17.26)$$

If $k \in \mathbb{N}$, the formula simplifies by the [property of the beta function](#) above.

From the general formula for the moments, it's straightforward to compute the mean, variance, skewness, and kurtosis.

The mean and variance of X are

$$\mathbb{E}(X) = \frac{a}{a+b} \quad (5.17.27)$$

$$\text{var}(X) = \frac{ab}{(a+b)^2(a+b+1)} \quad (5.17.28)$$

Proof

The formula for the mean and variance follow from the [formula for the moments](#) and the computational formula $\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2$

Note that the variance depends on the parameters a and b only through the product ab and the sum $a + b$.

Open the special distribution simulator and select the beta distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

$$\text{skew}(X) = \frac{2(b-a)\sqrt{a+b+1}}{(a+b+2)\sqrt{ab}} \quad (5.17.29)$$

$$\text{kurt}(X) = \frac{3a^3b + 3ab^3 + 6a^2b^2 + a^3 + b^3 + 13a^2b + 13ab^2 + a^2 + b^2 + 14ab}{ab(a+b+2)(a+b+3)} \quad (5.17.30)$$

Proof

These results follow from the computational formulas that give the skewness and kurtosis in terms of $\mathbb{E}(X^k)$ for $k \in \{1, 2, 3, 4\}$ and the [formula for the moments](#) above.

In particular, note that the distribution is positively skewed if $a < b$, unskewed if $a = b$ (the distribution is symmetric about $x = \frac{1}{2}$ in this case) and negatively skewed if $a > b$.

Open the special distribution simulator and select the beta distribution. Vary the parameters and note the shape of the probability density function in light of the previous result on skewness. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the true probability density function.

Related Distributions

The beta distribution is related to a number of other special distributions.

If X has the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$ then $Y = 1 - X$ has the beta distribution with left parameter b and right parameter a .

Proof

This follows from the standard change of variables formula. If f and g denote the PDFs of X and Y respectively, then

$$g(y) = f(1-y) = \frac{1}{B(a,b)}(1-y)^{a-1}y^{b-1} = \frac{1}{B(b,a)}y^{b-1}(1-y)^{a-1}, \quad y \in (0,1) \quad (5.17.31)$$

The beta distribution with right parameter 1 has a reciprocal relationship with the Pareto distribution.

Suppose that $a \in (0, \infty)$.

1. If X has the beta distribution with left parameter a and right parameter 1 then $Y = 1/X$ has the Pareto distribution with shape parameter a .
2. If Y has the Pareto distribution with shape parameter a then $X = 1/Y$ has the beta distribution with left parameter a and right parameter 1.

Proof

The two results are equivalent. In (a), suppose that X has the beta distribution with parameters a and 1. The transformation $y = 1/x$ maps $(0,1)$ one-to-one onto $(0,\infty)$. The inverse is $x = 1/y$ with $dx/dy = -1/y^2$. Recall also that $B(a,1) = 1/a$. By the change of variables formula, the PDF g of $Y = 1/X$ is given by

$$g(y) = f\left(\frac{1}{y}\right) \frac{1}{y^2} = a\left(\frac{1}{y}\right)^{a-1} \frac{1}{y^2} = \frac{a}{y^{a+1}}, \quad y \in (0,\infty) \quad (5.17.32)$$

We recognize g as the PDF of the Pareto distribution with shape parameter a .

The following result gives a connection between the beta distribution and the gamma distribution.

Suppose that X has the gamma distribution with shape parameter $a \in (0, \infty)$ and rate parameter $r \in (0, \infty)$, Y has the gamma distribution with shape parameter $b \in (0, \infty)$ and rate parameter r , and that X and Y are independent. Then $V = X/(X + Y)$ has the beta distribution with left parameter a and right parameter b .

Proof

Let $U = X + Y$ and $V = X/(X + Y)$. We will actually prove stronger results: U and V are independent, U has the gamma distribution with shape parameter $a + b$ and rate parameter r , and V has the beta distribution with parameters a and b . First note that (X, Y) has joint PDF f given by

$$f(x, y) = \frac{r^a}{\Gamma(a)} x^{a-1} e^{-rx} \frac{r^b}{\Gamma(b)} y^{b-1} e^{-ry} = \frac{r^{a+b}}{\Gamma(a)\Gamma(b)} x^{a-1} y^{b-1} e^{-r(x+y)}; \quad x, y \in (0, \infty) \quad (5.17.33)$$

The transformation $u = x + y$ and $v = x/(x + y)$ maps $(0, \infty) \times (0, \infty)$ one-to-one onto $(0, \infty) \times (0, 1)$. The inverse is $x = uv$, $y = u(1 - v)$ and the absolute value of the Jacobian is

$$\left| \det \frac{\partial(x, y)}{\partial(u, v)} \right| = u \quad (5.17.34)$$

Hence by the multivariate change of variables theorem, the PDF g of (U, V) is given by

$$g(u, v) = f[uv, u(1 - v)]u = \frac{r^{a+b}}{\Gamma(a)\Gamma(b)} (uv)^{a-1} [u(1 - v)]^{b-1} e^{-ru} u \quad (5.17.35)$$

$$= \frac{r^{a+b}}{\Gamma(a)\Gamma(b)} u^{a+b-1} e^{-ru} v^{a-1} (1 - v)^{b-1} \quad (5.17.36)$$

$$= \frac{r^{a+b}}{\Gamma(a+b)} u^{a+b-1} e^{-ru} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} v^{a-1} (1 - v)^{b-1}; \quad u \in (0, \infty), v \in (0, 1) \quad (5.17.37)$$

The results now follow from the factorization theorem. The factor in u is the gamma PDF with shape parameter $a + b$ and rate parameter r while the factor in v is the beta PDF with parameters a and b .

The following result gives a connection between the beta distribution and the F distribution. This connection is a minor variation of the previous result.

If X has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator then

$$Y = \frac{(n/d)X}{1 + (n/d)X} \quad (5.17.38)$$

has the beta distribution with left parameter $a = n/2$ and right parameter $b = d/2$.

Proof

If X has the F distribution with $n > 0$ degrees of freedom in the numerator and $d > 0$ degrees of freedom in the denominator then X can be written as

$$X = \frac{U/n}{V/d} \quad (5.17.39)$$

where U has the chi-square distribution with n degrees of freedom, V has the chi-square distribution with d degrees of freedom, and U and V are independent. Hence

$$Y = \frac{(n/d)X}{1 + (n/d)X} = \frac{U/V}{1 + U/V} = \frac{U}{U + V} \quad (5.17.40)$$

But the chi-square distribution is a special case of the gamma distribution. Specifically, U has the gamma distribution with shape parameter $n/2$ and rate parameter $1/2$, V has the gamma distribution with shape parameter $d/2$ and rate parameter $1/2$ and

again U and V are independent. Hence by the [previous result](#), Y has the beta distribution with left parameter $n/2$ and right parameter $d/2$.

Our next result is that the beta distribution is a member of the general exponential family of distributions.

Suppose that X has the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. Then the distribution is a two-parameter exponential family with natural parameters $a - 1$ and $b - 1$, and natural statistics $\ln(X)$ and $\ln(1 - X)$.

Proof

This follows from the definition of the general exponential distribution, since the PDF f of X can be written as

$$f(x) = \frac{1}{B(a, b)} \exp[(a - 1) \ln(x) + (b - 1) \ln(1 - x)], \quad x \in (0, 1) \quad (5.17.41)$$

The beta distribution is also the distribution of the order statistics of a random sample from the standard uniform distribution.

Suppose $n \in \mathbb{N}_+$ and that (X_1, X_2, \dots, X_n) is a sequence of independent variables, each with the standard uniform distribution. For $k \in \{1, 2, \dots, n\}$, the k th order statistics $X_{(k)}$ has the beta distribution with left parameter $a = k$ and right parameter $b = n - k + 1$.

Proof

See the section on order statistics.

One of the most important properties of the beta distribution, and one of the main reasons for its wide use in statistics, is that it forms a conjugate family for the success probability in the binomial and negative binomial distributions.

Suppose that P is a random probability having the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. Suppose also that X is a random variable such that the conditional distribution of X given $P = p \in (0, 1)$ is binomial with trial parameter $n \in \mathbb{N}_+$ and success parameter p . Then the conditional distribution of P given $X = k$ is beta with left parameter $a + k$ and right parameter $b + n - k$.

Proof

The joint PDF f of (P, X) on $(0, 1) \times \{0, 1, \dots, n\}$ is given by

$$f(p, k) = \frac{1}{B(a, b)} p^{a-1} (1-p)^{b-1} \binom{n}{k} p^k (1-p)^{n-k} = \frac{1}{B(a, b)} \binom{n}{k} p^{a+k-1} (1-p)^{b+n-k-1} \quad (5.17.42)$$

The conditional PDF of P given $X = k$ is simply the normalized version of the function $p \mapsto f(p, k)$. We can tell from the functional form that this distribution is beta with the parameters given in the theorem.

Suppose again that P is a random probability having the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. Suppose also that N is a random variable such that the conditional distribution of N given $P = p \in (0, 1)$ is negative binomial with stopping parameter $k \in \mathbb{N}_+$ and success parameter p . Then the conditional distribution of P given $N = n$ is beta with left parameter $a + k$ and right parameter $b + n - k$.

Proof

The joint PDF f of (P, N) on $(0, 1) \times \{k, k + 1, \dots\}$ is given by

$$f(p, n) = \frac{1}{B(a, b)} p^{a-1} (1-p)^{b-1} \binom{n-1}{k-1} p^k (1-p)^{n-k} = \frac{1}{B(a, b)} \binom{n-1}{k-1} p^{a+k-1} (1-p)^{b+n-k-1} \quad (5.17.43)$$

The conditional PDF of P given $N = n$ is simply the normalized version of the function $p \mapsto f(p, n)$. We can tell from the functional form that this distribution is beta with the parameters given in the theorem.

In both cases, note that in the posterior distribution of P , the left parameter is increased by the number of successes and the right parameter by the number of failures. For more on this, see the section on Bayesian estimation in the chapter on point estimation.

The General Beta Distribution

The beta distribution can be easily generalized from the support interval $(0, 1)$ to an arbitrary bounded interval using a linear transformation. Thus, this generalization is simply the location-scale family associated with the standard beta distribution.

Suppose that Z has the standard beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. For $c \in \mathbb{R}$ and $d \in (0, \infty)$ random variable $X = c + dZ$ has the beta distribution with left parameter a , right parameter b , location parameter c and scale parameter d .

For the remainder of this discussion, suppose that X has the distribution in the definition above.

X has probability density function

$$f(x) = \frac{1}{B(a, b)d^{a+b-1}}(x - c)^{a-1}(c + d - x)^{b-1}, \quad x \in (c, c + d) \quad (5.17.44)$$

Proof

This follows from a standard result for location-scale families. If g denotes the [standard beta PDF](#) of Z , then X has PDF f given by

$$f(x) = \frac{1}{d}g\left(\frac{x - c}{d}\right), \quad x \in (c, c + d) \quad (5.17.45)$$

Most of the results in the previous sections have simple extensions to the general beta distribution.

The mean and variance of X are

1. $\mathbb{E}(X) = c + d\frac{a}{a+b}$
2. $\text{var}(X) = d^2 \frac{ab}{(a+b)^2(a+b+1)}$

Proof

This follows from the [standard mean and variance](#) and basic properties of expected value and variance.

1. $\mathbb{E}(X) = c + d\mathbb{E}(Z)$
2. $\text{var}(X) = d^2\text{var}(Z)$

Recall that skewness and variance are defined in terms of standard scores, and hence are unchanged under location-scale transformations. Hence the skewness and kurtosis of X are just as for the [standard beta distribution](#).

This page titled [5.17: The Beta Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.18: The Beta Prime Distribution

Basic Theory

The beta prime distribution is the distribution of the odds ratio associated with a random variable with the [beta distribution](#). Since variables with beta distributions are often used to model random probabilities and proportions, the corresponding odds ratios occur naturally as well.

Definition

Suppose that U has the beta distribution with shape parameters $a, b \in (0, \infty)$. Random variable $X = U/(1 - U)$ has the *beta prime distribution* with shape parameters a and b .

The special case $a = b = 1$ is known as the *standard beta prime distribution*. Since U has a continuous distribution on the interval $(0, 1)$, random variable X has a continuous distribution on the interval $(0, \infty)$.

Distribution Functions

Suppose that X has the beta prime distribution with shape parameters $a, b \in (0, \infty)$, and as usual, let B denote the beta function.

X has probability density function f given by

$$f(x) = \frac{1}{B(a, b)} \frac{x^{a-1}}{(1+x)^{a+b}}, \quad x \in (0, \infty) \quad (5.18.1)$$

Proof

First, recall that the beta PDF g with parameters a and b is

$$g(u) = u^{a-1}(1-u)^{b-1}, \quad u \in (0, 1) \quad (5.18.2)$$

The transformation $x = u/(1-u)$ maps $(0, 1)$ onto $(0, \infty)$ and is increasing. The inverse transformation is $u = x/(x+1)$, and $1-u = 1/(x+1)$ and $du/dx = 1/(x+1)^2$. Thus, by the change of variables formula,

$$f(x) = g(u) \frac{du}{dx} = \frac{1}{B(a, b)} \left(\frac{x}{x+1} \right)^{a-1} \left(\frac{1}{x+1} \right)^{b-1} \frac{1}{(x+1)^2} = \frac{1}{B(a, b)} \frac{x^{a-1}}{(x+1)^{a+b}}, \quad x \in (0, \infty) \quad (5.18.3)$$

If $a \geq 1$, the probability density function is defined at $x = 0$, so in this case, it's customary add this endpoint to the domain. In particular, for the standard beta prime distribution,

$$f(x) = \frac{1}{(1+x)^2}, \quad x \in [0, \infty) \quad (5.18.4)$$

Qualitatively, the first order properties of the probability density function f depend only on a , and in particular on whether a is less than, equal to, or greater than 1.

The probability density function f satisfies the following properties:

1. If $0 < a < 1$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
2. If $a = 1$, f is decreasing with mode at $x = 0$.
3. If $a > 1$, f increases and then decreases with mode at $x = (a-1)/(b+1)$.

Proof

These properties follow from standard calculus. The first derivative of f is

$$f'(x) = \frac{1}{B(a, b)} \frac{x^{a-2}}{(1+x)^{a+b+1}} [(a-1) - x(b+1)], \quad x \in (0, \infty) \quad (5.18.5)$$

Qualitatively, the second order properties of f also depend only on a , with transitions at $a = 1$ and $a = 2$.

For $a > 1$, define

$$x_1 = \frac{(a-1)(b+2) - \sqrt{(a-1)(b+2)(a+b)}}{(b+1)(b+2)} \quad (5.18.6)$$

$$x_2 = \frac{(a-1)(b+2) + \sqrt{(a-1)(b+2)(a+b)}}{(b+1)(b+2)} \quad (5.18.7)$$

The probability density function f satisfies the following properties:

1. If $0 < a \leq 1$, f is concave upward.
2. If $1 < a \leq 2$, f is concave downward and then upward, with inflection point at x_2 .
3. If $a > 2$, f is concave upward, then downward, then upward again, with inflection points at x_1 and x_2 .

Proof

These results follow from standard calculus. The second derivative of f is

$$f''(x) = \frac{1}{B(a, b)} \frac{x^{a-3}}{(1+x)^{a+b+2}} [(a-1)(a-2) - 2(a-1)(b+2)x + (b+1)(b+2)x^2], \quad x \in (0, \infty) \quad (5.18.8)$$

Open the Special Distribution Simulator and select the beta prime distribution. Vary the parameters and note the shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Because of the definition of the beta prime variable, the distribution function of X has a simple expression in terms of the beta distribution function with the same parameters, which in turn is the regularized incomplete beta function. So let G denote the distribution function of the beta distribution with parameters $a, b \in (0, \infty)$, and recall that

$$G(x) = \frac{B(x; a, b)}{B(a, b)}, \quad x \in (0, 1) \quad (5.18.9)$$

X has distribution function F given by

$$F(x) = G\left(\frac{x}{x+1}\right), \quad x \in [0, \infty) \quad (5.18.10)$$

Proof

As noted in the proof of the [formula for the PDF](#), $x = u/(1-u)$ is strictly increasing with inverse $u = x/(x+1)$. Hence

$$F(x) = \mathbb{P}(X \leq x) = \mathbb{P}\left(\frac{U}{U-1} \leq x\right) = \mathbb{P}\left(U \leq \frac{x}{x+1}\right) = G\left(\frac{x}{x+1}\right), \quad x \in [0, \infty) \quad (5.18.11)$$

Similarly, the quantile function of X has a simple expression in terms of the beta quantile function G^{-1} with the same parameters.

X has quantile function F^{-1} given by

$$F^{-1}(p) = \frac{G^{-1}(p)}{1 - G^{-1}(p)}, \quad p \in [0, 1) \quad (5.18.12)$$

Proof

This follows from the [result for the CDF](#) by solving $p = F(x) = G\left(\frac{x}{x+1}\right)$ for x in terms of p .

Open the Special Distribution Calculator and choose the beta prime distribution. Vary the parameters and note the shape of the distribution function. For selected values of the parameters, find the median and the first and third quartiles.

For certain values of the parameters, the distribution and quantile functions have simple, closed form expressions.

If $a \in (0, \infty)$ and $b = 1$ then

1. $F(x) = \left(\frac{x}{x+1}\right)^a$ for $x \in [0, \infty)$
2. $F^{-1}(p) = \frac{p^{1/a}}{1-p^{1/a}}$ for $p \in [0, 1)$

Proof

For $a > 0$ and $b = 1$, $G(u) = u^a$ for $u \in [0, 1]$ and $G^{-1}(p) = p^{1/a}$ for $p \in [0, 1]$

If $a = 1$ and $b \in (0, \infty)$ then

1. $F(x) = 1 - \left(\frac{1}{x+1}\right)^b$ for $x \in [0, \infty)$
2. $F^{-1}(p) = \frac{1 - (1-p)^{1/b}}{(1-p)^{1/b}}$ for $p \in [0, 1]$

Proof

For $a = 1$ and $b > 0$, $G(u) = 1 - (1-u)^b$ for $u \in [0, 1]$ and $G^{-1}(p) = 1 - (1-p)^{1/b}$ for $p \in [0, 1]$.

If $a = b = \frac{1}{2}$ then

1. $F(x) = \frac{2}{\pi} \arcsin\left(\sqrt{\frac{x}{x+1}}\right)$ for $x \in [0, \infty)$
2. $F^{-1}(p) = \frac{\sin^2\left(\frac{\pi}{2}p\right)}{1 - \sin^2\left(\frac{\pi}{2}p\right)}$ for $p \in [0, 1]$

Proof

For $a = b = \frac{1}{2}$, $G(u) = \frac{2}{\pi} \arcsin(\sqrt{u})$ for $u \in (0, 1)$ and $G^{-1}(p) = \sin^2\left(\frac{\pi}{2}p\right)$ for $p \in [0, 1]$

When $a = b = \frac{1}{2}$, X is the odds ratio for a variable with the standard arcsine distribution.

Moments

As before, X denotes a random variable with the beta prime distribution, with parameters $a, b \in (0, \infty)$. The moments of X have a simple expression in terms of the beta function.

If $t \in (-a, b)$ then

$$\mathbb{E}(X^t) = \frac{B(a+t, b-t)}{B(a, b)} \quad (5.18.13)$$

If $t \in (-\infty, -a] \cup [b, \infty)$ then $\mathbb{E}(X^t) = \infty$.

Proof

Once again, let g denote the beta PDF with parameters a and b . With the transformation $x = u/(1-u)$, as in the proof [PDF formula](#), we have $f(x)dx = g(u)du$. Hence

$$\int_0^\infty x^t f(x)dx = \int_0^1 \left(\frac{u}{1-u}\right)^t g(u)du = \frac{1}{B(a, b)} \int_0^1 u^{a+t-1} (1-u)^{b-t-1} du \quad (5.18.14)$$

If $t \leq -a$ the improper integral diverges to ∞ at 0. If $t \geq b$ the improper integral diverges to ∞ at 1. If $-a < t < b$ the integral is $B(a+t, b-t)$ by definition of the beta function.

Of course, we are usually most interested in the integer moments of X . Recall that for $x \in \mathbb{R}$ and $n \in \mathbb{N}$, the *rising power* of x of order n is $x^{[n]} = x(x+1) \cdots (x+n-1)$.

Suppose that $n \in \mathbb{N}$. If $n < b$ Then

$$\mathbb{E}(X^n) = \prod_{k=1}^n \frac{a+k-1}{b-k} \quad (5.18.15)$$

If $n \geq b$ then $\mathbb{E}(X^n) = \infty$.

Proof

From the [general moment result](#),

$$\mathbb{E}(X^n) = \frac{B(a+n, b-n)}{B(a, b)} = \frac{\Gamma(a+n)\Gamma(a-n)}{\Gamma(a+b)} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} = \frac{\Gamma(a+n)}{\Gamma(a)} \frac{\Gamma(b-n)}{\Gamma(b)} = \frac{a^{[n]}}{(b-n)^{[n]}} \quad (5.18.16)$$

by a basic property of the gamma function.

As a corollary, we have the mean and variance.

If $b > 1$ then

$$\mathbb{E}(X) = \frac{a}{b-1} \quad (5.18.17)$$

If $b > 2$ then

$$\text{var}(X) = \frac{a(a+b-1)}{(b-1)^2(b-2)} \quad (5.18.18)$$

Proof

This follows from the [general moment result](#) above and the computational formula $\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2$.

Open the Special Distribution Simulator and select the beta prime distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Finally, the general moment result leads to the skewness and kurtosis of X .

If $b > 3$ then

$$\text{skew}(X) = \frac{2(2a+b-1)}{b-3} \sqrt{\frac{b-2}{a(a+b-1)}} \quad (5.18.19)$$

Proof

This follows from the usual computational formula for skewness in terms of the moments $\mathbb{E}(X^n)$ for $n \in \{1, 2, 3\}$ and the [general moment result](#) above.

In particular, the distribution is positively skewed for all $a > 0$ and $b > 3$.

If $b > 4$ then

$$\begin{aligned} \text{kurt}(X) &= \frac{3a^3b^2 + 69a^3b - 30a^3 + 6a^2b^3 + 12a^2b^2 - 78a^2b + 60a^2 + 3ab^4 + 9ab^3 - 69ab^2 + 99ab - 42a + 6b^4 - 30b^3}{(a+b-1)(b-3)(b-4)} \end{aligned} \quad (5.18.20)$$

Proof

This follows from the usual computational formula for kurtosis in terms of the moments $\mathbb{E}(X^n)$ for $n \in \{1, 2, 3, 4\}$ and the [general moment result](#) above.

Related Distributions

The most important connection is the one between the beta prime distribution and the beta distribution given in the [definition](#). We repeat this for emphasis.

Suppose that $a, b \in (0, \infty)$.

1. If U has the beta distribution with parameters a and b , then $X = U/(1-U)$ has the beta prime distribution with parameters a and b .
2. If X has the beta prime distribution with parameters a and b , then $U = X/(X+1)$ has the beta distribution with parameters a and b .

The beta prime family is closed under the reciprocal transformation.

If X has the beta prime distribution with parameters $a, b \in (0, \infty)$ then $1/X$ has the beta prime distribution with parameters b and a .

Proof

A direct proof using the change of variables formula is possible, of course, but a better proof uses a corresponding property of the beta distribution. By definition, we can take $X = U/(1-U)$ where U has the beta distribution with parameters a and b . But then $1/X = (1-U)/U$, and $1-U$ has the beta distribution with parameters b and a . By another application of the definition, $1/X$ has the beta prime distribution with parameters b and a .

The beta prime distribution is closely related to the F distribution by a simple scale transformation.

Connections with the F distributions.

1. If X has the beta prime distribution with parameters $a, b \in (0, \infty)$ then $Y = \frac{b}{a}X$ has the F distribution with $2a$ degrees of freedom in the numerator and $2b$ degrees of freedom in the denominator.
2. If Y has the F distribution with $n \in (0, \infty)$ degrees of freedom in the numerator and $d \in (0, \infty)$ degrees of freedom in the denominator, then $X = \frac{n}{d}Y$ has the beta prime distribution with parameters $n/2$ and $d/2$.

Proof

Let f denote the PDF of X and g the PDF of Y .

1. By the change of variables formula,

$$g(y) = \frac{a}{b} f\left(\frac{a}{b}y\right), \quad x \in (0, \infty) \quad (5.18.21)$$

Substituting into the beta prime PDF shows that Y has the appropriate F distribution.

2. Again using the change of variables formula,

$$f(x) = \frac{d}{n} g\left(\frac{d}{n}x\right), \quad x \in (0, \infty) \quad (5.18.22)$$

Substituting into the F PDF shows that X has the appropriate beta prime PDF.

The beta prime is the distribution of the ratio of independent variables with standard gamma distributions. (Recall that *standard* here means that the scale parameter is 1.)

Suppose that Y and Z are independent and have standard gamma distributions with shape parameters $a \in (0, \infty)$ and $b \in (0, \infty)$, respectively. Then $X = Y/Z$ has the beta prime distribution with parameters a and b .

Proof

Of course, a direct proof can be constructed, but a better approach is to use the previous result. Thus suppose that Y and Z are as stated in the theorem. Then $2Y$ and $2Z$ are independent chi-square variables with $2a$ and $2b$ degrees of freedom, respectively. Hence

$$W = \frac{Y/2a}{Z/2b} \quad (5.18.23)$$

has the F distribution with $2a$ degrees of freedom in the numerator and $2b$ degrees of freedom in the denominator. By the previous result,

$$X = \frac{2a}{2b} W = \frac{Y}{Z} \quad (5.18.24)$$

has the beta prime distribution with parameters a and b .

The standard beta prime distribution is the same as the standard log-logistic distribution.

Proof

The PDF of the standard beta prime distribution is $f(x) = 1/(1+x)^2$ for $x \in [0, \infty)$, which is the same as the PDF of the standard log-logistic distribution.

Finally, the beta prime distribution is a member of the general exponential family of distributions.

Suppose that X has the beta prime distribution with parameters $a, b \in (0, \infty)$. Then X has a two-parameter general exponential distribution with natural parameters $a-1$ and $-(a+b)$ and natural statistics $\ln(X)$ and $\ln(1+X)$.

Proof

This follows from the definition of the general exponential family, since the PDF can be written in the form

$$f(x) = \frac{1}{B(a, b)} \exp[(a-1)\ln(x) - (a+b)\ln(1+x)], \quad x \in (0, \infty) \quad (5.18.25)$$

This page titled [5.18: The Beta Prime Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.19: The Arcsine Distribution

The *arcsine distribution* is important in the study of Brownian motion and prime numbers, among other applications.

The Standard Arcsine Distribution

Distribution Functions

The *standard arcsine distribution* is a continuous distribution on the interval $(0, 1)$ with probability density function g given by

$$g(x) = \frac{1}{\pi\sqrt{x(1-x)}}, \quad x \in (0, 1) \quad (5.19.1)$$

Proof

There are a couple of ways to see that g is a valid PDF. First, it's the beta PDF with parameters $a = b = \frac{1}{2}$:

$$g(x) = \frac{1}{B(1/2, 1/2)} x^{-1/2} (1-x)^{-1/2}, \quad x \in (0, 1) \quad (5.19.2)$$

since we recall that $B(\frac{1}{2}, \frac{1}{2}) = \pi$. A direct proof is also easy: The substitution $u = \sqrt{x}$, $x = u^2$, $dx = 2u du$ gives

$$\int_0^1 \frac{1}{\pi\sqrt{x(1-x)}} dx = \int_0^1 \frac{2}{\pi\sqrt{1-u^2}} du = \frac{2}{\pi} \arcsin u \Big|_0^1 = \frac{2}{\pi} \left(\frac{\pi}{2} - 0 \right) = 1 \quad (5.19.3)$$

The occurrence of the arcsine function in the proof that g is a probability density function explains the name.

The standard arcsine probability density function g satisfies the following properties:

1. g is symmetric about $x = \frac{1}{2}$.
2. g decreases and then increases with minimum value at $x = \frac{1}{2}$.
3. g is concave upward
4. $g(x) \rightarrow \infty$ as $x \downarrow 0$ and as $x \uparrow 1$.

Proof

1. Note that g is a function of x only through $x(1-x)$.
2. This follows from standard calculus:

$$g'(x) = \frac{2x-1}{2\pi[x(1-x)]^{3/2}} \quad (5.19.4)$$

3. This also follows from standard calculus:

$$g''(x) = \frac{3-8x+8x^2}{4\pi[x(1-x)]^{5/2}} \quad (5.19.5)$$

4. The limits are clear.

In particular, the standard arcsine distribution is U-shaped and has no mode.

Open the Special Distribution Simulator and select the arcsine distribution. Keep the default parameter values and note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function has a simple expression in terms of the arcsine function, again justifying the name of the distribution.

The standard arcsine distribution function G is given by $G(x) = \frac{2}{\pi} \arcsin(\sqrt{x})$ for $x \in [0, 1]$.

Proof

Again, using the substitution $u = \sqrt{t}$, $t = u^2$, $dt = 2u du$:

$$G(x) = \int_0^x \frac{1}{\pi \sqrt{t(1-t)}} dt = \int_0^{\sqrt{x}} \frac{2}{\pi \sqrt{1-u^2}} du = \frac{2}{\pi} \arcsin(t) \Big|_0^{\sqrt{x}} = \frac{2}{\pi} \arcsin(\sqrt{x}) \quad (5.19.6)$$

Not surprisingly, the quantile function has a simple expression in terms of the sine function.

The standard arcsine quantile function G^{-1} is given by $G^{-1}(p) = \sin^2\left(\frac{\pi}{2}p\right)$ for $p \in [0, 1]$. In particular, the quartiles are

1. $q_1 = \sin^2\left(\frac{\pi}{8}\right) = \frac{1}{4}(2 - \sqrt{2}) \approx 0.1464$, the first quartile
2. $q_2 = \frac{1}{2}$, the median
3. $q_3 = \sin^2\left(\frac{3\pi}{8}\right) = \frac{1}{4}(2 + \sqrt{2}) \approx 0.8536$, the third quartile

Proof

The formula for the quantile function follows from the distribution function by solving $p = G(x)$ for x in terms of $p \in [0, 1]$.

Open the Special Distribution Calculator and select the arcsine distribution. Keep the default parameter values and note the shape of the distribution function. Compute selected values of the distribution function and the quantile function.

Moments

Suppose that random variable Z has the standard arcsine distribution. First we give the mean and variance.

The mean and variance of Z are

1. $\mathbb{E}(Z) = \frac{1}{2}$
2. $\text{var}(Z) = \frac{1}{8}$

Proof

1. The mean is $\frac{1}{2}$ by symmetry.
2. Using the usual substitution $u = \sqrt{x}$, $x = u^2$, $dx = 2u du$ and then the substitution $u = \sin \theta$, $du = \cos \theta d\theta$ gives

$$\mathbb{E}(Z^2) = \int_0^1 \frac{1}{\pi \sqrt{x(1-x)}} dx = \int_0^1 \frac{2u^4}{\pi \sqrt{1-u^2}} = \int_0^{\pi/2} \frac{2}{\pi} \sin^4(\theta) d\theta = \frac{2}{\pi} \frac{3\pi}{16} = \frac{3}{8} \quad (5.19.7)$$

Open the Special Distribution Simulator and select the arcsine distribution. Keep the default parameter values. Run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The general moments about 0 can be expressed as products.

For $n \in \mathbb{N}$,

$$\mathbb{E}(Z^n) = \prod_{j=0}^{n-1} \frac{2j+1}{2j+2} \quad (5.19.8)$$

Proof

The same integral substitutions as [before](#) gives

$$\mathbb{E}(Z^n) = \int_0^{\pi/2} \frac{2}{\pi} \sin^{2n}(\theta) d\theta = \prod_{j=0}^{n-1} \frac{2j+1}{2j+2} \quad (5.19.9)$$

Of course, the moments can be used to give a formula for the moment generating function, but this formula is not particularly helpful since it is not in closed form.

Z has moment generating function m given by

$$m(t) = \mathbb{E}(e^{tZ}) = \sum_{n=0}^{\infty} \left(\prod_{j=0}^{n-1} \frac{2j+1}{2j+2} \right) \frac{t^n}{n!}, \quad t \in \mathbb{R} \quad (5.19.10)$$

Finally we give the skewness and kurtosis.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = \frac{3}{2}$

Proof

1. The skewness is 0 by the symmetry of the distribution.
2. The result for the kurtosis follows from the standard formula for kurtosis in terms of the [moments](#): $\mathbb{E}(Z) = \frac{1}{2}$, $\mathbb{E}(Z^2) = \frac{3}{8}$, $\mathbb{E}(Z^3) = \frac{5}{16}$, and $\mathbb{E}(Z^4) = \frac{35}{128}$.

Related Distributions

As noted earlier, the standard arcsine distribution is a special case of the beta distribution.

The standard arcsine distribution is the beta distribution with left parameter $\frac{1}{2}$ and right parameter $\frac{1}{2}$.

Proof

The beta distribution with parameters $a = b = \frac{1}{2}$ has PDF

$$x \mapsto \frac{1}{B(1/2, 1/2)} x^{-1/2} (1-x)^{-1/2}, \quad x \in (0, 1) \quad (5.19.11)$$

But $B(\frac{1}{2}, \frac{1}{2}) = \pi$, so this is the standard arcsine PDF.

Since the [quantile function](#) is in closed form, the standard arcsine distribution can be simulated by the random quantile method.

Connections with the standard uniform distribution.

1. If U has the standard uniform distribution (a random number) then $X = \sin^2(\frac{\pi}{2}U)$ has the standard arcsine distribution.
2. If X has the standard arcsine distribution then $U = \frac{2}{\pi} \arcsin(\sqrt{X})$ has the standard uniform distribution.

Open the random quantile simulator and select the arcsine distribution. Keep the default parameters. Run the experiment 1000 times and compare the empirical probability density function, mean, and standard deviation to their distributional counterparts. Note how the random quantiles simulate the distribution.

The following exercise illustrates the connection between the Brownian motion process and the standard arcsine distribution.

Open the Brownian motion simulator. Keep the default time parameter and select the last zero random variable. Note that this random variable has the standard arcsine distribution. Run the experiment 1000 time and compare the empirical probability density function, mean, and standard deviation to their distributional counterparts. Note how the last zero simulates the distribution.

The General Arcsine Distribution

The standard arcsine distribution is generalized by adding location and scale parameters.

Definition

If Z has the standard arcsine distribution, and if $a \in \mathbb{R}$ and $b \in (0, \infty)$, then $X = a + bZ$ has the *arcsine distribution* with location parameter a and scale parameter b .

So X has a continuous distribution on the interval $(a, a + b)$.

Distribution Functions

Suppose that X has the arcsine distribution with location parameter $a \in \mathbb{R}$ and scale parameter $w \in (0, \infty)$.

X has probability density function f given by

$$f(x) = \frac{1}{\pi \sqrt{(x-a)(a+w-x)}}, \quad x \in (a, a+w) \quad (5.19.12)$$

1. f is symmetric about $a + \frac{1}{2}w$.
2. f decreases and then increases with minimum value at $x = a + \frac{1}{2}w$.
3. f is concave upward.
4. $f(x) \rightarrow \infty$ as $x \downarrow a$ and as $x \uparrow a+w$.

Proof

Recall that $f(x) = \frac{1}{w} g\left(\frac{x-a}{w}\right)$ where g is the [PDF of the standard arcsine distribution](#).

An alternate parameterization of the general arcsine distribution is by the endpoints of the support interval: the left endpoint (location parameter) a and the right endpoint $b = a + w$.

Open the Special Distribution Simulator and select the arcsine distribution. Vary the location and scale parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Once again, the distribution function has a simple representation in terms of the arcsine function.

X has distribution function F given by

$$F(x) = \frac{2}{\pi} \arcsin\left(\sqrt{\frac{x-a}{w}}\right), \quad x \in [a, a+w] \quad (5.19.13)$$

Proof

Recall that $F(x) = G[(x-a)/w]$ where G is the [CDF of the standard arcsine distribution](#).

As before, the quantile function has a simple representation in terms of the sine function

X has quantile function F^{-1} given by $F^{-1}(p) = a + w \sin^2\left(\frac{\pi}{2}p\right)$ for $p \in [0, 1]$ In particular, the quantiles of X are

1. $q_1 = a + w \sin^2\left(\frac{\pi}{8}\right) = a + \frac{1}{4}(2 - \sqrt{2})w$, the first quartile
2. $q_2 = a + \frac{1}{2}w$, the median
3. $q_3 = a + w \sin^2\left(\frac{3\pi}{8}\right) = a + \frac{1}{4}(2 + \sqrt{2})w$, the third quartile

Proof

Recall that $F^{-1}(p) = a + wG^{-1}(p)$ where G^{-1} is the [quantile function of the standard arcsine distribution](#).

Open the Special Distribution Calculator and select the arcsine distribution. Vary the parameters and note the shape and location of the distribution function. For various values of the parameters, compute selected values of the distribution function and the quantile function.

Moments

Again, we assume that X has the arcsine distribution with location parameter $a \in \mathbb{R}$ and scale parameter $w \in (0, \infty)$. First we give the mean and variance.

The mean and variance of X are

1. $\mathbb{E}(X) = a + \frac{1}{2}w$
2. $\text{var}(X) = \frac{1}{8}w^2$

Proof

These results from the representation $X = a + wZ$ and the results for the [mean and variance of \$Z\$](#) .

Open the Special Distribution Simulator and select the arcsine distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The moments of X can be obtained from the moments of Z , but the results are messy, except when the location parameter is 0.

Suppose the location parameter $a = 0$. For $n \in \mathbb{N}$,

$$\mathbb{E}(X^n) = w^n \prod_{j=0}^{n-1} \frac{2j+1}{2j+2} \quad (5.19.14)$$

Proof

This follows from the representation $X = wZ$ and the results for the [moments of \$Z\$](#) .

The moment generating function can be expressed as a series with product coefficients, and so is not particularly helpful.

X has moment generating function M given by

$$M(t) = \mathbb{E}(e^{tX}) = e^{at} \sum_{n=0}^{\infty} \left(\prod_{j=0}^{n-1} \frac{2j+1}{2j+2} \right) \frac{w^n t^n}{n!}, \quad t \in \mathbb{R} \quad (5.19.15)$$

Proof

Recall that $M(t) = e^{at} m(wt)$ where m is the [moment generating function of \$Z\$](#) .

Finally, the skewness and kurtosis are unchanged.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = \frac{3}{2}$

Proof

Recall that the skewness and kurtosis are defined in terms of the standard score of X and hence are invariant under a location-scale transformation.

Related Distributions

By construction, the general arcsine distribution is a location-scale family, and so is closed under location-scale transformations.

If X has the arcsine distribution with location parameter $a \in \mathbb{R}$ and scale parameter $w \in (0, \infty)$ and if $c \in \mathbb{R}$ and $d \in (0, \infty)$ then $c + dX$ has the arcsine distribution with location parameter $c + ad$ scale parameter dw .

Proof

By [definition](#) we can take $X = a + wZ$ where Z has the standard arcsine distribution. Hence $c + dX = (c + da) + (dw)Z$.

Since the [quantile function](#) is in closed form, the arcsine distribution can be simulated by the random quantile method.

Suppose that $a \in \mathbb{R}$ and $w \in (0, \infty)$.

1. If U has the standard uniform distribution (a random number) then $X = a + w \sin^2\left(\frac{\pi}{2}U\right)$ has the arcsine distribution with location parameter a and scale parameter w .
2. If X has the arcsine distribution with location parameter a and scale parameter w then $U = \frac{2}{\pi} \arcsin\left(\sqrt{\frac{X-a}{w}}\right)$ has the standard uniform distribution.

Open the random quantile simulator and select the arcsine distribution. Vary the parameters and note the location and shape of the probability density function. For selected parameter values, run the experiment 1000 times and compare the empirical probability density function, mean, and standard deviation to their distributional counterparts. Note how the random quantiles simulate the distribution.

The following exercise illustrates the connection between the Brownian motion process and the arcsine distribution.

Open the Brownian motion simulator and select the last zero random variable. Vary the time parameter t and note that the last zero has the arcsine distribution on the interval $(0, t)$. Run the experiment 1000 time and compare the empirical probability density function, mean, and standard deviation to their distributional counterparts. Note how the last zero simulates the distribution.

This page titled [5.19: The Arcsine Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.20: General Uniform Distributions

This section explores uniform distributions in an abstract setting. If you are a new student of probability, or are not familiar with measure theory, you may want to skip this section and read the sections on the uniform distribution on an interval and the discrete uniform distributions.

Basic Theory

Definition

Suppose that $(S, \mathcal{S}, \lambda)$ is a measure space. That is, S is a set, \mathcal{S} a σ -algebra of subsets of S , and λ a positive measure on \mathcal{S} . Suppose also that $0 < \lambda(S) < \infty$, so that λ is a finite, positive measure.

Random variable X with values in S has the *uniform distribution* on S (with respect to λ) if

$$\mathbb{P}(X \in A) = \frac{\lambda(A)}{\lambda(S)}, \quad A \in \mathcal{S} \quad (5.20.1)$$

Thus, the probability assigned to a set $A \in \mathcal{S}$ depends only on the size of A (as measured by λ).

The most common special cases are as follows:

1. *Discrete*: The set S is finite and non-empty, \mathcal{S} is the σ -algebra of all subsets of S , and $\lambda = \#$ (counting measure).
2. *Euclidean*: For $n \in \mathbb{N}_+$, let \mathcal{R}_n denote the σ -algebra of Borel measurable subsets of \mathbb{R}^n and let λ_n denote Lebesgue measure on $(\mathbb{R}^n, \mathcal{R}_n)$. In this setting, $S \in \mathcal{R}_n$ with $0 < \lambda_n(S) < \infty$, $\mathcal{S} = \{A \in \mathcal{R}_n : A \subseteq S\}$, and the measure is λ_n restricted to (S, \mathcal{S}) .

In the Euclidean case, recall that λ_1 is length measure on \mathbb{R} , λ_2 is area measure on \mathbb{R}^2 , λ_3 is volume measure on \mathbb{R}^3 , and in general λ_n is sometimes referred to as n -dimensional volume. Thus, $S \in \mathcal{R}_n$ is a set with positive, finite volume.

Properties

Suppose $(S, \mathcal{S}, \lambda)$ is a finite, positive measure space, as above, and that X is uniformly distributed on S .

The probability density function f of X (with respect to λ) is

$$f(x) = \frac{1}{\lambda(S)}, \quad x \in S \quad (5.20.2)$$

Proof

This follows directly from the definition of probability density function:

$$\int_A \frac{1}{\lambda(S)} d\lambda(x) = \frac{\lambda(A)}{\lambda(S)}, \quad A \in \mathcal{S} \quad (5.20.3)$$

Thus, the defining property of the uniform distribution on a set is constant density on that set. Another basic property is that uniform distributions are preserved under conditioning.

Suppose that $R \in \mathcal{S}$ with $\lambda(R) > 0$. The conditional distribution of X given $X \in R$ is uniform on R .

Proof

For $A \in \mathcal{S}$ with $A \subseteq R$,

$$\mathbb{P}(X \in A \mid X \in R) = \frac{\mathbb{P}(X \in A)}{\mathbb{P}(X \in R)} = \frac{\lambda(A)/\lambda(S)}{\lambda(R)/\lambda(S)} = \frac{\lambda(A)}{\lambda(R)} \quad (5.20.4)$$

In the setting of previous result, suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent variables, each uniformly distributed on S . Let $N = \min\{n \in \mathbb{N}_+ : X_n \in R\}$. Then N has the geometric distribution on \mathbb{N}_+ with success parameter $p = \mathbb{P}(X \in R)$. More importantly, the distribution of X_N is the same as the conditional distribution of X given $X \in R$, and hence is uniform on R . This is the basis of the *rejection method* of simulation. If we can simulate a uniform distribution on S , then we can simulate a uniform distribution on R .

If h is a real-valued function on S , then $\mathbb{E}[h(X)]$ is the average value of h on S , as measured by λ :

If $h : S \rightarrow \mathbb{R}$ is integrable with respect to λ Then

$$\mathbb{E}[h(X)] = \frac{1}{\lambda(S)} \int_S h(x) d\lambda(x) \quad (5.20.5)$$

Proof

This result follows from the change of variables theorem for expected value, since

$$\mathbb{E}[h(X)] = \int_S h(x) f(x) d\lambda(x) = \frac{1}{\lambda(S)} \int_S h(x) d\lambda(x) \quad (5.20.6)$$

The entropy of the uniform distribution on S depends only on the size of S , as measured by λ :

The entropy of X is $H(X) = \ln[\lambda(S)]$.

Proof

$$H(X) = \mathbb{E}\{-\ln[f(X)]\} = \int_S -\ln\left(\frac{1}{\lambda(S)}\right) \frac{1}{\lambda(S)} = -\ln\left(\frac{1}{\lambda(S)}\right) = \ln[\lambda(S)] \quad (5.20.7)$$

Product Spaces

Suppose now that $(S, \mathcal{S}, \lambda)$ and (T, \mathcal{T}, μ) are finite, positive measure spaces, so that $0 < \lambda(S) < \infty$ and $0 < \mu(T) < \infty$. Recall the product space $(S \times T, \mathcal{S} \otimes \mathcal{T}, \lambda \otimes \mu)$. The *product σ -algebra* $\mathcal{S} \otimes \mathcal{T}$ is the σ -algebra of subsets of $S \times T$ generated by product sets $A \times B$ where $A \in \mathcal{S}$ and $B \in \mathcal{T}$. The *product measure* $\lambda \otimes \mu$ is the unique positive measure on $(S \times T, \mathcal{S} \otimes \mathcal{T})$ that satisfies $(\lambda \otimes \mu)(A \times B) = \lambda(A)\mu(B)$ for $A \in \mathcal{S}$ and $B \in \mathcal{T}$.

(X, Y) is uniformly distributed on $S \times T$ if and only if X is uniformly distributed on S , Y is uniformly distributed on T , and X and Y are independent.

Proof

Suppose first that (X, Y) is uniformly distributed on $S \times T$. If $A \in \mathcal{S}$ and $B \in \mathcal{T}$ then

$$\mathbb{P}(X \in A, Y \in B) = \mathbb{P}[(X, Y) \in A \times B] = \frac{(\lambda \otimes \mu)(A \times B)}{(\lambda \otimes \mu)(S \times T)} = \frac{\lambda(A)\mu(B)}{\lambda(S)\mu(T)} = \frac{\lambda(A)}{\lambda(S)} \frac{\mu(B)}{\mu(T)} \quad (5.20.8)$$

Taking $B = T$ in the displayed equation gives $\mathbb{P}(X \in A) = \lambda(A)/\lambda(S)$ for $A \in \mathcal{S}$, so X is uniformly distributed on S . Taking $A = S$ in the displayed equation gives $\mathbb{P}(Y \in B) = \mu(B)/\mu(T)$ for $B \in \mathcal{T}$, so Y is uniformly distributed on T . Returning to the displayed equation generally gives $\mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B)$ for $A \in \mathcal{S}$ and $B \in \mathcal{T}$, so X and Y are independent.

Conversely, suppose that X is uniformly distributed on S , Y is uniformly distributed on T , and X and Y are independent. Then for $A \in \mathcal{S}$ and $B \in \mathcal{T}$,

$$\mathbb{P}[(X, Y) \in A \times B] = \mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B) = \frac{\lambda(A)}{\lambda(S)} \frac{\mu(B)}{\mu(T)} = \frac{\lambda(A)\mu(B)}{\lambda(S)\mu(T)} = \frac{(\lambda \otimes \mu)(A \times B)}{(\lambda \otimes \mu)(S \times T)} \quad (5.20.9)$$

It then follows (see the section on existence and uniqueness of measures) that $\mathbb{P}[(X, Y) \in C] = (\lambda \otimes \mu)(C)/(\lambda \otimes \mu)(S \times T)$ for every $C \in \mathcal{S} \otimes \mathcal{T}$, so (X, Y) is uniformly distributed on $S \times T$.

This page titled 5.20: General Uniform Distributions is shared under a CC BY 2.0 license and was authored, remixed, and/or curated by Kyle Siegrist (Random Services) via source content that was edited to the style and standards of the LibreTexts platform.

5.21: The Uniform Distribution on an Interval

The *continuous uniform distribution* on an interval of \mathbb{R} is one of the simplest of all probability distributions, but nonetheless very important. In particular, continuous uniform distributions are the basic tools for simulating other probability distributions. The uniform distribution corresponds to picking a point *at random* from the interval. The uniform distribution on an interval is a special case of the general uniform distribution with respect to a measure, in this case Lebesgue measure (length measure) on \mathbb{R} .

The Standard Uniform Distribution

Definition

The continuous uniform distribution on the interval $[0, 1]$ is known as the *standard uniform distribution*. Thus if U has the standard uniform distribution then

$$\mathbb{P}(U \in A) = \lambda(A) \quad (5.21.1)$$

for every (Borel measurable) subset A of $[0, 1]$, where λ is Lebesgue (length) measure.

A simulation of a random variable with the standard uniform distribution is known in computer science as a *random number*. All programming languages have functions for computing random numbers, as do calculators, spreadsheets, and mathematical and statistical software packages.

Distribution Functions

Suppose that U has the standard uniform distribution. By definition, the probability density function is constant on $[0, 1]$.

U has probability density function g given by $g(u) = 1$ for $u \in [0, 1]$.

Since the density function is constant, the mode is not meaningful.

Open the Special Distribution Simulator and select the continuous uniform distribution. Keep the default parameter values. Run the simulation 1000 times and compare the empirical density function and to the probability density function.

The distribution function is simply the identity function on $[0, 1]$.

U has distribution function G given by $G(u) = u$ for $u \in [0, 1]$.

Proof

Note that $\mathbb{P}(U \leq u) = \lambda[0, u] = u$ for $u \in [0, 1]$. Recall again that λ is length measure.

The quantile function is the same as the distribution function.

U has quantile function G^{-1} given by $G^{-1}(p) = p$ for $p \in [0, 1]$. The quartiles are

1. $q_1 = \frac{1}{4}$, the first quartile
2. $q_2 = \frac{1}{2}$, the median
3. $q_3 = \frac{3}{4}$, the third quartile

Proof

G^{-1} is the ordinary inverse of G on the interval $[0, 1]$, which is G itself since G is the identity function.

Open the Special Distribution Calculator and select the continuous uniform distribution. Keep the default parameter values. Compute a few values of the distribution function and the quantile function.

Moments

Suppose again that U has the standard uniform distribution. The moments (about 0) are simple.

For $n \in \mathbb{N}$,

$$\mathbb{E}(U^n) = \frac{1}{n+1} \quad (5.21.2)$$

Proof

Since the PDF is 1 on $[0, 1]$,

$$\mathbb{E}(U^n) = \int_0^1 u^n du = \frac{1}{n+1} \quad (5.21.3)$$

The mean and variance follow easily from the general moment formula.

The mean and variance of U are

1. $\mathbb{E}(U) = \frac{1}{2}$
2. $\text{var}(U) = \frac{1}{12}$

Open the Special Distribution Simulator and select the continuous uniform distribution. Keep the default parameter values. Run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

Next are the skewness and kurtosis.

The skewness and kurtosis of U are

1. $\text{skew}(U) = 0$
2. $\text{kurt}(U) = \frac{9}{5}$

Proof

1. This follows from the symmetry of the distribution about the mean $\frac{1}{2}$.
2. This follows from the usual formula for kurtosis in terms of the [moments](#), or directly, since $\sigma^4 = \frac{1}{144}$ and

$$\mathbb{E} \left[\left(U - \frac{1}{2} \right)^4 \right] = \int_0^1 \left(x - \frac{1}{2} \right)^4 dx = \frac{1}{80} \quad (5.21.4)$$

Thus, the *excess kurtosis* is $\text{kurt}(U) - 3 = -\frac{6}{5}$

Finally, we give the moment generating function.

The moment generating function m of U is given by $m(0) = 1$ and

$$m(t) = \frac{e^t - 1}{t}, \quad t \in \mathbb{R} \setminus \{0\} \quad (5.21.5)$$

Proof

Again, since the PDF is 1 on $[0, 1]$

$$\mathbb{E}(e^{tU}) = \int_0^1 e^{tu} du = \frac{e^t - 1}{t}, \quad t \neq 0 \quad (5.21.6)$$

Trivially $m(0) = 1$.

Related Distributions

The standard uniform distribution is connected to every other probability distribution on \mathbb{R} by means of the quantile function of the other distribution. When the quantile function has a simple closed form expression, this result forms the primary method of simulating the other distribution with a random number.

Suppose that F is the distribution function for a probability distribution on \mathbb{R} , and that F^{-1} is the corresponding quantile function. If U has the standard uniform distribution, then $X = F^{-1}(U)$ has distribution function F .

Proof

A basic property of quantile functions is that $F(x) \leq p$ if and only if $x \leq F^{-1}(p)$ for $x \in \mathbb{R}$ and $p \in (0, 1)$. Hence from the [distribution function of \$U\$](#) ,

$$\mathbb{P}(X \leq x) = \mathbb{P}[F^{-1}(U) \leq x] = \mathbb{P}[U \leq F(x)] = F(x), \quad x \in \mathbb{R} \quad (5.21.7)$$

Open the Random Quantile Experiment. For each distribution, run the simulation 1000 times and compare the empirical density function to the probability density function of the selected distribution. Note how the random quantiles simulate the distribution.

For a continuous distribution on an interval of \mathbb{R} , the connection goes the other way.

Suppose that X has a continuous distribution on an interval $I \subseteq \mathbb{R}$, with distribution function F . Then $U = F(X)$ has the standard uniform distribution.

Proof

For $u \in (0, 1)$ recall that $F^{-1}(u)$ is a quantile of order u . Since X has a continuous distribution,

$$\mathbb{P}(U \geq u) = \mathbb{P}[F(X) \geq u] = \mathbb{P}[X \geq F^{-1}(u)] = 1 - F[F^{-1}(u)] = 1 - u \quad (5.21.8)$$

Hence U is uniformly distributed on $(0, 1)$.

The standard uniform distribution is a special case of the beta distribution.

The beta distribution with left parameter $a = 1$ and right parameter $b = 1$ is the standard uniform distribution.

Proof

The beta distribution with parameters $a > 0$ and $b > 0$ has PDF

$$x \mapsto \frac{1}{B(a, b)} x^{a-1} (1-x)^{b-1}, \quad x \in (0, 1) \quad (5.21.9)$$

where B is the beta function. With $a = b = 1$, the PDF is the standard uniform PDF.

The standard uniform distribution is also the building block of the Irwin-Hall distributions.

The Uniform Distribution on a General Interval

Definition

The standard uniform distribution is generalized by adding location-scale parameters.

Suppose that U has the standard uniform distribution. For $a \in \mathbb{R}$ and $w \in (0, \infty)$ random variable $X = a + wU$ has the *uniform distribution* with location parameter a and scale parameter w .

Distribution Functions

Suppose that X has the uniform distribution with location parameter $a \in \mathbb{R}$ and scale parameter $w \in (0, \infty)$.

X has probability density function f given by $f(x) = 1/w$ for $x \in [a, a + w]$.

Proof

Recall that $f(x) = \frac{1}{w} g\left(\frac{x-a}{w}\right)$ for $x \in [a, a + w]$, where g is the [standard uniform PDF](#). But $g(u) = 1$ for $u \in [0, 1]$, so the result follows.

The last result shows that X really does have a *uniform distribution*, since the probability density function is constant on the support interval. Moreover, we can clearly parameterize the distribution by the endpoints of this interval, namely a and $b = a + w$, rather than by the location, scale parameters a and w . In fact, the distribution is more commonly known as the *uniform distribution* on the interval $[a, b]$. Nonetheless, it is useful to know that the distribution is the location-scale family associated with the standard uniform distribution. In terms of the endpoint parameterization,

$$f(x) = \frac{1}{b-a}, \quad x \in [a, b] \quad (5.21.10)$$

Open the Special Distribution Simulator and select the uniform distribution. Vary the location and scale parameters and note the graph of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has distribution function F given by

$$F(x) = \frac{x-a}{w}, \quad x \in [a, a+w] \quad (5.21.11)$$

Proof

Recall that $F(x) = G\left(\frac{x-a}{w}\right)$ for $x \in [a, a+w]$, where G is the [standard uniform CDF](#). But $G(u) = u$ for $u \in [0, 1]$ so the result follows. Of course, a direct proof using the PDF is also easy.

In terms of the endpoint parameterization,

$$F(x) = \frac{x-a}{b-a}, \quad x \in [a, b] \quad (5.21.12)$$

X has quantile function F^{-1} given by $F^{-1}(p) = a + pw = (1-p)a + pb$ for $p \in [0, 1]$. The quartiles are

1. $q_1 = a + \frac{1}{4}w = \frac{3}{4}a + \frac{1}{4}b$, the first quartile
2. $q_2 = a + \frac{1}{2}w = \frac{1}{2}a + \frac{1}{2}b$, the median
3. $q_3 = a + \frac{3}{4}w = \frac{1}{4}a + \frac{3}{4}b$, the third quartile

Proof

Recall that $F^{-1}(p) = a + wG^{-1}(p)$ where G^{-1} is the [standard uniform quantile function](#). But $G^{-1}(p) = p$ for $p \in [0, 1]$ so the result follows. Of course a direct proof from the CDF is also easy.

Open the Special Distribution Calculator and select the uniform distribution. Vary the parameters and note the graph of the distribution function. For selected values of the parameters, compute a few values of the distribution function and the quantile function.

Moments

Again we assume that X has the uniform distribution on the interval $[a, b]$ where $a, b \in \mathbb{R}$ and $a < b$. Thus the location parameter is a and the scale parameter $w = b - a$.

The moments of X are

$$\mathbb{E}(X^n) = \frac{b^{n+1} - a^{n+1}}{(n+1)(b-a)}, \quad n \in \mathbb{N} \quad (5.21.13)$$

Proof

For $n \in \mathbb{N}$,

$$\mathbb{E}(X^n) = \int_a^b x^n \frac{1}{b-a} dx = \frac{b^{n+1} - a^{n+1}}{(n+1)(b-a)} \quad (5.21.14)$$

The mean and variance of X are

1. $\mathbb{E}(X) = \frac{1}{2}(a+b)$
2. $\text{var}(X) = \frac{1}{12}(b-a)^2$

Open the Special Distribution Simulator and select the uniform distribution. Vary the parameters and note the location and size of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = \frac{9}{5}$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score and hence are invariant under location-scale transformations.

Once again, the excess kurtosis is $\text{kurt}(X) - 3 = -\frac{6}{5}$.

The moment generating function M of X is given by $M(0) = 1$ and

$$M(t) = \frac{e^{bt} - e^{at}}{t(b-a)}, \quad t \in \mathbb{R} \setminus \{0\} \quad (5.21.15)$$

Proof

Recall that $M(t) = e^{at}m(wt)$ where m is the [standard uniform MGF](#). Substituting gives the result.

If h is a real-valued function on $[a, b]$, then $\mathbb{E}[h(X)]$ is the *average value* of h on $[a, b]$, as defined in calculus:

If $h : [a, b] \rightarrow \mathbb{R}$ is integrable, then

$$\mathbb{E}[h(X)] = \frac{1}{b-a} \int_a^b h(x) dx \quad (5.21.16)$$

Proof

This follows from the change of variables formula for expected value: $\mathbb{E}[h(X)] = \int_a^b h(x)f(x) dx$.

The entropy of the uniform distribution on an interval depends only on the length of the interval.

The entropy of X is $H(X) = \ln(b-a)$.

Proof

$$H(X) = \mathbb{E}\{-\ln[f(X)]\} = \int_a^b -\ln\left(\frac{1}{b-a}\right) \frac{1}{b-a} dx = -\ln\left(\frac{1}{b-a}\right) = \ln(b-a) \quad (5.21.17)$$

Related Distributions

Since the uniform distribution is a location-scale family, it is trivially closed under location-scale transformations.

If X has the uniform distribution with location parameter a and scale parameter w , and if $c \in \mathbb{R}$ and $d \in (0, \infty)$, then $Y = c + dX$ has the uniform distribution with location parameter $c + da$ and scale parameter dw .

Proof

From the [definition](#), we can take $X = a + wU$ where U has the standard uniform distribution. Hence $Y = c + dX = (c + da) + (dw)U$.

As we saw above, the standard uniform distribution is a basic tool in the random quantile method of simulation. Uniform distributions on intervals are also basic in the *rejection method* of simulation. We sketch the method in the next paragraph; see the

section on general uniform distributions for more theory.

Suppose that h is a probability density function for a continuous distribution with values in a bounded interval $(a, b) \subseteq \mathbb{R}$. Suppose also that h is bounded, so that there exists $c > 0$ such that $h(x) \leq c$ for all $x \in (a, b)$. Let $\mathbf{X} = (X_1, X_2, \dots)$ be a sequence of independent variables, each uniformly distributed on (a, b) , and let $\mathbf{Y} = (Y_1, Y_2, \dots)$ be a sequence of independent variables, each uniformly distributed on $(0, c)$. Finally, assume that \mathbf{X} and \mathbf{Y} are independent. Then $((X_1, Y_1), (X_2, Y_2), \dots)$ is a sequence of independent variables, each uniformly distributed on $(a, b) \times (0, c)$. Let $N = \min\{n \in \mathbb{N}_+ : 0 < Y_n < h(X_n)\}$. Then (X_N, Y_N) is uniformly distributed on $R = \{(x, y) \in (a, b) \times (0, c) : y < h(x)\}$ (the region under the graph of h), and therefore X_N has probability density function h . In words, we generate uniform points in the rectangular region $(a, b) \times (0, c)$ until we get a point under the graph of h . The x -coordinate of that point is our simulated value. The rejection method can be used to approximately simulate random variables when the region under the density function is unbounded.

Open the rejection method simulator. For each distribution, select a set of parameter values. Run the experiment 2000 times and observe how the rejection method works. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

This page titled [5.21: The Uniform Distribution on an Interval](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.22: Discrete Uniform Distributions

Uniform Distributions on a Finite Set

Suppose that S is a nonempty, finite set. A random variable X taking values in S has the *uniform distribution* on S if

$$\mathbb{P}(X \in A) = \frac{\#(A)}{\#(S)}, \quad A \subseteq S \quad (5.22.1)$$

The discrete uniform distribution is a special case of the general uniform distribution with respect to a measure, in this case counting measure. The distribution corresponds to picking an element of S *at random*. Most classical, combinatorial probability models are based on underlying discrete uniform distributions. The chapter on Finite Sampling Models explores a number of such models.

The probability density function f of X is given by

$$f(x) = \frac{1}{\#(S)}, \quad x \in S \quad (5.22.2)$$

Proof

This follows from the definition of the (discrete) probability density function: $\mathbb{P}(X \in A) = \sum_{x \in A} f(x)$ for $A \subseteq S$. Or more simply, $f(x) = \mathbb{P}(X = x) = 1/\#(S)$.

Like all uniform distributions, the discrete uniform distribution on a finite set is characterized by the property of constant density on the set. Another property that all uniform distributions share is invariance under conditioning on a subset.

Suppose that R is a nonempty subset of S . Then the conditional distribution of X given $X \in R$ is uniform on R .

Proof

For $A \subseteq R$,

$$\mathbb{P}(X \in A \mid X \in R) = \frac{\mathbb{P}(X \in A)}{\mathbb{P}(X \in R)} = \frac{\#(A)/\#(S)}{\#(R)/\#(S)} = \frac{\#(A)}{\#(R)} \quad (5.22.3)$$

If $h : S \rightarrow \mathbb{R}$ then the expected value of $h(X)$ is simply the arithmetic average of the values of h :

$$\mathbb{E}[h(X)] = \frac{1}{\#(S)} \sum_{x \in S} h(x) \quad (5.22.4)$$

Proof

This follows from the change of variables theorem for expected value:

$$\mathbb{E}[h(X)] = \sum_{x \in S} f(x)h(x) = \frac{1}{\#(S)} \sum_{x \in S} h(x) \quad (5.22.5)$$

The entropy of X depends only on the number of points in S .

The entropy of X is $H(X) = \ln[\#(S)]$.

Proof

Let $n = \#(S)$. Then

$$H(X) = \mathbb{E}\{-\ln[f(X)]\} = \sum_{x \in S} -\ln\left(\frac{1}{n}\right) \frac{1}{n} = -\ln\left(\frac{1}{n}\right) = \ln(n) \quad (5.22.6)$$

Uniform Distributions on Finite Subsets of \mathbb{R}

Without some additional structure, not much more can be said about discrete uniform distributions. Thus, suppose that $n \in \mathbb{N}_+$ and that $S = \{x_1, x_2, \dots, x_n\}$ is a subset of \mathbb{R} with n points. We will assume that the points are indexed in order, so that $x_1 < x_2 < \dots < x_n$. Suppose that X has the uniform distribution on S .

The probability density function f of X is given by $f(x) = \frac{1}{n}$ for $x \in S$.

The distribution function F of X is given by

1. $F(x) = 0$ for $x < x_1$
2. $F(x) = \frac{k}{n}$ for $x_k \leq x < x_{k+1}$ and $k \in \{1, 2, \dots, n-1\}$
3. $F(x) = 1$ for $x > x_n$

Proof

This follows from the definition of the distribution function: $F(x) = \mathbb{P}(X \leq x)$ for $x \in \mathbb{R}$.

The quantile function F^{-1} of X is given by $F^{-1}(p) = x_{[np]}$ for $p \in (0, 1]$.

Proof

By definition, $F^{-1}(p) = x_k$ for $\frac{k-1}{n} < p \leq \frac{k}{n}$ and $k \in \{1, 2, \dots, n\}$. It follows that $k = [np]$ in this formulation.

The moments of X are ordinary arithmetic averages.

For $k \in \mathbb{N}$

$$\mathbb{E}(X^k) = \frac{1}{n} \sum_{i=1}^n x_i^k \quad (5.22.7)$$

In particular,

The mean and variance of X are

1. $\mu = \frac{1}{n} \sum_{i=1}^n x_i$
2. $\sigma^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2$

Uniform Distributions on Discrete Intervals

We specialize further to the case where the finite subset of \mathbb{R} is a *discrete interval*, that is, the points are uniformly spaced.

The Standard Distribution

Suppose that $n \in \mathbb{N}_+$ and that Z has the discrete uniform distribution on $S = \{0, 1, \dots, n-1\}$. The distribution of Z is the *standard discrete uniform distribution* with n points.

Of course, the results in the previous subsection apply with $x_i = i-1$ and $i \in \{1, 2, \dots, n\}$.

The probability density function g of Z is given by $g(z) = \frac{1}{n}$ for $z \in S$.

Open the Special Distribution Simulation and select the discrete uniform distribution. Vary the number of points, but keep the default values for the other parameters. Note the graph of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function G of Z is given by $G(z) = \frac{1}{n}(\lfloor z \rfloor + 1)$ for $z \in [0, n-1]$.

Proof

Note that $G(z) = \frac{k}{n}$ for $k-1 \leq z < k$ and $k \in \{1, 2, \dots, n\}$. Thus $k-1 = \lfloor z \rfloor$ in this formulation.

The quantile function G^{-1} of Z is given by $G^{-1}(p) = \lceil np \rceil - 1$ for $p \in (0, 1]$. In particular

1. $G^{-1}(1/4) = \lceil n/4 \rceil - 1$ is the first quartile.
2. $G^{-1}(1/2) = \lceil n/2 \rceil - 1$ is the median.
3. $G^{-1}(3/4) = \lceil 3n/4 \rceil - 1$ is the third quartile.

Proof

Note that $G^{-1}(p) = k - 1$ for $\frac{k-1}{n} < p \leq \frac{k}{n}$ and $k \in \{1, 2, \dots, n\}$. Thus $k = \lceil np \rceil$ in this formulation.

Open the special distribution calculator and select the discrete uniform distribution. Vary the number of points, but keep the default values for the other parameters. Note the graph of the distribution function. Compute a few values of the distribution function and the quantile function.

For the standard uniform distribution, results for the moments can be given in closed form.

The mean and variance of Z are

1. $\mathbb{E}(Z) = \frac{1}{2}(n-1)$
2. $\text{var}(Z) = \frac{1}{12}(n^2-1)$

Proof

Recall that

$$\sum_{k=0}^{n-1} k = \frac{1}{2}n(n-1) \quad (5.22.8)$$

$$\sum_{k=0}^{n-1} k^2 = \frac{1}{6}n(n-1)(2n-1) \quad (5.22.9)$$

Hence $\mathbb{E}(Z) = \frac{1}{2}(n-1)$ and $\mathbb{E}(Z^2) = \frac{1}{6}(n-1)(2n-1)$. Part (b) follows from $\text{var}(Z) = \mathbb{E}(Z^2) - [\mathbb{E}(Z)]^2$.

Open the Special Distribution Simulation and select the discrete uniform distribution. Vary the number of points, but keep the default values for the other parameters. Note the size and location of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = \frac{3}{5} \frac{3n^2-7}{n^2-1}$

Proof

Recall that

$$\sum_{k=1}^{n-1} k^3 = \frac{1}{4}(n-1)^2 n^2 \quad (5.22.10)$$

$$\sum_{k=1}^{n-1} k^4 = \frac{1}{30}(n-1)(2n-1)(3n^2-3n-1) \quad (5.22.11)$$

Hence $\mathbb{E}(Z^3) = \frac{1}{4}(n-1)^2 n$ and $\mathbb{E}(Z^4) = \frac{1}{30}(n-1)(2n-1)(3n^2-3n-1)$. The results now follow from the results on the [mean and variance](#) and the standard formulas for skewness and kurtosis. Of course, the fact that $\text{skew}(Z) = 0$ also follows from the symmetry of the distribution.

Note that $\text{skew}(Z) \rightarrow \frac{9}{5}$ as $n \rightarrow \infty$. The limiting value is the skewness of the uniform distribution on an interval.

Z has probability generating function P given by $P(1) = 1$ and

$$P(t) = \frac{1}{n} \frac{1-t^n}{1-t}, \quad t \in \mathbb{R} \setminus \{1\} \quad (5.22.12)$$

Proof

$$P(t) = \mathbb{E}(t^Z) = \frac{1}{n} \sum_{k=0}^{n-1} t^k = \frac{1}{n} \frac{1-t^n}{1-t} \quad (5.22.13)$$

The General Distribution

We now generalize the standard discrete uniform distribution by adding location and scale parameters.

Suppose that Z has the standard discrete uniform distribution on $n \in \mathbb{N}_+$ points, and that $a \in \mathbb{R}$ and $h \in (0, \infty)$. Then $X = a + hZ$ has the *uniform distribution* on n points with *location parameter* a and *scale parameter* h .

Note that X takes values in

$$S = \{a, a+h, a+2h, \dots, a+(n-1)h\} \quad (5.22.14)$$

so that S has n elements, starting at a , with step size h , a *discrete interval*. In the further special case where $a \in \mathbb{Z}$ and $h = 1$, we have an *integer interval*. Note that the last point is $b = a + (n-1)h$, so we can clearly also parameterize the distribution by the endpoints a and b , and the step size h . With this parametrization, the number of points is $n = 1 + (b-a)/h$. For the remainder of this discussion, we assume that X has the distribution in the definition. Our first result is that the distribution of X really is uniform.

X has probability density function f given by $f(x) = \frac{1}{n}$ for $x \in S$

Proof

Recall that $f(x) = g\left(\frac{x-a}{h}\right)$ for $x \in S$, where g is the [PDF of \$Z\$](#) .

Open the Special Distribution Simulation and select the discrete uniform distribution. Vary the parameters and note the graph of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function F of x is given by

$$F(x) = \frac{1}{n} \left(\left\lfloor \frac{x-a}{h} \right\rfloor + 1 \right), \quad x \in [a, b] \quad (5.22.15)$$

Proof

Recall that $F(x) = G\left(\frac{x-a}{h}\right)$ for $x \in S$, where G is the [CDF of \$Z\$](#) .

The quantile function F^{-1} of X is given by $G^{-1}(p) = a + h(\lceil np \rceil - 1)$ for $p \in (0, 1]$. In particular

1. $F^{-1}(1/4) = a + h(\lceil n/4 \rceil - 1)$ is the first quartile.
2. $F^{-1}(1/2) = a + h(\lceil n/2 \rceil - 1)$ is the median.
3. $F^{-1}(3/4) = a + h(\lceil 3n/4 \rceil - 1)$ is the third quartile.

Proof

Recall that $F^{-1}(p) = a + hG^{-1}(p)$ for $p \in (0, 1]$, where G^{-1} is the [quantile function of \$Z\$](#) .

Open the special distribution calculator and select the discrete uniform distribution. Vary the parameters and note the graph of the distribution function. Compute a few values of the distribution function and the quantile function.

The mean and variance of X are

1. $\mathbb{E}(X) = a + \frac{1}{2}(n-1)h = \frac{1}{2}(a+b)$

$$2. \text{var}(X) = \frac{1}{12}(n^2 - 1)h^2 = \frac{1}{12}(b - a)(b - a + 2h)$$

Proof

Recall that $\mathbb{E}(X) = a + h\mathbb{E}(Z)$ and $\text{var}(X) = h^2\text{var}(Z)$, so the results follow from the corresponding results for the [standard distribution](#).

Note that the mean is the average of the endpoints (and so is the midpoint of the interval $[a, b]$) while the variance depends only on the number of points and the step size.

Open the Special Distribution Simulator and select the discrete uniform distribution. Vary the parameters and note the shape and location of the mean/standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The skewness and kurtosis of Z are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = \frac{3}{5} \frac{3n^2 - 7}{n^2 - 1}$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are the skewness and kurtosis of X are the same as the [skewness and kurtosis of \$Z\$](#) .

X has moment generating function M given by $M(0) = 1$ and

$$M(t) = \frac{1}{n} e^{ta} \frac{1 - e^{nth}}{1 - e^{th}}, \quad t \in \mathbb{R} \setminus \{0\} \quad (5.22.16)$$

Proof

Note that $M(t) = \mathbb{E}(e^{tX}) = e^{ta} \mathbb{E}(e^{thZ}) = e^{ta} P(e^{th})$ where P is the [probability generating function of \$Z\$](#) .

Related Distributions

Since the discrete uniform distribution on a discrete interval is a location-scale family, it is trivially closed under location-scale transformations.

Suppose that X has the discrete uniform distribution on $n \in \mathbb{N}_+$ points with location parameter $a \in \mathbb{R}$ and scale parameter $h \in (0, \infty)$. If $c \in \mathbb{R}$ and $w \in (0, \infty)$ then $Y = c + wX$ has the discrete uniform distribution on n points with location parameter $c + wa$ and scale parameter wh .

Proof

By [definition](#) we can take $X = a + hZ$ where Z has the standard uniform distribution on n points. Then $Y = c + wX = (c + wa) + (wh)Z$.

In terms of the endpoint parameterization, X has left endpoint a , right endpoint $a + (n - 1)h$, and step size h while Y has left endpoint $c + wa$, right endpoint $(c + wa) + (n - 1)wh$, and step size wh .

The uniform distribution on a discrete interval converges to the continuous uniform distribution on the interval with the same endpoints, as the step size decreases to 0.

Suppose that X_n has the discrete uniform distribution with endpoints a and b , and step size $(b - a)/n$, for each $n \in \mathbb{N}_+$. Then the distribution of X_n converges to the continuous uniform distribution on $[a, b]$ as $n \rightarrow \infty$.

Proof

The CDF F_n of X_n is given by

$$F_n(x) = \frac{1}{n} \left\lfloor n \frac{x - a}{b - a} \right\rfloor, \quad x \in [a, b] \quad (5.22.17)$$

But $ny - 1 \leq \lfloor ny \rfloor \leq ny$ for $y \in \mathbb{R}$ so $\lfloor ny \rfloor / n \rightarrow y$ as $n \rightarrow \infty$. Hence $F_n(x) \rightarrow (x - a)/(b - a)$ as $n \rightarrow \infty$ for $x \in [a, b]$, and this is the CDF of the continuous uniform distribution on $[a, b]$.

This page titled [5.22: Discrete Uniform Distributions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.23: The Semicircle Distribution

The Semicircle Distribution

The *semicircle distribution* plays a very important role in the study of random matrices. It is also known as the *Wigner distribution* in honor of the physicist Eugene Wigner, who did pioneering work on random matrices.

The Standard Semicircle Distribution

Distribution Functions

The *standard semicircle distribution* is a continuous distribution on the interval $[-1, 1]$ with probability density function g given by

$$g(x) = \frac{2}{\pi} \sqrt{1-x^2}, \quad x \in [-1, 1] \quad (5.23.1)$$

Proof

The graph of $x \mapsto \sqrt{1-x^2}$ for $x \in [-1, 1]$ is the upper half of the circle of radius 1 centered at the origin. Hence the area under this graph is $\pi/2$ and therefore g is a valid PDF—the constant $2/\pi$ in g is the *normalizing constant*

As noted in the proof, $x \mapsto \sqrt{1-x^2}$ for $x \in [-1, 1]$ is the upper half of the circle of radius 1 centered at the origin, hence the name.

The standard semicircle probability density function g satisfies the following properties:

1. g is symmetric about $x = 0$.
2. g increases and then decreases with mode at $x = 0$.
3. g is concave downward.

Proof

As noted earlier, except for the normalizing constant, the graph of g is the upper half of the circle of radius 1 centered at the origin, and so these properties are obvious.

Open special distribution simulator and select the semicircle distribution. With the default parameter value, note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The standard semicircle distribution function G is given by

$$G(x) = \frac{1}{2} + \frac{1}{\pi} x \sqrt{1-x^2} + \frac{1}{\pi} \arcsin x, \quad x \in [-1, 1] \quad (5.23.2)$$

Proof

Of course $G(x) = \int_{-1}^x g(t) dt$ for $-1 \leq x \leq 1$. The integral is evaluated by using the trigonometric substitution $t = \sin \theta$.

We cannot give the quantile function G^{-1} in closed form, but values of this function can be approximated. Clearly by symmetry, $G^{-1}(\frac{1}{2} - p) = -G^{-1}(\frac{1}{2} + p)$ for $0 \leq p \leq \frac{1}{2}$. In particular, the median is 0.

Open the special distribution simulator and select the semicircle distribution. With the default parameter value, note the shape of the distribution function. Compute the first and third quartiles.

Moments

Suppose that X has the standard semicircle distribution. The moments of X about 0 can be computed explicitly. In particular, the odd order moments are 0 by symmetry.

For $n \in \mathbb{N}$, the moment of order $2n + 1$ is $\mathbb{E}(X^{2n+1}) = 0$ and the moment of order $2n$ is

$$\mathbb{E}(X^{2n}) = \left(\frac{1}{2}\right)^{2n} \frac{1}{n+1} \binom{2n}{n} \quad (5.23.3)$$

Proof

Clearly X has moments of all orders since the PDF g is bounded and the support interval is bounded. So by symmetry, the odd order moments are 0, and we just need to prove the result for the even order moments. Note that

$$\mathbb{E}(X^{2n}) = \int_{-1}^1 x^{2n} \frac{2}{\pi} \sqrt{1-x^2} dx \quad (5.23.4)$$

We use the substitution $x = \sin \theta$ to get

$$\mathbb{E}(X^{2n}) = \int_{-\pi/2}^{\pi/2} \frac{2}{\pi} \sin^{2n}(\theta) \cos^2(\theta) d\theta \quad (5.23.5)$$

This integral can be evaluated by standard calculus methods to give the result above.

The numbers $C_n = \frac{1}{n+1} \binom{2n}{n}$ for $n \in \mathbb{N}$ are known as the *Catalan numbers*, and are named for the Belgian mathematician Eugene Catalan. In particular, we can compute the mean, variance, skewness, and kurtosis.

The mean and variance of X are

1. $\mathbb{E}(X) = 0$
2. $\text{var}(X) = \frac{1}{4}$

Open the special distribution simulator and select the semicircle distribution. With the default parameter value, note the size and location of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = 2$

Proof

The standard score of X is $2X$. Hence $\text{skew}(X) = E(2^3 X^3) = 0$. Of course, this is also clear from the symmetry of the distribution of X . Similarly, by the [moment formula](#),

$$\text{kurt}(X) = \mathbb{E}(2^4 X^4) = 2^4 \left(\frac{1}{2}\right)^4 \frac{1}{3} \binom{4}{2} = 2 \quad (5.23.6)$$

It follows that the *excess kurtosis* is $\text{kurt}(X) - 3 = -1$.

Related Distributions

The semicircle distribution has simple connections to the continuous uniform distribution.

If (X, Y) is uniformly distributed on the circular region in \mathbb{R}^2 centered at the origin with radius 1, then X and Y each have the standard semicircular distribution.

Proof

(X, Y) has joint PDF $(x, y) \mapsto 1/\pi$ on $C = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. Hence X has PDF

$$g(x) = \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} \frac{1}{\pi} dy = \frac{2}{\pi} \sqrt{1-x^2}, \quad x \in [-1, 1] \quad (5.23.7)$$

It's easy to simulate a random point that is uniformly distributed on circular region in the previous theorem, and this provides a way of simulating a standard semicircle distribution. This is important since we can't use the random quantile method of simulation.

Suppose that U , V , and W are independent random variables, each with the standard uniform distribution (random numbers). Let $R = \max\{U, V\}$ and $\Theta = 2\pi W$, and then let $X = R \cos \Theta$, $Y = R \sin \Theta$. Then (X, Y) is uniformly distributed on the circular region of radius 1 centered at the origin, and hence X and Y each have the standard semicircle distribution.

Proof

U and V have CDF $u \mapsto u$ for $u \in [0, 1]$ and therefore R has CDF $r \mapsto r^2$ for $r \in [0, 1]$. Hence R has PDF $r \mapsto 2r$ for $r \in [0, 1]$. On the other hand, Θ is uniformly distributed on $[0, 2\pi)$ and hence has density $\theta \mapsto 1/2\pi$ on $[0, 2\pi)$. By independence, the Joint PDF of (R, Θ) is $(r, \theta) \mapsto (2r)(1/2\pi) = r/\pi$ on $\{(r, \theta) : 0 \leq r \leq 1, 0 \leq \theta \leq 2\pi\}$. For the polar coordinate transformation $(x, y) \mapsto (r \cos \theta, r \sin \theta)$, the Jacobian is r . Hence by the change of variables theorem, (X, Y) has PDF

$$(x, y) \mapsto \frac{r}{\pi} \frac{1}{r} = \frac{1}{\pi} \text{ on } \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\} \quad (5.23.8)$$

Of course, note that X and Y in the previous theorem are not independent. Another method of simulation is to use the rejection method. This method works well since the semicircle distribution has a bounded support interval and a bounded probability density function.

Open the rejection method app and select the semicircle distribution. Keep the default parameters to get the standard semicircle distribution. Run the simulation 1000 times and note the points in the scatterplot. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The General Semicircle Distribution

Like so many *standard distributions*, the standard semicircle distribution is usually generalized by adding location and scale parameters.

Definition

Suppose that Z has the standard semicircle distribution. For $a \in \mathbb{R}$ and $r \in (0, \infty)$, $X = a + rZ$ has the *semicircle distribution* with *center* (location parameter) a and *radius* (scale parameter) r .

Distribution Functions

Suppose that X has the semicircle distribution with center $a \in \mathbb{R}$ and radius $r \in (0, \infty)$.

X has probability density function f given by

$$f(x) = \frac{2}{\pi r^2} \sqrt{r^2 - (x - a)^2}, \quad x \in [a - r, a + r] \quad (5.23.9)$$

Proof

This follows from a standard result for location-scale families. Recall that

$$f(x) = \frac{1}{r} g\left(\frac{x - a}{r}\right), \quad \frac{x - a}{r} \in [-1, 1] \quad (5.23.10)$$

where g is the [standard semicircle PDF](#).

The graph of $x \mapsto \sqrt{r^2 - (x - a)^2}$ for $x \in [a - r, a + r]$ is the upper half of the circle of radius r centered at a . The area under this semicircle is $\pi r^2/2$ so as a check on our work, we see that f is a valid probability density function.

The probability density function f of X satisfies the following properties:

1. f is symmetric about $x = a$.
2. f increases and then decreases with mode at $x = a$.

3. f is concave downward.

Open special distribution simulator and select the semicircle distribution. Vary the center a and the radius r , and note the shape of the probability density function. For selected values of a and r , run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function F of X is

$$F(x) = \frac{1}{2} + \frac{x-a}{\pi r^2} \sqrt{r^2 - (x-a)^2} + \frac{1}{\pi} \arcsin\left(\frac{x-a}{r}\right), \quad x \in [a-r, a+r] \quad (5.23.11)$$

Proof

This follows from a standard result for location-scale families:

$$F(x) = G\left(\frac{x-a}{r}\right), \quad \frac{x-a}{r} \in [-1, 1] \quad (5.23.12)$$

where G is the [standard semicircle CDF](#).

As in the standard case, we cannot give the quantile function F^{-1} in closed form, but values of this function can be approximated. Recall that $F^{-1}(p) = a + rG^{-1}(p)$ where G^{-1} is the standard semicircle quantile function. In particular, $F^{-1}\left(\frac{1}{2} - p\right) = 2a - F^{-1}\left(\frac{1}{2} + p\right)$ for $0 \leq p \leq \frac{1}{2}$. The median is a .

Open the special distribution simulator and select the semicircle distribution. Vary the center a and the radius r , and note the shape of the distribution function. For selected values of a and r , compute the first and third quartiles.

Moments

Suppose again that X has the semicircle distribution with center $a \in \mathbb{R}$ and radius $r \in (0, \infty)$, so by [definition](#) we can assume $X = a + rZ$ where Z has the standard semicircle distribution. The moments of X can be computed from the [moments of \$Z\$](#) . Using the binomial theorem and the linearity of expected value we have

$$\mathbb{E}(X^n) = \sum_{k=0}^n \binom{n}{k} r^k a^{n-k} \mathbb{E}(Z^k), \quad n \in \mathbb{N} \quad (5.23.13)$$

In particular,

The mean and variance of X are

1. $\mathbb{E}(X) = a$
2. $\text{var}(X) = r^2/4$

When the center is 0, the general moments have a simple form:

Suppose that $a = 0$. For $n \in \mathbb{N}$ the moment of order $2n+1$ is $\mathbb{E}(X^{2n+1}) = 0$ and the moment of order $2n$ is

$$\mathbb{E}(X^{2n}) = \left(\frac{r}{2}\right)^{2n} \frac{1}{n+1} \binom{2n}{n} \quad (5.23.14)$$

Proof

This follows from the [moment results for \$Z\$](#) since $X^m = r^m Z^m$ for $m \in \mathbb{N}$.

Open the special distribution simulator and select the semicircle distribution. Vary the center a and the radius r , and note the size and location of the mean \pm standard deviation bar. For selected values of a and r , run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = 2$

Proof

These results follow immediately from the [skewness and kurtosis of the standard distribution](#). Recall that skewness and kurtosis are defined in terms of the standard score, which is independent of the location and scale parameters..

Once again, the excess kurtosis is $\text{kurt}(X) - 3 = -1$.

Related Distributions

Since the semicircle distribution is a location-scale family, it's invariant under location-scale transformations.

Suppose that X has the semicircle distribution with center $a \in \mathbb{R}$ and radius $r \in (0, \infty)$. If $b \in \mathbb{R}$ and $c \in (0, \infty)$ then $b + cX$ has the semicircle distribution with center $b + ca$ and radius cr .

Proof

Again from the [definition](#) we can take $X = a + rZ$ where Z has the standard semicircle distribution. Then $b + cX = (b + ca) + (cr)Z$.

One member of the beta family of distributions is a semicircle distribution:

The beta distribution with left parameter $3/2$ and right parameter $3/2$ is the semicircle distribution with center $1/2$ and radius $1/2$.

Proof

By definition, the beta distribution with left and right parameters $3/2$ has PDF

$$f(x) = \frac{1}{B(3/2, 3/2)} x^{1/2} (1-x)^{1/2}, \quad x \in [0, 1] \quad (5.23.15)$$

But $B(3/2, 3/2) = \pi/8$ and $x^{1/2}(1-x)^{1/2} = \sqrt{x-x^2}$. Completing the square gives

$$f(x) = \frac{8}{\pi} \sqrt{\frac{1}{4} - \left(x - \frac{1}{2}\right)^2}, \quad x \in [0, 1] \quad (5.23.16)$$

which is the PDF of the semicircle distribution with center $1/2$ and radius $1/2$

Since we can simulate a variable Z with the standard semicircle distribution by the [method above](#), we can simulate a variable with the semicircle distribution with center $a \in \mathbb{R}$ and radius $r \in (0, \infty)$ by our very definition: $X = a + rZ$. Once again, the rejection method also works well since the support and probability density function of X are bounded.

Open the rejection method app and select the semicircle distribution. For selected values of a and r , run the simulation 1000 times and note the points in the scatterplot. Compare the empirical density function, mean and standard deviation to their distributional counterparts.

This page titled [5.23: The Semicircle Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.24: The Triangle Distribution

Like the semicircle distribution, the *triangle distribution* is based on a simple geometric shape. The distribution arises naturally when uniformly distributed random variables are transformed in various ways.

The Standard Triangle Distribution

Distribution Functions

The *standard triangle distribution* with vertex at $p \in [0, 1]$ (equivalently, *shape parameter* p) is a continuous distribution on $[0, 1]$ with probability density function g described as follows:

1. If $p = 0$ then $g(x) = 2(1 - x)$ for $x \in [0, 1]$
2. If $p = 1$ then $g(x) = 2x$ for $x \in [0, 1]$.
3. If $p \in (0, 1)$ then

$$g(x) = \begin{cases} \frac{2x}{p}, & x \in [0, p] \\ \frac{2(1-x)}{1-p}, & x \in [p, 1] \end{cases} \quad (5.24.1)$$

The shape of the probability density function justifies the name *triangle distribution*.

The graph of g , together with the domain $[0, 1]$, forms a triangle with vertices $(0, 0)$, $(1, 0)$, and $(p, 2)$. The mode of the distribution is $x = p$.

1. If $p = 0$, g is decreasing.
2. If $p = 1$, g is increasing.
3. If $p \in (0, 1)$, g increases and then decreases.

Proof

Using $[0, 1]$ as the base, we can compute the area of the triangle as $\frac{1}{2} \cdot 2 = 1$ so we see immediately that g is a valid probability density function. The properties are obvious.

Open special distribution simulator and select the triangle distribution. Vary p (but keep the default values for the other parameters) and note the shape of the probability density function. For selected values of p , run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function G is given as follows:

1. If $p = 0$, $G(x) = 1 - (1 - x)^2$ for $x \in [0, 1]$.
2. If $p = 1$, $G(x) = x^2$ for $x \in [0, 1]$.
3. If $p \in (0, 1)$

$$G(x) = \begin{cases} \frac{x^2}{p}, & x \in [0, p] \\ 1 - \frac{(1-x)^2}{1-p}, & x \in [p, 1] \end{cases} \quad (5.24.2)$$

Proof

This result follows from standard calculus since $G(x) = \int_0^x g(t) dt$.

The quantile function G^{-1} is given by

$$G^{-1}(u) = \begin{cases} \sqrt{up}, & u \in [0, p] \\ 1 - \sqrt{(1-u)(1-p)}, & u \in [p, 1] \end{cases} \quad (5.24.3)$$

1. The first quartile is $\sqrt{\frac{1}{4}p}$ if $p \in [\frac{1}{4}, 1]$ and is $1 - \sqrt{\frac{3}{4}(1-p)}$ if $p \in [0, \frac{1}{4}]$

2. The median is $\sqrt{\frac{1}{2}p}$ if $p \in [\frac{1}{2}, 1]$ and is $1 - \sqrt{\frac{1}{2}(1-p)}$ if $p \in [0, \frac{1}{2}]$.
3. The third quartile is $\sqrt{\frac{3}{4}p}$ if $p \in [\frac{3}{4}, 1]$ and is $1 - \sqrt{\frac{1}{4}(1-p)}$ if $p \in [0, \frac{3}{4}]$.

Open the special distribution calculator and select the triangle distribution. Vary p (but keep the default values for the other parameters) and note the shape of the distribution function. For selected values of p , compute the first and third quartiles.

Moments

Suppose that X has the standard triangle distribution with vertex $p \in [0, 1]$. The moments are easy to compute.

Suppose that $n \in \mathbb{N}$.

1. If $p = 1$, $\mathbb{E}(X^n) = 2/(n+2)$.
2. If $p \in [0, 1)$,

$$\mathbb{E}(X^n) = \frac{2}{n+1}p^{n+2} + \frac{2}{n+1} \frac{1-p^{n+1}}{1-p} + \frac{2}{n+2} \frac{1-p^{n+2}}{1-p} \quad (5.24.4)$$

Proof

This follows from standard calculus, since $\mathbb{E}(X^n) = \int_0^1 x^n g(x) dx$.

From the general moment formula, we can compute the mean, variance, skewness, and kurtosis.

The mean and variance of X are

1. $\mathbb{E}(X) = \frac{1}{3}(1+p)$
2. $\text{var}(X) = \frac{1}{18}[1-p(1-p)]$

Proof

This follows from the [general moment result](#). Recall that $\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2$.

Note that $\mathbb{E}(X)$ increases from $\frac{1}{3}$ to $\frac{2}{3}$ as p increases from 0 to 1. The graph of $\text{var}(X)$ is a parabola opening downward; the largest value is $\frac{1}{18}$ when $p = 0$ or $p = 1$ and the smallest value is $\frac{1}{24}$ when $p = \frac{1}{2}$.

Open the special distribution simulator and select the triangle distribution. Vary p (but keep the default values for the other parameters) and note the size and location of the mean \pm standard deviation bar. For selected values of p , run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness of X is

$$\text{skew}(X) = \frac{\sqrt{2}(1-2p)(1+p)(2-p)}{5[1-p(1-p)]^{3/2}} \quad (5.24.5)$$

The kurtosis of X is $\text{kurt}(X) = \frac{12}{5}$.

Proof

These results follow from the [general moment result](#) and the computational formulas for skewness and kurtosis.

Note that X is positively skewed for $p < \frac{1}{2}$, negatively skewed for $p > \frac{1}{2}$, and symmetric for $p = \frac{1}{2}$. More specifically, if we indicate the dependence on the parameter p then $\text{skew}_{1-p}(X) = -\text{skew}_p(X)$. Note also that the kurtosis is independent of p , and the excess kurtosis is $\text{kurt}(X) - 3 = -\frac{3}{5}$.

Open the special distribution simulator and select the triangle distribution. Vary p (but keep the default values for the other parameters) and note the degree of symmetry and the degree to which the distribution is peaked. For selected values of p , run the simulation 1000 times and compare the empirical density function to the probability density function.

Related Distributions

If X has the standard triangle distribution with parameter p , then $1 - X$ has the standard triangle distribution with parameter $1 - p$.

Proof

For $x \in [0, 1]$, $\mathbb{P}(1 - X \leq x) = \mathbb{P}(X \geq 1 - x) = 1 - G(1 - x)$, where G is the CDF of X . The result now follows from the formula for the CDF.

The standard triangle distribution has a number of connections with the standard uniform distribution. Recall that a simulation of a random variable with a standard uniform distribution is a *random number* in computer science.

Suppose that U_1 and U_2 are independent random variables, each with the standard uniform distribution. Then

1. $X = \min\{U_1, U_2\}$ has the standard triangle distribution with $p = 0$.
2. $Y = \max\{U_1, U_2\}$ has the standard triangle distribution with $p = 1$.

Proof

U_1 and U_2 have CDF $u \mapsto u$ for $u \in [0, 1]$

1. X has CDF $x \mapsto 1 - (1 - x)^2$ for $x \in [0, 1]$
2. Y has CDF $y \mapsto y^2$ for $y \in [0, 1]$.

Suppose again that U_1 and U_2 are independent random variables, each with the standard uniform distribution. Then

1. $X = |U_2 - U_1|$ has the standard triangle distribution with $p = 0$.
2. $Y = (U_1 + U_2) / 2$ has the standard triangle distribution with $p = \frac{1}{2}$.

Proof

1. Let $x \in [0, 1]$. Note that the event $\{X > x\} = \{|U_2 - U_1| > x\}$ is simply the union of two disjoint triangular regions, each with base and height of length $1 - x$. Hence $\mathbb{P}(X \leq x) = 1 - (1 - x)^2$.
2. Let $y \in [0, \frac{1}{2}]$. The event $\{Y \leq y\} = \{U_1 + U_2 \leq 2y\}$ is a triangular region with height and base of length $2y$. Hence $\mathbb{P}(Y \leq y) = 2y^2$. For $y \in [\frac{1}{2}, 1]$, the event $\{Y > y\}$ is a triangular region with height and base of length $2 - 2y$. Hence $\mathbb{P}(Y \leq y) = 1 - 2(1 - y)^2$.

In the previous result, note that Y is the sample mean from a random sample of size 2 from the standard uniform distribution. Since the [quantile function](#) has a simple closed-form expression, the standard triangle distribution can be simulated using the random quantile method.

Suppose that U has the standard uniform distribution and $p \in [0, 1]$. Then the random variable below has the standard triangle distribution with parameter p :

$$X = \begin{cases} \sqrt{pU}, & U \leq p \\ 1 - \sqrt{(1-p)(1-U)}, & p < U \leq 1 \end{cases} \quad (5.24.6)$$

Open the random quantile experiment and select the triangle distribution. Vary p (but keep the default values for the other parameters) and note the shape of the distribution function/quantile function. For selected values of p , run the experiment 1000 times and watch the random quantiles. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The standard triangle distribution can also be simulated using the rejection method, which also works well since the region R under the probability density function g is bounded. Recall that this method is based on the following fact: if (X, Y) is uniformly distributed on the rectangular region $S = \{(x, y) : 0 \leq x \leq 1, 0 \leq y \leq 2\}$ which contains R , then the conditional distribution of (X, Y) given $(X, Y) \in R$ is uniformly distributed on R , and hence X has probability density function g .

Open the rejection method experiment and select the triangle distribution. Vary p (but keep the default values for the other parameters) and note the shape of the probability density function. For selected values of p , run the experiment 1000 times and

watch the scatterplot. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

For the extreme values of the shape parameter, the standard triangle distributions are also beta distributions.

Connections to the beta distribution:

1. The standard triangle distribution with shape parameter $p = 0$ is the beta distribution with left parameter $a = 1$ and right parameter $b = 2$.
2. The standard triangle distribution with shape parameter $p = 1$ is the beta distribution with left parameter $a = 2$ and right parameter $b = 1$.

Proof

These results follow directly from the form of the [standard triangle PDF](#).

Open the special distribution simulator and select the beta distribution. For parameter values given below, run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

1. $a = 1, b = 2$
2. $a = 2, b = 1$

The General Triangle Distribution

Like so many *standard distributions*, the standard triangle distribution is usually generalized by adding location and scale parameters.

Definition

Suppose that Z has the standard triangle distribution with vertex at $p \in [0, 1]$. For $a \in \mathbb{R}$ and $w \in (0, \infty)$, random variable $X = a + wZ$ has the *triangle distribution* with *location parameter* a , and *scale parameter* w , and *shape parameter* p

Distribution Functions

Suppose that X has the general triangle distribution given in the definition above.

X has probability density function f given as follows:

1. If $p = 0$, $f(x) = \frac{2}{w^2}(a + w - x)$ for $x \in [a, a + w]$.
2. If $p = 1$, $f(x) = \frac{2}{w^2}(x - a)$ for $x \in [a, a + w]$.
3. If $p \in (0, 1)$,

$$f(x) = \begin{cases} \frac{2}{pw^2}(x - a), & x \in [a, a + pw] \\ \frac{2}{w^2(1-p)}(a + w - x), & x \in [a + pw, a + w] \end{cases} \quad (5.24.7)$$

Proof

This follows from a standard result for location-scale families. Recall that

$$f(x) = \frac{1}{w}g\left(\frac{x - a}{w}\right), \quad \frac{x - a}{w} \in [0, 1] \quad (5.24.8)$$

where g is the [standard triangle PDF](#) with parameter p .

Once again, the shape of the probability density function justifies the name *triangle distribution*.

The graph of f , together with the domain $[a, a + w]$, forms a triangle with vertices $(a, 0)$, $(a + w, 0)$, and $(a + pw, 2/w)$. The mode of the distribution is $x = a + pw$.

1. If $p = 0$, f is decreasing.
2. If $p = 1$, f is increasing.
3. If $p \in (0, 1)$, f increases and then decreases.

Clearly the general triangle distribution could be parameterized by the left endpoint a , the right endpoint $b = a + w$ and the location of the vertex $c = a + pw$, but the location-scale-shape parameterization is better.

Open special distribution simulator and select the triangle distribution. Vary the parameters a , w , and p , and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function F of X is given as follows:

1. If $p = 0$, $F(x) = 1 - \frac{1}{w^2}(a + w - x)^2$ for $x \in [a, a + w]$
2. If $p = 1$, $F(x) = \frac{1}{w^2}(x - a)^2$ for $x \in [a, a + w]$
3. If $p \in (0, 1)$,

$$F(x) = \begin{cases} \frac{1}{pw^2}(x - a)^2, & x \in [a, a + pw] \\ 1 - \frac{1}{w^2(1-p)}(a + w - x)^2, & x \in [a + pw, a + w] \end{cases} \quad (5.24.9)$$

Proof

This follows from a standard result for location-scale families:

$$F(x) = G\left(\frac{x - a}{w}\right), \quad x \in [a, a + w] \quad (5.24.10)$$

where G is the [standard triangle CDF](#) with parameter p .

X has quantile function F^{-1} given by

$$F^{-1}(u) = a + \begin{cases} w\sqrt{up}, & 0 \leq u \leq p \\ w[1 - \sqrt{(1-u)(1-p)}], & p \leq u \leq 1 \end{cases} \quad (5.24.11)$$

1. The first quartile is $a + w\sqrt{\frac{1}{4}p}$ if $p \in [\frac{1}{4}, 1]$ and is $a + w\left(1 - \sqrt{\frac{3}{4}(1-p)}\right)$ if $p \in [0, \frac{1}{4}]$
2. The median is $a + w\sqrt{\frac{1}{2}p}$ if $p \in [\frac{1}{2}, 1]$ and is $a + w\left(1 - \sqrt{\frac{1}{2}(1-p)}\right)$ if $p \in [0, \frac{1}{2}]$.
3. The third quartile is $a + w\sqrt{\frac{3}{4}p}$ if $p \in [\frac{3}{4}, 1]$ and is $a + w\left(1 - \sqrt{\frac{1}{4}(1-p)}\right)$ if $p \in [0, \frac{3}{4}]$.

Proof

This follows from a standard result for location-scale families: $F^{-1}(u) = a + wG^{-1}(u)$ for $u \in [0, 1]$, where G^{-1} is the [standard triangle quantile function](#) with parameter p .

Open the special distribution simulator and select the triangle distribution. Vary the parameters a , w , and p , and note the shape and location of the distribution function. For selected values of parameters, compute the median and the first and third quartiles.

Moments

Suppose again that X has the triangle distribution with location parameter $a \in \mathbb{R}$, scale parameter $w \in (0, \infty)$ and shape parameter $p \in [0, 1]$. Then we can take $X = a + wZ$ where Z has the standard triangle distribution with parameter p . Hence the moments of X can be computed from the [moments of \$Z\$](#) . Using the binomial theorem and the linearity of expected value we have

$$\mathbb{E}(X^n) = \sum_{k=0}^n \binom{n}{k} w^k a^{n-k} \mathbb{E}(Z^k), \quad n \in \mathbb{N} \quad (5.24.12)$$

The general results are rather messy.

The mean and variance of X are

1. $\mathbb{E}(X) = a + \frac{w}{3}(1 + p)$
2. $\text{var}(X) = \frac{w^2}{18}[1 - p(1 - p)]$

Proof

This follows from the results for the [mean and variance of the standard triangle distribution](#), and simple properties of expected value and variance.

Open the special distribution simulator and select the triangle distribution. Vary the parameters a , w , and p , and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness of X is

$$\text{skew}(X) = \frac{\sqrt{2}(1-2p)(1+p)(2-p)}{5[1-p(1-p)]^{3/2}} \quad (5.24.13)$$

The kurtosis of X is $\text{kurt}(X) = \frac{12}{5}$.

Proof

These results follow immediately from the [skewness and kurtosis of the standard triangle distribution](#). Recall that skewness and kurtosis are defined in terms of the standard score, which is independent of the location and scale parameters.

As before, the excess kurtosis is $\text{kurt}(X) - 3 = -\frac{3}{5}$.

Related Distributions

Since the triangle distribution is a location-scale family, it's invariant under location-scale transformations. More generally, the family is closed under linear transformations with nonzero slope.

Suppose that X has the triangle distribution with shape parameter $a \in \mathbb{R}$, scale parameter $w \in (0, \infty)$, and shape parameter $p \in [0, 1]$. If $b \in \mathbb{R}$ and $c \in (0, \infty)$ then

1. $b + cX$ has the triangle distribution with location parameter $b + ca$, scale parameter cw , and shape parameter p .
2. $b - cX$ has the triangle distribution with location parameter $b - c(a + w)$, scale parameter cw , and shape parameter $1 - p$.

Proof

From the [definition](#) we can take $X = a + wZ$ where Z has the standard triangle distribution with parameter p .

1. Note that $b + cX = (b + ca) + cwZ$.
2. Note that $b - cX = b - c(a + w) + cw(1 - Z)$, and recall from the [result above](#) that $1 - Z$ has the basic triangle distribution with parameter $1 - p$.

As with the standard distribution, there are several connections between the triangle distribution and the continuous uniform distribution.

Suppose that V_1 and V_2 are independent and are uniformly distributed on the interval $[a, a + w]$, where $a \in \mathbb{R}$ and $w \in (0, \infty)$. Then

1. $\min\{V_1, V_2\}$ has the triangle distribution with location parameter a , scale parameter w , and shape parameter $p = 0$.
2. $\max\{V_1, V_2\}$ has the triangle distribution with location parameter a , scale parameter w , and shape parameter $p = 1$.

Proof

The uniform distribution is itself a location-scale family, so we can write $V_1 = a + wU_1$ and $V_2 = a + wU_2$, where U_1 and U_2 are independent and each has the standard uniform distribution. Then $\min\{V_1, V_2\} = a + w \min\{U_1, U_2\}$ and $\max\{V_1, V_2\} = a + w \max\{U_1, U_2\}$ so the result follows from the [corresponding result](#) for the standard triangle distribution.

Suppose again that V_1 and V_2 are independent and are uniformly distributed on the interval $[a, a + w]$, where $a \in \mathbb{R}$ and $w \in (0, \infty)$. Then

1. $|V_2 - V_1|$ has the triangle distribution with location parameter 0, scale parameter w , and shape parameter $p = 0$.
2. $V_1 + V_2$ has the triangle distribution with location parameter $2a$, scale parameter $2w$, and shape parameter $p = \frac{1}{2}$.

3. $V_2 - V_1$ has the triangle distribution with location parameter $-w$, scale parameter $2w$, and shape parameter $p = \frac{1}{2}$

Proof

As before, we can write $V_1 = a + wU_1$ and $V_2 = a + wU_2$, where U_1 and U_2 are independent and each has the standard uniform distribution.

1. $|V_2 - V_1| = w|U_2 - U_1|$ and by the [result above](#), $|U_2 - U_1|$ has the standard triangle distribution with parameter $p = 0$.
2. $V_1 + V_2 = 2a + 2w\left[\frac{1}{2}(U_1 + U_2)\right]$ and by the [result above](#), $\frac{1}{2}(U_1 + U_2)$ has the standard triangle distribution with parameter $p = \frac{1}{2}$.
3. Let $Z = \frac{1}{2} + \frac{1}{2}(U_2 - U_1) = \frac{1}{2}U_2 + \frac{1}{2}(1 - U_1)$. Since $1 - U_1$ also has the standard uniform distribution and is independent of U_2 , it follows from the [result above](#) that Z has the basic triangle distribution with parameter $p = \frac{1}{2}$. But $V_2 - V_1 = w(U_2 - U_1) = w(2Z - 1) = 2wZ - w$ and hence the result follows.

A special case of (b) leads to a connection between the triangle distribution and the Irwin-Hall distribution.

Suppose that U_1 and U_2 are independent random variables, each with the standard uniform distribution. Then $U_1 + U_2$ has the triangle distribution with location parameter 0, scale parameter 2 and shape parameter $\frac{1}{2}$. But this is also the Irwin-Hall distribution of order $n = 2$.

Open the special distribution simulator and select the Irwin-Hall distribution. Set $n = 2$ and note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

Since we can simulate a variable Z with the basic triangle distribution with parameter $p \in [0, 1]$ by the [random quantile method above](#), we can simulate a variable with the triangle distribution that has location parameter $a \in \mathbb{R}$, scale parameter $w \in (0, \infty)$, and shape parameter p by our very [definition](#): $X = a + wZ$. Equivalently, we could compute a random quantile using the [quantile function of \$X\$](#) .

Open the random quantile experiment and select the triangle distribution. Vary the location parameter a , the scale parameter w , and the shape parameter p , and note the shape of the distribution function. For selected values of the parameters, run the experiment 1000 times and watch the random quantiles. Compare the empirical density function, mean and standard deviation to their distributional counterparts.

As with the standard distribution, the general triangle distribution has a bounded probability density function on a bounded interval, and hence can be simulated easily via the rejection method.

Open the rejection method experiment and select the triangle distribution. Vary the parameters and note the shape of the probability density function. For selected values of the parameters, run the experiment 1000 times and watch the scatterplot. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

This page titled [5.24: The Triangle Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.25: The Irwin-Hall Distribution

The *Irwin-Hall distribution*, named for Joseph Irwin and Phillip Hall, is the distribution that governs the sum of independent random variables, each with the standard uniform distribution. It is also known as the *uniform sum distribution*. Since the standard uniform is one of the simplest and most basic distributions (and corresponds in computer science to a *random number*), the Irwin-Hall is a natural family of distributions. It also serves as a nice example of the central limit theorem, conceptually easy to understand.

Basic Theory

Definition

Suppose that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent random variables, each with the uniform distribution on the interval $[0, 1]$ (the *standard uniform distribution*). For $n \in \mathbb{N}_+$, let

$$X_n = \sum_{i=1}^n U_i \quad (5.25.1)$$

Then X_n has the *Irwin-Hall distribution* of order n .

So X_n has a continuous distribution on the interval $[0, n]$ for $n \in \mathbb{N}_+$.

Distribution Functions

Let f denote the probability density function of the standard uniform distribution, so that $f(x) = 1$ for $0 \leq x \leq 1$ (and is 0 otherwise). It follows immediately that the probability density function f_n of X_n satisfies $f_n = f^{*n}$, where of course f^{*n} is the n -fold convolution power of f . We can compute f_2 and f_3 by hand.

The probability density function f_2 of X_2 is given by

$$f_2(x) = \begin{cases} x, & 0 \leq x \leq 1 \\ x - 2(x - 1), & 1 \leq x \leq 2 \end{cases} \quad (5.25.2)$$

Proof

Note that X_2 takes values in $[0, 2]$ and $f_2(x) = \int_{\mathbb{R}} f(u)f(x-u) du$ for $x \in [0, 2]$. The integral reduces to $\int_0^x 1 du = x$ for $0 \leq x \leq 1$ and the integral reduces to $\int_{x-1}^1 1 du = 2 - x$ for $1 \leq x \leq 2$.

Note that the graph of f_2 on $[0, 2]$ consists of two lines, pieced together in a continuous way at $x = 1$. The form given above is not the simplest, but makes the continuity clear, and will be helpful when we generalize.

In the special distribution simulator, select the Irwin-Hall distribution and set $n = 2$. Note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The probability density function f_3 of X_3 is given by

$$f_3(x) = \begin{cases} \frac{1}{2}x^2, & 0 \leq x \leq 1 \\ \frac{1}{2}x^2 - \frac{3}{2}(x-1)^2, & 1 \leq x \leq 2 \\ \frac{1}{2}x^2 - \frac{3}{2}(x-1)^2 + \frac{3}{2}(x-2)^2, & 2 \leq x \leq 3 \end{cases} \quad (5.25.3)$$

Note that the graph of f_3 on $[0, 3]$ consists of three parabolas pieced together in a continuous way at $x = 1$ and $x = 2$. The expressions for $f_3(x)$ for $1 \leq x \leq 2$ and $2 \leq x \leq 3$ can be expanded and simplified, but again the form given above makes the continuity clear, and will be helpful when we generalize.

In the special distribution simulator, select the Irwin-Hall distribution and set $n = 3$. Note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

Naturally, we don't want to perform the convolutions one at a time; we would like a general formula. To state the formula succinctly, we need to recall the *floor function*:

$$\lfloor x \rfloor = \max\{n \in \mathbb{Z} : n \leq x\}, \quad x \in \mathbb{R} \quad (5.25.4)$$

so that $\lfloor x \rfloor = j$ if $j \in \mathbb{Z}$ and $j \leq x < j+1$.

For $n \in \mathbb{N}_+$, the probability density function f_n of X_n is given by

$$f_n(x) = \frac{1}{(n-1)!} \sum_{k=0}^{\lfloor x \rfloor} (-1)^k \binom{n}{k} (x-k)^{n-1}, \quad x \in \mathbb{R} \quad (5.25.5)$$

Proof

Let f_n denote the function given by the formula above. Clearly X_n takes values in $[0, n]$, so first let's note that f_n gives the correct value outside of this interval. If $x < 0$, the sum is over an empty index set and hence is 0. Suppose $x > n$. Since $\binom{n}{k} = 0$ for $k > n$, we have

$$f_n(x) = \frac{1}{(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} (x-k)^{n-1}, \quad x \in \mathbb{R} \quad (5.25.6)$$

Using the binomial theorem,

$$\begin{aligned} \sum_{k=0}^n (-1)^k \binom{n}{k} (x-k)^{n-1} &= \sum_{k=0}^n (-1)^k \binom{n}{k} \sum_{j=0}^{n-1} \binom{n-1}{j} x^j (-k)^{n-1-j} \\ &= \sum_{j=0}^{n-1} (-1)^{n-1-j} \binom{n-1}{j} x^j \sum_{k=0}^n (-1)^k \binom{n}{k} k^{n-1-j} \end{aligned}$$

The second sum in the last expression is 0 for $j \in \{0, 1, \dots, n-1\}$ by the alternating series identity for binomial coefficients. We will see this identity again.

To show that the formula is correct on $[0, n]$ we use induction on n . Suppose that $n = 1$. If $0 < x < 1$, then $\lfloor x \rfloor = 0$ so

$$f_1(x) = \frac{1}{0!} (-1)^0 \binom{1}{0} x^0 = 1 = f(x) \quad (5.25.7)$$

Suppose now that the formula is correct for a given $n \in \mathbb{N}_+$. We need to show that $f_n * f = f_{n+1}$. Note that

$$(f_n * f)(x) = \int_{\mathbb{R}} f_n(y) f(x-y) dy = \int_{x-1}^x f_n(y) dy \quad (5.25.8)$$

As often with convolutions, we must take cases. Suppose that $j \leq x < j+1$ where $j \in \{0, 1, \dots, n\}$. Then

$$(f_n * f)(x) = \int_{x-1}^x f_n(y) dy = \int_{x-1}^j f_n(y) dy + \int_j^x f_n(y) dy \quad (5.25.9)$$

Substituting the formula for $f_n(y)$ and integrating gives

$$\begin{aligned} \int_{x-1}^j f_n(y) dy &= \frac{1}{n!} \sum_{k=0}^{j-1} (-1)^k \binom{n}{k} (j-k)^n - \frac{1}{n!} \sum_{k=0}^{j-1} (-1)^k \binom{n}{k} (x-1-k)^n \\ \int_j^x f_n(y) dy &= \frac{1}{n!} \sum_{k=0}^j (-1)^k \binom{n}{k} (x-k)^n - \frac{1}{n!} \sum_{k=0}^j (-1)^k \binom{n}{k} (j-k)^n \end{aligned}$$

Adding these together, note that the first sum in the first equation cancels the second sum in the second equation. Re-indexing the second sum in the first equation we have

$$(f_n * f)(x) = \frac{1}{n!} \sum_{k=1}^j (-1)^k \binom{n}{k-1} (x-k)^n + \frac{1}{n!} \sum_{k=0}^n (-1)^k \binom{n}{k} (x-k)^n \quad (5.25.10)$$

Finally, using the famous binomial identity $\binom{n}{k-1} + \binom{n}{k} = \binom{n+1}{k}$ for $k \in \{1, 2, \dots, n\}$ we have

$$(f_n * f)(x) = \frac{1}{n!} \sum_{k=0}^j (-1)^k \binom{n+1}{k} (x-k)^n = f_{n+1}(x) \quad (5.25.11)$$

Note that for $n \in \mathbb{N}_+$, the graph of f_n on $[0, n]$ consists of n polynomials of degree $n-1$ pieced together in a continuous way. Such a construction is known as a *polynomial spline*. The points where the polynomials are connected are known as *knots*. So f_n is a polynomial spline of degree $n-1$ with knots at $x \in \{1, 2, \dots, n-1\}$. There is another representation of f_n as a sum. To state this one succinctly, we need to recall the *sign function*:

$$\text{sgn}(x) = \begin{cases} -1, & x < 0 \\ 0, & x = 0 \\ 1, & x > 0 \end{cases} \quad (5.25.12)$$

For $n \in \mathbb{N}_+$, the probability density function f_n of X_n is given by

$$f_n(x) = \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} \text{sgn}(x-k)(x-k)^{n-1}, \quad x \in \mathbb{R} \quad (5.25.13)$$

Direct Proof

Let g_n denote the function defined in the theorem. We will show directly that $g_n = f_n$, the probability density function given in the [previous theorem](#). Suppose that $j \leq x < j+1$, so that $\lfloor x \rfloor = j$. Note that $\text{sgn}(x-k) = 1$ for $k < j$ and $\text{sgn}(x-k) = -1$ for $k > j$. Hence

$$g_n(x) = \frac{1}{2(n-1)!} \sum_{k=0}^j (-1)^k \binom{n}{k} (x-k)^{n-1} - \frac{1}{2(n-1)!} \sum_{k=j+1}^n (-1)^k \binom{n}{k} (x-k)^{n-1} \quad (5.25.14)$$

Adding and subtracting a copy of the first term gives

$$\begin{aligned} g_n(x) &= \frac{1}{(n-1)!} \sum_{k=0}^j (-1)^k \binom{n}{k} (x-k)^{n-1} - \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} (x-k)^{n-1} \\ &= f_n(x) - \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} (x-k)^{n-1} \end{aligned}$$

The last sum is identically 0, from the proof of the previous theorem.

Proof by induction

For $n = 1$ the displayed formula is

$$\frac{1}{2} [\text{sgn}(x)x^0 - \text{sgn}(x-1)(x-1)^0] = \frac{1}{2} [\text{sgn}(x) - \text{sgn}(x-1)] = \begin{cases} 1, & 0 < x < 1 \\ 0, & \text{otherwise} \end{cases} \quad (5.25.15)$$

So the formula is correct for $n = 1$. Assume now that the formula is correct for $n \in \mathbb{N}_+$. Then

$$f_{n+1}(x) = (f_n * f)(x) = \int_{\mathbb{R}} \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} \text{sgn}(u-k)(u-k)^{n-1} f(x-u) du \quad (5.25.16)$$

$$= \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} \int_{x-1}^x \text{sgn}(u-k)(u-k)^{n-1} du \quad (5.25.17)$$

But $\int_{x-1}^x \text{sgn}(u-k)(u-k)^{n-1} du = \frac{1}{n} [\text{sgn}(x-k)(x-k)^n - \text{sgn}(x-k-1)(x-k-1)^n]$ for $k \in \{0, 1, \dots, n\}$. So substituting and re-indexing one of the sums gives

$$f_{n+1}(x) = \frac{1}{2n!} \sum_{k=0}^n (-1)^k \binom{n}{k} \text{sgn}(x-k)(x-k)^n + \frac{1}{2n!} \sum_{k=1}^{n+1} (-1)^k \binom{n}{k-1} \text{sgn}(x-k)(x-k)^n \quad (5.25.18)$$

Using the famous identity $\binom{n}{k} + \binom{n}{k-1} = \binom{n+1}{k}$ for $k \in \{1, 2, \dots, n\}$ we finally get

$$f_{n+1}(x) = \frac{1}{2n!} \sum_{k=0}^{n+1} (-1)^k \binom{n+1}{k} \operatorname{sgn}(x-k)(x-k)^n \quad (5.25.19)$$

which verifies the formula for $n+1$.

Open the special distribution simulator and select the Irwin-Hall distribution. Start with $n=1$ and increase n successively to the maximum $n=10$. Note the shape of the probability density function. For various values of n , run the simulation 1000 times and compare the empirical density function to the probability density function.

For $n \in \{2, 3, \dots\}$, the Irwin-Hall distribution is symmetric and unimodal, with mode at $n/2$.

The distribution function F_n of X_n is given by

$$F_n(x) = \frac{1}{n!} \sum_{k=0}^{\lfloor x \rfloor} (-1)^k \binom{n}{k} (x-k)^n, \quad x \in [0, n] \quad (5.25.20)$$

Proof

This follows from the [first form of the PDF](#) and integration.

So F_n is a polynomial spline of degree n with knots at $\{1, 2, \dots, n-1\}$. The [alternate form of the probability density function](#) leads to an alternate form of the distribution function.

The distribution function F_n of X_n is given by

$$F_n(x) = \frac{1}{2} + \frac{1}{2n!} \sum_{k=0}^n (-1)^k \binom{n}{k} \operatorname{sgn}(x-k)(x-k)^n, \quad x \in [0, n] \quad (5.25.21)$$

Proof

The result follows from the [second form of the PDF](#) and integration.

The quantile function F_n^{-1} does not have a simple representation, but of course by symmetry, the median is $n/2$.

Open the special distribution calculator and select the Irwin-Hall distribution. Vary n from 1 to 10 and note the shape of the distribution function. For each value of n compute the first and third quartiles.

Moments

The moments of the Irwin-Hall distribution are easy to obtain from the representation as a sum of independent standard uniform variables. Once again, we assume that X_n has the Irwin-Hall distribution of order $n \in \mathbb{N}_+$.

The mean and variance of X_n are

1. $\mathbb{E}(X_n) = n/2$
2. $\operatorname{var}(X_n) = n/12$

Proof

This follows immediately from the representation $X_n = \sum_{i=1}^n U_i$ where $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent, standard uniform variables, since $\mathbb{E}(U_i) = 1/2$ and $\operatorname{var}(U_i) = 1/12$

Open the special distribution simulator and select the Irwin-Hall distribution. Vary n and note the shape and location of the mean \pm standard deviation bar. For selected values of n run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X_n are

1. $\operatorname{skew}(X_n) = 0$

$$2. \text{kurt}(X_n) = 3 - \frac{6}{5n}$$

Proof

The fact that the skewness is 0 follows immediately from the symmetry of the distribution (once we know that X_n has moments of all orders). The kurtosis result follows from the usual formula and the moments of the standard uniform distribution.

Note that $\text{kurt}(X_n) \rightarrow 3$, the kurtosis of the normal distribution, as $n \rightarrow \infty$. That is, the *excess kurtosis* $\text{kurt}(X_n) - 3 \rightarrow 0$ as $n \rightarrow \infty$.

Open the special distribution simulator and select the Irwin-Hall distribution. Vary n and note the shape and of the probability density function in light of the previous results on skewness and kurtosis. For selected values of n run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The moment generating function M_n of X_n is given by $M_n(0) = 1$ and

$$M_n(t) = \left(\frac{e^t - 1}{t} \right)^n, \quad t \in \mathbb{R} \setminus \{0\} \quad (5.25.22)$$

Proof

This follows immediately from the representation $X_n = \sum_{i=1}^n U_i$ where $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent standard uniform variables. Recall that the standard uniform distribution has MGF $t \mapsto (e^t - 1)/t$, and the MGF of a sum of independent variables is the product of the MGFs.

Related Distributions

The most important connection is to the standard uniform distribution in the [definition](#): The Irwin-Hall distribution of order $n \in \mathbb{N}_+$ is the distribution of the sum of n independent variables, each with the standard uniform distribution. The Irwin-Hall distribution of order 2 is also a triangle distribution:

The Irwin-Hall distribution of order 2 is the triangle distribution with location parameter 0, scale parameter 2, and shape parameter $\frac{1}{2}$.

Proof

This follows immediately from the [PDF](#) f_2 .

The Irwin-Hall distribution is connected to the normal distribution via the central limit theorem.

Suppose that X_n has the Irwin-Hall distribution of order n for each $n \in \mathbb{N}_+$. Then the distribution of

$$Z_n = \frac{X_n - n/2}{\sqrt{n/12}} \quad (5.25.23)$$

converges to the standard normal distribution as $n \rightarrow \infty$.

Proof

This follows immediately from the central limit theorem, since $X_n = \sum_{i=1}^n U_i$ where (U_1, U_2, \dots) is a sequence of independent variables, each with the standard uniform distribution. Note that Z_n is the standard score of X_n .

Thus, if n is large, X_n has approximately a normal distribution with mean $n/2$ and variance $n/12$.

Open the special distribution simulator and select the Irwin-Hall distribution. Start with $n = 1$ and increase n successively to the maximum $n = 10$. Note how the probability density function becomes more “normal” as n increases. For various values of n , run the simulation 1000 times and compare the empirical density function to the probability density function.

The Irwin-Hall distribution of order n is trivial to simulate, as the sum of n random numbers. Since the probability density function is bounded on a bounded support interval, the distribution can also be simulated via the rejection method. Computationally, this is a

dumb thing to do, of course, but it can still be a fun exercise.

Open the rejection method experiment and select the Irwin-Hall distribution. For various values of n , run the simulation 2000 times. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

This page titled [5.25: The Irwin-Hall Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.26: The U-Power Distribution

The *U-power distribution* is a U-shaped family of distributions based on a simple family of power functions.

The Standard U-Power Distribution

Distribution Functions

The *standard U-power distribution* with shape parameter $k \in \mathbb{N}_+$ is a continuous distribution on $[-1, 1]$ with probability density function g given by

$$g(x) = \frac{2k+1}{2} x^{2k}, \quad x \in [-1, 1] \quad (5.26.1)$$

Proof

From simple calculus, g is a probability density function:

$$\int_{-1}^1 x^{2k} dx = \frac{2}{2k+1} \quad (5.26.2)$$

The algebraic form of the probability density function explains the name of the distribution. The most common of the standard U-power distributions is the *U-quadratic distribution*, which corresponds to $k = 1$.

The standard U-power probability density function g satisfies the following properties:

1. g is symmetric about $x = 0$.
2. g decreases and then increases with minimum value at $x = 0$.
3. The modes are $x = \pm 1$.
4. g is concave upward.

Proof

Again, these properties follow from basic calculus since

$$g'(x) = \frac{1}{2}(2k+1)(2k)x^{2k-1}, \quad x \in [-1, 1] \quad (5.26.3)$$

$$g''(x) = \frac{1}{2}(2k+1)(2k)(2k-1)x^{2k-2}, \quad x \in [-1, 1] \quad (5.26.4)$$

Open the Special Distribution Simulator and select the U-power distribution. Vary the shape parameter but keep the default values for the other parameters. Note the graph of the probability density function. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function G given by

$$G(x) = \frac{1}{2}(1 + x^{2k+1}), \quad x \in [-1, 1] \quad (5.26.5)$$

Proof

This follows from the [PDF above](#) and simple calculus.

The quantile function G^{-1} given by $G^{-1}(p) = (2p - 1)^{1/(2k+1)}$ for $p \in [0, 1]$.

1. $G^{-1}(1 - p) = -G^{-1}(p)$ for $p \in [0, 1]$.
2. The first quartile is $q_1 = -\frac{1}{2^{1/(2k+1)}}$.
3. The median is 0.
4. The third quartile is $q_3 = \frac{1}{2^{1/(2k+1)}}$.

Proof

The formula for the quantile function follows immediately from the [CDF above](#) by solving $p = G(x)$ for x in terms of $p \in [0, 1]$. Property (a) follows from the symmetry of the distribution about 0.

Open the Special Distribution Calculator and select the U-power distribution. Vary the shape parameter but keep the default values for the other parameters. Note the shape of the distribution function. For various values of the shape parameter, compute a few quantiles.

Moments

Suppose that Z has the standard U-power distribution with parameter $k \in \mathbb{N}_+$. The moments (about 0) are easy to compute.

Let $n \in \mathbb{N}$. The moment of order $2n + 1$ is $\mathbb{E}(Z^{2n+1}) = 0$. The moment of order $2n$ is

$$\mathbb{E}(Z^{2n}) = \frac{2k+1}{2(n+k)+1} \quad (5.26.6)$$

Proof

This result follows from simple calculus. The fact that the even order moments are 0 also follows from the symmetry of the distribution about 0.

Since the mean is 0, the moments about 0 are also the central moments.

The mean and variance of Z are

1. $\mathbb{E}(Z) = 0$
2. $\text{var}(Z) = \frac{2k+1}{2k+3}$

Proof

These results follow from the previous [general moment result](#).

Note that $\text{var}(Z) \rightarrow 1$ as $k \rightarrow \infty$.

Open the Special Distribution Simulator and select the U-power distribution. Vary the shape parameter but keep the default values for the other parameters. Note the position and size of the mean \pm standard deviation bar. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = \frac{(2k+3)^2}{(2k+5)(2k+1)}$

Proof

The skewness is 0 by the symmetry of the distribution. Since the mean is 0, the kurtosis is $\mathbb{E}(Z^4)/[\mathbb{E}(Z^2)]^2$ and so the result follows from the [general moment result](#) above

Note that $\text{kurt}(Z) \rightarrow 1$ as $k \rightarrow \infty$. The *excess kurtosis* is $\text{kurt}(Z) - 3 = \frac{(2k+3)^2}{(2k+5)(2k+1)} - 3$ and so $\text{kurt}(Z) - 3 \rightarrow -2$ as $k \rightarrow \infty$.

Related Distributions

The [U-power probability density function](#) g actually makes sense for $k = 0$ as well, and in this case the distribution reduces to the uniform distribution on the interval $[-1, 1]$. But of course, this distribution is not U-shaped, except in a degenerate sense. There are other connections to the uniform distribution. The first is a standard result since the [U-power quantile function](#) has a simple, closed representation:

Suppose that $k \in \mathbb{N}_+$.

1. If U has the standard uniform distribution then $Z = (2U - 1)^{1/(2k+1)}$ has the standard U-power distribution with parameter k .
2. If Z has the standard U-power distribution with parameter k then $U = \frac{1}{2}(1 + Z^{2k+1})$ has the standard uniform distribution.

Part (a) of course leads to the random quantile method of simulation.

Open the random quantile simulator and select the U-power distribution. Vary the shape parameter but keep the default values for the other parameters. Note the shape of the distribution and density functions. For selected values of the parameter, run the simulation 1000 times and note the random quantiles. Compare the empirical density function to the probability density function.

The standard U-power distribution with shape parameter $k \in \mathbb{N}_+$ converges to the discrete uniform distribution on $\{-1, 1\}$ as $k \rightarrow \infty$.

Proof

This follows from the definition of convergence in distribution. The U-power distribution function G is 0 on $(-\infty, -1]$, is 1 on $[1, \infty)$, and is given by the [formula above](#) on $[-1, 1]$. As $k \rightarrow \infty$, $G(x) \rightarrow 0$ for $x \in (-\infty, -1)$, $G(x) \rightarrow \frac{1}{2}$ for $x \in (-1, 1)$, and $G(x) \rightarrow 1$ for $x \in (1, \infty)$. This agrees with the distribution function of the discrete uniform distribution on $\{-1, 1\}$ except at the points of discontinuity ± 1 .

The General U-Power Distribution

Like so many standard distributions, the standard U-power distribution is generalized by adding location and scale parameters.

Definition

Suppose that Z has the standard U-power distribution with shape parameter $k \in \mathbb{N}_+$. If $\mu \in \mathbb{R}$ and $c \in (0, \infty)$ then $X = \mu + cZ$ has the *U-power distribution* with shape parameter k , location parameter μ and scale parameter c .

Note that X has a continuous distribution on the interval $[a, b]$ where $a = \mu - c$ and $b = \mu + c$, so the distribution can also be parameterized by the shape parameter k and the endpoints a and b . With this parametrization, the location parameter is $\mu = \frac{a+b}{2}$ and the scale parameter is $c = \frac{b-a}{2}$.

Distribution Functions

Suppose that X has the U-power distribution with shape parameter $k \in \mathbb{N}_+$, location parameter $\mu \in \mathbb{R}$, and scale parameter $c \in (0, \infty)$.

X has probability density function f given by

$$f(x) = \frac{2k+1}{2c} \left(\frac{x-\mu}{c} \right)^{2k}, \quad x \in [\mu - c, \mu + c] \quad (5.26.7)$$

1. f is symmetric about μ .
2. f decreases and then increases with minimum value at $x = \mu$.
3. The modes are at $x = \mu \pm c$.
4. f is concave upward.

Proof

Recall that $f(x) = \frac{1}{c} g\left(\frac{x-\mu}{c}\right)$ where g is the [PDF of \$Z\$](#) .

Open the Special Distribution Simulator and select the U-power distribution. Vary the parameters and note the shape and location of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has distribution function F given by

$$F(x) = \frac{1}{2} \left[1 + \left(\frac{x - \mu}{c} \right)^{2k+1} \right], \quad x \in [\mu - c, \mu + c] \quad (5.26.8)$$

Proof

Recall that $F(x) = G\left(\frac{x - \mu}{c}\right)$ where G is the [CDF of \$Z\$](#) .

X has quantile function F^{-1} given by $F^{-1}(p) = \mu + c(2p - 1)^{1/(2k+1)}$ for $p \in [0, 1]$.

1. $F^{-1}(1 - p) = \mu - cF^{-1}(p)$
2. The first quartile is $q_1 = \mu - c \frac{1}{2^{1/(2k+1)}}$
3. The median is μ .
4. The third quartile is $q_3 = \mu + c \frac{1}{2^{1/(2k+1)}}$

Proof

Recall that $F^{-1}(p) = \mu + cG^{-1}(p)$ where G^{-1} is the [quantile function of \$Z\$](#) .

Open the Special Distribution Calculator and select the U-power distribution. Vary the parameters and note the graph of the distribution function. For various values of the parameters, compute selected values of the distribution function and the quantile function.

Moments

Suppose again that X has the U-power distribution with shape parameter $k \in \mathbb{N}_+$, location parameter $\mu \in \mathbb{R}$, and scale parameter $c \in (0, \infty)$.

The mean and variance of X are

1. $\mathbb{E}(X) = \mu$
2. $\text{var}(X) = c^2 \frac{2k+1}{2k+3}$

Proof

These results follow from the representation $X = \mu + cZ$ where Z has the standard U-power distribution with shape parameter k , and from the [mean and variance of \$Z\$](#) .

Note that $\text{var}(Z) \rightarrow c^2$ as $k \rightarrow \infty$

Open the Special Distribution Simulator and select the U-power distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The moments about 0 are messy, but the central moments are simple.

Let $n \in \mathbb{N}_+$. The central moment of order $2n + 1$ is $\mathbb{E}[(X - \mu)^{2n+1}] = 0$. The moment of order $2n$ is

$$\mathbb{E}[(x - \mu)^{2n}] = c^{2n} \frac{2k + 1}{2(n + k) + 1} \quad (5.26.9)$$

Proof

This follows from the representation $X = \mu + cZ$ where Z has the standard U-power distribution with shape parameter k , and the [central moments of \$Z\$](#) .

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$

$$2. \text{kurt}(X) = \frac{(2k+3)^2}{(2k+5)(2k+1)}$$

Proof

Recall that the skewness and kurtosis are defined in terms of the standard score of X and hence are invariant under a location-scale transformation. Thus, the results are the same as for the [standard distribution](#).

Again, $\text{kurt}(X) \rightarrow 1$ as $k \rightarrow \infty$ and the excess kurtosis is $\text{kurt}(X) - 3 = \frac{(2k+3)^2}{(2k+5)(2k+1)} - 3$

Related Distributions

Since the U-power distribution with a given shape parameter is a location-scale family, it is trivially closed under location-scale transformations.

Suppose that X has the U-power distribution with shape parameter $k \in \mathbb{N}_+$, location parameter $\mu \in \mathbb{R}$, and scale parameter $c \in (0, \infty)$. If $\alpha \in \mathbb{R}$ and $\beta \in (0, \infty)$, then $Y = \alpha + \beta X$ has the U-power distribution with shape parameter k , location parameter $\alpha + \beta\mu$, and scale parameter βc .

Proof

From the [definition](#), we can take $X = \mu + cZ$ where Z has the standard U-power distribution with shape parameter k . Then $Y = \alpha + \beta X = (\alpha + \beta\mu) + (\beta c)Z$.

As before, since the [U-power distribution function](#) and the [U-power quantile function](#) have simple forms, we have the usual connections with the standard uniform distribution.

Suppose that $k \in \mathbb{N}_+$, $\mu \in \mathbb{R}$ and $c \in (0, \infty)$.

1. If U has the standard uniform distribution then $X = \mu + c(2U - 1)^{1/(2k+1)}$ has the U-power distribution with shape parameter k , location parameter μ , and scale parameter c .
2. If X has the U-power distribution with shape parameter k , location parameter μ , and scale parameter c , then $U = \frac{1}{2} \left[1 + \left(\frac{X - \mu}{c} \right)^{2k+1} \right]$ has the standard uniform distribution.

Again, part (a) of course leads to the random quantile method of simulation.

Open the random quantile simulator and select the U-power distribution. Vary the parameters and note the shape of the distribution and density functions. For selected values of the parameters, run the simulation 1000 times and note the random quantiles. Compare the empirical density function to the probability density function.

The U-power distribution with given location and scale parameters converges to the discrete uniform distribution at the endpoints as the shape parameter increases.

The U-power distribution with shape parameter $k \in \mathbb{N}_+$, location parameter $\mu \in \mathbb{R}$, and scale parameter $c \in (0, \infty)$ converges to the discrete uniform distribution on $\{\mu - c, \mu + c\}$ as $k \rightarrow \infty$.

Proof

This follows from the [convergence result for the standard distribution](#) and basic properties of convergence in distribution.

The U-power distribution is a general exponential family in the shape parameter, if the location and scale parameters are fixed.

Suppose that X has the U-power distribution with unspecified shape parameter $k \in \mathbb{N}_+$, but with specified location parameter $\mu \in \mathbb{R}$ and scale parameter $c \in (0, \infty)$. Then X has a one-parameter exponential distribution with natural parameter $2k$ and natural statistics $\ln\left(\frac{X - \mu}{c}\right)$.

Proof

This follows from the definition of the general exponential family, since the PDF of the U-power distribution can be written as

$$f(x) = \frac{2k+1}{2c} \exp \left[2k \ln \left(\frac{x-\mu}{c} \right) \right], \quad x \in [\mu - c, \mu + c] \quad (5.26.10)$$

Since the U-power distribution has a bounded probability density function on a bounded support interval, it can also be simulated via the rejection method.

Open the rejection method experiment and select the U-power distribution. Vary the parameters and note the shape of the probability density function. For selected values of the parameters, run the experiment 1000 times and watch the scatterplot. Compare the empirical density function, mean, and standard deviation to their distributional counterparts.

This page titled [5.26: The U-Power Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.27: The Sine Distribution

The *sine distribution* is a simple probability distribution based on a portion of the sine curve. It is also known as *Gilbert's sine distribution*, named for the American geologist [Grove Karl \(GK\) Gilbert](#) who used the distribution in 1892 to study craters on the moon.

The Standard Sine Distribution

Distribution Functions

The *standard sine distribution* is a continuous distribution on $[0, 1]$ with probability density function g given by

$$g(z) = \frac{\pi}{2} \sin(\pi z), \quad z \in [0, 1] \quad (5.27.1)$$

1. g is symmetric about $z = \frac{1}{2}$.
2. g increases and then decreases with mode at $z = \frac{1}{2}$.
3. g is concave downward.

Proof

From simple calculus, g is a probability density function: $\sin(\pi x) \geq 0$ for $x \in [0, 1]$ and

$$\int_0^1 \sin(\pi z) dz = \frac{2}{\pi} \quad (5.27.2)$$

The properties follow from basic calculus since

$$g'(z) = \frac{\pi^2}{2} \cos(\pi z), \quad z \in [0, 1] \quad (5.27.3)$$

$$g''(z) = -\frac{\pi^3}{2} \sin(\pi z), \quad z \in [0, 1] \quad (5.27.4)$$

Open the [Special Distribution Simulator](#) and select the sine distribution. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function G is given by $G(z) = \frac{1}{2}[1 - \cos(\pi z)]$ for $z \in [0, 1]$.

Proof

This follows from the [PDF above](#) and simple calculus.

The quantile function G^{-1} is given by $G^{-1}(p) = \frac{1}{\pi} \arccos(1 - 2p)$ for $p \in [0, 1]$.

1. The first quartile is $q_1 = \frac{1}{3}$.
2. The median is $\frac{1}{2}$.
3. The third quartile is $q_3 = \frac{2}{3}$.

Proof

The formula for the quantile function follows immediately from the [CDF above](#) by solving $p = G(z)$ for z in terms of $p \in [0, 1]$.

Open the [Special Distribution Calculator](#) and select the sine distribution. Compute a few quantiles.

Moments

Suppose that Z has the standard sine distribution. The moment generating function can be given in closed form.

The moment generating function m of Z is given by

$$m(t) = \mathbb{E}(e^{tZ}) = \frac{\pi^2(1+e^t)}{2(t^2 + \pi^2)}, \quad t \in \mathbb{R} \quad (5.27.5)$$

Proof

Note first that

$$m(t) = \frac{\pi}{2} \int_0^1 e^{tz} \sin(\pi z) dz \quad (5.27.6)$$

Integrating by parts with $u = e^{tz}$ and $dv = \sin(\pi z)dz$ gives

$$m(t) = \frac{t}{2}(1+e^t) + \frac{t}{2} \int_0^1 e^{tz} \cos(\pi z) dz \quad (5.27.7)$$

Integrating by parts again with $u = e^{tz}$ and $dv = \cos(\pi z)dz$ gives

$$m(t) = \frac{t}{2}(1+e^t) - \frac{t^2}{\pi^2} m(t) \quad (5.27.8)$$

Solving for $m(t)$ gives the result.

The moments of all orders exist, but a general formula is complicated and involves special functions. However, the mean and variance are easy to compute.

The mean and variance of Z are

1. $\mathbb{E}(Z) = 1/2$
2. $\text{var}(Z) = 1/4 - 2/\pi^2$

Proof

1. We know that the mean exists since the PDF is continuous on a bounded interval. By symmetry, the mean must be $1/2$
2. Integration by parts (twice) gives

$$\mathbb{E}(Z^2) = \int_0^1 z^2 \frac{\pi}{2} \sin(\pi z) dz = \frac{1}{2} - \frac{2}{\pi^2} \quad (5.27.9)$$

The variance then follows from the usual computational formula $\text{var}(Z) = \mathbb{E}(Z^2) - [\mathbb{E}(Z)]^2$.

Of course, the mean and variance could also be obtained by differentiating the MGF.

Numerically, $\text{sd}(Z) \approx 0.2176$.

Open the Special Distribution Simulator and select the sine distribution. Note the position and size of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = (384 - 48\pi^2 + \pi^4)/(\pi^2 - 8)^2$

Proof

1. The skewness is 0 by the symmetry of the distribution.
2. The formula for the kurtosis follows from the usual computational formula and the first four moments: $\mathbb{E}(Z) = 1/2$, $\mathbb{E}(Z^2) = 1/2 - 2/\pi^2$, $\mathbb{E}(Z^3) = 1/2 - 3/\pi^2$, $\mathbb{E}(Z^4) = 1/2 + 24/\pi^4 - 6/\pi^2$.

Numerically, $\text{kurt}(Z) \approx 2.1938$.

Related Distributions

Since the [distribution function](#) and the [quantile function](#) have closed form representations, the standard sine distribution has the usual connection to the standard uniform distribution.

1. If U has the standard uniform distribution then $Z = G^{-1}(U) = \frac{1}{\pi} \arccos(1 - 2U)$ has the standard sine distribution.
2. If Z has the standard sine distribution then $U = G(Z) = \frac{1}{2}[1 - \cos(\pi Z)]$ has the standard uniform distribution.

Part (a) of course leads to the random quantile method of simulation.

Open the random quantile simulator and select the sine distribution. Note the shape of the distribution and density functions. Run the simulation 1000 times and note the random quantiles. Compare the empirical density function to the probability density function.

Since the [probability density function](#) is continuous and is defined on a closed, bounded interval, the standard sine distribution can also be simulated using the rejection method.

Open the rejection method app and select the sine distribution. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The General Sine Distribution

As with so many other “standard distributions”, the standard sine distribution is generalized by adding location and scale parameters.

Suppose that Z has the standard sine distribution. For $a \in \mathbb{R}$ and $b \in (0, \infty)$, random variable $X = a + bZ$ has the *sine distribution* with *location parameter* a and *scale parameter* b .

Distribution Functions

Analogies of the results above for the standard sine distribution follow easily from basic properties of the location-scale transformation. Suppose that X has the sine distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. So X has a continuous distribution on the interval $[a, a + b]$.

The probability density function f of X is given by

$$f(x) = \frac{\pi}{2b} \sin\left(\pi \frac{x-a}{b}\right), \quad x \in [a, a+b] \quad (5.27.10)$$

1. f is symmetric about $x = a + b/2$.
2. f increases and then decreases, with mode $x = a + b/2$.
3. f is concave downward.

Proof

Recall that

$$f(x) = \frac{1}{b} g\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.27.11)$$

where g is the [standard PDF](#).

Pure scale transformations ($a = 0$ and $b > 0$) are particularly common, since X often represents a *random angle*. The scale transformation with $b = \pi$ gives the angle in radians. In this case the probability density function is $f(x) = \frac{1}{2} \sin(x)$ for $x \in [0, \pi]$. Since the radian is the standard angle unit, this distribution could also be considered the “standard one”. The scale transformation with $b = 90$ gives the angle in degrees. In this case, the probability density function is $f(x) = \frac{\pi}{180} \sin\left(\frac{\pi}{90}x\right)$ for $x \in [0, 90]$. This was Gilbert's original formulation.

In the special distribution simulator, select the sine distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical

density function to the probability density function.

The distribution function F of X is given by

$$F(x) = \frac{1}{2} \left[1 - \cos \left(\pi \frac{x-a}{b} \right) \right], \quad x \in [a, a+b] \quad (5.27.12)$$

Proof

Recall that

$$F(x) = G \left(\frac{x-a}{b} \right), \quad x \in \mathbb{R} \quad (5.27.13)$$

where G is the [standard CDF](#).

The quantile function F^{-1} of X is given by

$$F^{-1}(p) = a + \frac{b}{\pi} \arccos(1-2p), \quad p \in (0, 1) \quad (5.27.14)$$

1. The first quartile is $a + b/3$.
2. The median is $a + b/2$.
3. The third quartile is $a + 2b/3$

Proof

Recall that $F^{-1}(p) = a + bG^{-1}(p)$ for $p \in (0, 1)$, where G^{-1} is the [standard quantile function](#).

In the special distribution calculator, select the sine distribution. Vary the parameters and note the shape and location of the probability density function and the distribution function. For selected values of the parameters, find the quantiles of order 0.1 and 0.9.

Moments

Suppose again that X has the sine distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$.

The moment generating function M of X is given by

$$M(t) = \frac{\pi^2 (e^{at} + e^{(a+b)t})}{2(b^2 t^2 + \pi^2)}, \quad t \in \mathbb{R} \quad (5.27.15)$$

Proof

Recall that $M(t) = e^{at} m(bt)$ where m is the [standard MGF](#).

The mean and variance of X are

1. $\mathbb{E}(X) = a + b/2$
2. $\text{var}(X) = b^2(1/4 - 2/\pi^2)$

Proof

By [definition](#) we can assume $X = a + bZ$ where Z has the standard sine distribution. Using the [mean and variance of \$Z\$](#) we have

1. $\mathbb{E}(X) = a + b\mathbb{E}(Z) = a + b/2$
2. $\text{var}(X) = b^2 \text{var}(Z) = b^2(1/4 - 2/\pi^2)$

In the special distribution simulator, select the sine distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = (384 - 48\pi^2 + \pi^4)/(\pi^2 - 8)^2$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are invariant under location-scale transformations. So the skewness and kurtosis of X are the same as the [skewness and kurtosis of \$Z\$](#) .

Related Distributions

The general sine distribution is a location-scale family, so it is trivially closed under location-scale transformations.

Suppose that X has the sine distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, and that $c \in \mathbb{R}$ and $d \in (0, \infty)$. Then $Y = c + dX$ has the sine distribution with location parameter $c + ad$ and scale parameter bd .

Proof

Again by [definition](#) we can take $X = a + bZ$ where Z has the standard sine distribution. Then $Y = c + dX = (c + ad) + (bd)Z$.

This page titled [5.27: The Sine Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.28: The Laplace Distribution

The *Laplace distribution*, named for Pierre Simon Laplace arises naturally as the distribution of the difference of two independent, identically distributed exponential variables. For this reason, it is also called the *double exponential distribution*.

The Standard Laplace Distribution

Distribution Functions

The *standard Laplace distribution* is a continuous distribution on \mathbb{R} with probability density function g given by

$$g(u) = \frac{1}{2}e^{-|u|}, \quad u \in \mathbb{R} \quad (5.28.1)$$

Proof

It's easy to see that g is a valid PDF. By symmetry

$$\int_{-\infty}^{\infty} \frac{1}{2}e^{-|u|} du = \int_0^{\infty} e^{-u} du = 1 \quad (5.28.2)$$

The probability density function g satisfies the following properties:

1. g is symmetric about 0.
2. g increases on $(-\infty, 0]$ and decreases on $[0, \infty)$, with mode $u = 0$.
3. g is concave upward on $(-\infty, 0]$ and on $[0, \infty)$ with a cusp at $u = 0$

Proof

These results follow from standard calculus, since $g(u) = \frac{1}{2}e^{-u}$ for $u \in [0, \infty)$ and $g(u) = \frac{1}{2}e^u$ for $u \in (-\infty, 0]$.

Open the Special Distribution Simulator and select the Laplace distribution. Keep the default parameter value and note the shape of the probability density function. Run the simulation 1000 times and compare the empirical density function and the probability density function.

The standard Laplace distribution function G is given by

$$G(u) = \begin{cases} \frac{1}{2}e^u, & u \in (-\infty, 0] \\ 1 - \frac{1}{2}e^{-u}, & u \in [0, \infty) \end{cases} \quad (5.28.3)$$

Proof

Again this follows from basic calculus, since $g(u) = \frac{1}{2}e^u$ for $u \leq 0$ and $g(u) = \frac{1}{2}e^{-u}$ for $u \geq 0$. Of course $G(u) = \int_{-\infty}^u g(t) dt$.

The quantile function G^{-1} given by

$$G^{-1}(p) = \begin{cases} \ln(2p), & p \in [0, \frac{1}{2}] \\ -\ln[2(1-p)], & p \in [\frac{1}{2}, 1] \end{cases} \quad (5.28.4)$$

1. $G^{-1}(1-p) = -G^{-1}(p)$ for $p \in (0, 1)$
2. The first quartile is $q_1 = -\ln 2 \approx -0.6931$.
3. The median is $q_2 = 0$
4. The third quartile is $q_3 = \ln 2 \approx 0.6931$.

Proof

The formula for the quantile function follows immediately from the CDF by solving $p = G(u)$ for u in terms of $p \in (0, 1)$. Part (a) is due to the symmetry of g about 0.

Open the Special Distribution Calculator and select the Laplace distribution. Keep the default parameter value. Compute selected values of the distribution function and the quantile function.

Moments

Suppose that U has the standard Laplace distribution.

U has moment generating function m given by

$$m(t) = \mathbb{E}(e^{tU}) = \frac{1}{1-t^2}, \quad t \in (-1, 1) \quad (5.28.5)$$

Proof

For $t \in (-1, 1)$,

$$m(t) = \int_{-\infty}^{\infty} e^{tu} g(u) du = \int_{-\infty}^0 \frac{1}{2} e^{(t+1)u} du + \int_0^{\infty} \frac{1}{2} e^{(t-1)u} du = \frac{1}{2(t+1)} - \frac{1}{2(t-1)} = \frac{1}{1-t^2} \quad (5.28.6)$$

The moments of U are

1. $\mathbb{E}(U^n) = 0$ if $n \in \mathbb{N}$ is odd.
2. $\mathbb{E}(U^n) = n!$ if $n \in \mathbb{N}$ is even.

Proof

This result can be obtained from the [moment generating function](#) or directly. That the odd order moments are 0 follows from the symmetry of the distribution. For the even order moments, symmetry and an integration by parts (or using the gamma function) gives

$$\mathbb{E}(U^n) = \frac{1}{2} \int_{-\infty}^0 u^n e^u du + \frac{1}{2} \int_0^{\infty} u^n e^{-u} du = \int_0^{\infty} u^n e^{-u} du = n! \quad (5.28.7)$$

The mean and variance of U are

1. $\mathbb{E}(U) = 0$
2. $\text{var}(U) = 2$

Open the [Special Distribution Simulator](#) and select the Laplace distribution. Keep the default parameter value. Run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of U are

1. $\text{skew}(U) = 0$
2. $\text{kurt}(U) = 6$

Proof

1. This follows from the symmetry of the distribution.
2. Since $\mathbb{E}(U) = 0$, we have

$$\text{kurt}(U) = \frac{\mathbb{E}(U^4)}{[\mathbb{E}(U^2)]^2} = \frac{4!}{(2!)^2} = 6 \quad (5.28.8)$$

It follows that the *excess kurtosis* is $\text{kurt}(U) - 3 = 3$.

Related Distributions

Of course, the standard Laplace distribution has simple connections to the standard exponential distribution.

If U has the standard Laplace distribution then $V = |U|$ has the standard exponential distribution.

Proof

Using the [CDF of U](#) we have $\mathbb{P}(V \leq v) = \mathbb{P}(-v \leq U \leq v) = G(v) - G(-v) = 1 - e^{-v}$ for $v \in [0, \infty)$. This function is the CDF of the standard exponential distribution.

If V and W are independent and each has the standard exponential distribution, then $U = V - W$ has the standard Laplace distribution.

Proof using PDFs

Let h denote the standard exponential PDF, extended to all of \mathbb{R} , so that $h(v) = e^{-v}$ if $v \geq 0$ and $h(v) = 0$ if $v < 0$. Using convolution, the PDF of $V - W$ is $g(u) = \int_{\mathbb{R}} h(v)h(v-u)dv$. If $v \geq 0$,

$$g(u) = \int_u^{\infty} e^{-v} e^{-(v-u)} dv = e^u \int_u^{\infty} e^{-2v} dv = \frac{1}{2} e^{-u} \quad (5.28.9)$$

If $u < 0$ then

$$g(u) = \int_0^{\infty} e^{-v} e^{-(v-u)} dv = e^u \int_0^{\infty} e^{-2v} dv = \frac{1}{2} e^u \quad (5.28.10)$$

Proof using MGFs

The MGF of V is $t \mapsto 1/(1-t)$ for $t < 1$. The MGF of $-W$ is $t \mapsto 1/(1+t)$ for $t > -1$. Hence the MGF of U is $t \mapsto 1/(1-t)(1+t) = 1/(1-t^2)$ for $-1 < t < 1$, which is the [standard Laplace MGF](#).

If V has the standard exponential distribution, I has the standard Bernoulli distribution, and V and I are independent, then $U = (2I - 1)V$ has the standard Laplace distribution.

Proof

If $u \geq 0$ then

$$\mathbb{P}(U \leq u) = \mathbb{P}(I = 0) + \mathbb{P}(I = 1, V \leq u) = \mathbb{P}(I = 0) + \mathbb{P}(I = 1)\mathbb{P}(V \leq u) = \frac{1}{2} + \frac{1}{2}(1 - e^{-u}) = 1 - \frac{1}{2}e^{-u} \quad (5.28.11)$$

If $u < 0$,

$$\mathbb{P}(U \leq u) = \mathbb{P}(I = 0, V > -u) = \mathbb{P}(I = 0)\mathbb{P}(V > -u) = \frac{1}{2}e^u \quad (5.28.12)$$

The standard Laplace distribution has a curious connection to the standard normal distribution.

Suppose that (Z_1, Z_2, Z_3, Z_4) is a random sample of size 4 from the standard normal distribution. Then $U = Z_1 Z_2 + Z_3 Z_4$ has the standard Laplace distribution.

Proof

$Z_1 Z_2$ and $Z_3 Z_4$ are independent, and each has a distribution known as the *product normal distribution*. The MGF of this distribution is

$$m_0(t) = \mathbb{E}(e^{tZ_1 Z_2}) = \int_{\mathbb{R}^2} e^{txy} \frac{1}{2\pi} e^{-(x^2+y^2)/2} dx dy \quad (5.28.13)$$

Changing to polar coordinates gives

$$m_0(t) = \frac{1}{2\pi} \int_0^{2\pi} \int_0^{\infty} e^{tr^2 \cos \theta \sin \theta} e^{-r^2/2} r dr d\theta = \frac{1}{2\pi} \int_0^{2\pi} \int_0^{\infty} \exp \left[r^2 \left(t \cos \theta \sin \theta - \frac{1}{2} \right) \right] r dr d\theta \quad (5.28.14)$$

The inside integral can be done with a simple substitution for $|t| < 1$, yielding

$$m_0(t) = \frac{1}{2\pi} \int_0^{2\pi} \frac{1}{1 - t \sin(2\theta)} d\theta = \frac{1}{\sqrt{1-t^2}} \quad (5.28.15)$$

Hence U has MGF $m_0^2(t) = \frac{1}{1-t^2}$ for $|t| < 1$, which again is the [standard Laplace MGF](#).

The standard Laplace distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) computed above.

Connections to the standard uniform distribution.

1. If V has the standard uniform distribution then $U = \ln(2V)\mathbf{1}(V < \frac{1}{2}) - \ln[2(1-V)]\mathbf{1}(V \geq \frac{1}{2})$ has the standard Laplace distribution.
2. If U has the standard Laplace distribution then $V = \frac{1}{2}e^U\mathbf{1}(U < 0) + (1 - \frac{1}{2}e^{-U})\mathbf{1}(U \geq 0)$ has the standard uniform distribution.

From part (a), the standard Laplace distribution can be simulated with the usual random quantile method.

Open the random quantile experiment and select the Laplace distribution. Keep the default parameter values and note the shape of the probability density and distribution functions. Run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The General Laplace Distribution

The standard Laplace distribution is generalized by adding location and scale parameters.

Suppose that U has the standard Laplace distribution. If $a \in \mathbb{R}$ and $b \in (0, \infty)$, then $X = a + bU$ has the *Laplace distribution* with location parameter a and scale parameter b .

Distribution Functions

Suppos that X has the Laplace distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$.

X has probability density function f given by

$$f(x) = \frac{1}{2b} \exp\left(-\frac{|x-a|}{b}\right), \quad x \in \mathbb{R} \quad (5.28.16)$$

1. f is symmetric about a .
2. f increases on $[0, a]$ and decreases on $[a, \infty)$ with mode $x = a$.
3. f is concave upward on $[0, a]$ and on $[a, \infty)$ with a cusp at $x = a$.

Proof

Recall that $f(x) = \frac{1}{b} g\left(\frac{x-a}{b}\right)$ where g is the [standard Laplace PDF](#).

Open the Special Distribution Simulator and select the Laplace distribution. Vary the parameters and note the shape and location of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the emprical density function to the probability density function.

X has distribution function F given by

$$F(x) = \begin{cases} \frac{1}{2} \exp\left(\frac{x-a}{b}\right), & x \in (-\infty, a] \\ 1 - \frac{1}{2} \exp\left(-\frac{x-a}{b}\right), & x \in [a, \infty) \end{cases} \quad (5.28.17)$$

Proof

Recall that $F(x) = G\left(\frac{x-a}{b}\right)$ where G is the [standard Laplace CDF](#).

X has quantile function F^{-1} given by

$$F^{-1}(p) = \begin{cases} a + b \ln(2p), & 0 \leq p \leq \frac{1}{2} \\ a - b \ln[2(1-p)], & \frac{1}{2} \leq p < 1 \end{cases} \quad (5.28.18)$$

1. $F^{-1}(1-p) = a - bF^{-1}(p)$ for $p \in (0, 1)$
2. The first quartile is $q_1 = a - b \ln 2$.
3. The median is $q_2 = a$
4. The third quartile is $q_3 = a + b \ln 2$.

Proof

Recall that $F^{-1}(p) = a + bG^{-1}(p)$ where G^{-1} is the [standard Laplace quantile function](#).

Open the Special Distribution Calculator and select the Laplace distribution. For various values of the scale parameter, compute selected values of the distribution function and the quantile function.

Moments

Again, we assume that X has the Laplace distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, so that by [definition](#), $X = a + bU$ where U has the standard Lapalce distribution.

X has moment generating function M given by

$$M(t) = \mathbb{E}(e^{tX}) = \frac{e^{at}}{1 - b^2 t^2}, \quad t \in (-1/b, 1/b) \quad (5.28.19)$$

Proof

Recall that $M(t) = e^{at}m(bt)$ where m is the [standard Laplace MGF](#).

The moments of X about the location parameter have a simple form.

The moments of X about a are

1. $\mathbb{E}[(X - a)^n] = 0$ if $n \in \mathbb{N}$ is odd.
2. $\mathbb{E}[(X - a)^n] = b^n n!$ if $n \in \mathbb{N}$ is even.

Proof

Note that $\mathbb{E}[(X - a)^n] = b^n \mathbb{E}(U^n)$ so the results follow the [moments of \$U\$](#) .

The mean and variance of X are

1. $\mathbb{E}(X) = a$
2. $\text{var}(X) = 2b^2$

Proof

Recall that $\mathbb{E}(X) = a + b\mathbb{E}(U)$ and $\text{var}(X) = b^2 \text{var}(U)$, so the results follow from the [mean and variance of \$U\$](#) .

Open the Special Distribution Simulator and select the Laplace distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For various values of the scale parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = 6$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are unchanged by a location-scale transformation. Thus the results from the [skewness and kurtosis of \$U\$](#) .

As before, the excess kurtosis is $\text{kurt}(X) - 3 = 3$.

Related Distributions

By construction, the Laplace distribution is a location-scale family, and so is closed under location-scale transformations.

Suppose that X has the Laplace distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, and that $c \in \mathbb{R}$ and $d \in (0, \infty)$. Then $Y = c + dX$ has the Laplace distribution with location parameter $c + ad$ scale parameter bd .

Proof

Again by [definition](#), we can take $X = a + bU$ where U has the standard Laplace distribution. Hence $Y = c + dX = (c + ad) + (bd)U$.

Once again, the Laplace distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) computed above. The latter leads to the usual random quantile method of simulation.

Suppose that $a \in \mathbb{R}$ and $b \in (0, \infty)$.

1. If V has the standard uniform distribution then

$$U = [a + b \ln(2V)] \mathbf{1}\left(V < \frac{1}{2}\right) + (a - b \ln[2(1 - V)]) \mathbf{1}\left(V \geq \frac{1}{2}\right) \quad (5.28.20)$$

has the Laplace distribution with location parameter a and scale parameter b .

2. If X has the Laplace distribution with location parameter a and scale parameter b , then

$$V = \frac{1}{2} \exp\left(\frac{X-a}{b}\right) \mathbf{1}(X < a) + \left[1 - \frac{1}{2} \exp\left(-\frac{X-a}{b}\right)\right] \mathbf{1}(X \geq a) \quad (5.28.21)$$

has the standard uniform distribution.

Open the random quantile experiment and select the Laplace distribution. Vary the parameter values and note the shape of the probability density and distribution functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The Laplace distribution is also a member of the general exponential family of distributions.

Suppose that X has the Laplace distribution with known location parameter $a \in \mathbb{R}$ and unspecified scale parameter $b \in (0, \infty)$. Then X has a general exponential distribution in the scale parameter b , with natural parameter $-1/b$ and natural statistics $|X - a|$.

Proof

This follows from the definition of the general exponential family and the form of the [probability density function](#) f

This page titled [5.28: The Laplace Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.29: The Logistic Distribution

The logistic distribution is used for various growth models, and is used in a certain type of regression, known appropriately as *logistic regression*.

The Standard Logistic Distribution

Distribution Functions

The *standard logistic distribution* is a continuous distribution on \mathbb{R} with distribution function G given by

$$G(z) = \frac{e^z}{1 + e^z}, \quad z \in \mathbb{R} \quad (5.29.1)$$

Proof

Note that G is continuous, and $G(z) \rightarrow 0$ as $z \rightarrow -\infty$ and $G(z) \rightarrow 1$ as $z \rightarrow \infty$. Moreover,

$$G'(z) = \frac{e^z}{(1 + e^z)^2} > 0, \quad z \in \mathbb{R} \quad (5.29.2)$$

so G is increasing.

The probability density function g of the standard logistic distribution is given by

$$g(z) = \frac{e^z}{(1 + e^z)^2}, \quad z \in \mathbb{R} \quad (5.29.3)$$

1. g is symmetric about $x = 0$.
2. g increases and then decreases with the mode $x = 0$.
3. g is concave upward, then downward, then upward again with inflection points at $x = \pm \ln(2 + \sqrt{3}) \approx \pm 1.317$.

Proof

These result follow from standard calculus. First recall that $g = G'$.

1. The symmetry of g is not obvious at first, but note that

$$g(-z) = \frac{e^{-z}}{(1 + e^{-z})^2} \frac{e^{2z}}{e^{2z}} = \frac{e^z}{(1 + e^z)^2} = g(z) \quad (5.29.4)$$

2. The first derivative of g is

$$g'(z) = \frac{e^z(1 - e^z)}{(1 + e^z)^3} \quad (5.29.5)$$

3. The second derivative of g is

$$g''(z) = \frac{e^z(1 - 4e^z + e^{2z})}{(1 + e^z)^4} \quad (5.29.6)$$

In the special distribution simulator, select the logistic distribution. Keep the default parameter values and note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function G^{-1} of the standard logistic distribution is given by

$$G^{-1}(p) = \ln\left(\frac{p}{1-p}\right), \quad p \in (0, 1) \quad (5.29.7)$$

1. The first quartile is $-\ln 3 \approx -1.0986$.

2. The median is 0.
3. The third quartile is $\ln 3 \approx 1.0986$

Proof

The formula for G^{-1} follows by solving $p = G(z)$ for z in terms of p .

Recall that $p : 1 - p$ are the odds in favor of an event with probability p . Thus, the logistic distribution has the interesting property that the quantiles are the logarithms of the corresponding odds ratios. Indeed, this function of p is sometimes called the *logit function*. The fact that the median is 0 also follows from symmetry, of course.

In the special distribution calculator, select the logistic distribution. Keep the default parameter values and note the shape and location of the probability density function and the distribution function. Find the quantiles of order 0.1 and 0.9.

Moments

Suppose that Z has the standard logistic distribution. The moment generating function of Z has a simple representation in terms of the beta function B , and hence also in terms of the gamma function Γ

The moment generating function m of Z is given by

$$m(t) = B(1+t, 1-t) = \Gamma(1+t) \Gamma(1-t), \quad t \in (-1, 1) \quad (5.29.8)$$

Proof

Note that

$$m(t) = \int_{-\infty}^{\infty} e^{tz} \frac{e^z}{(1+e^z)^2} dz \quad (5.29.9)$$

Let $u = \frac{e^z}{1+e^z}$ so that $du = \frac{e^z}{(1+e^z)^2} dz$ and $e^z = \frac{u}{1-u}$. Hence

$$m(t) = \int_0^1 \left(\frac{u}{1-u} \right)^t du = \int_0^1 u^t (1-u)^{-t} du \quad (5.29.10)$$

The last integral, by definition, is $B(1+t, 1-t)$ for $t \in (-1, 1)$

Since the moment generating function is finite on an open interval containing 0, random variable Z has moments of all orders. By symmetry, the odd order moments are 0. The even order moments can be represented in terms of *Bernoulli numbers*, named of course for Jacob Bernoulli. Let β_n Bernoulli number of order $n \in \mathbb{N}$.

Let $n \in \mathbb{N}$

1. If n is odd then $\mathbb{E}(Z^n) = 0$.
2. If n is even then $\mathbb{E}(Z^n) = (2^n - 2)\pi^n |\beta_n|$

Proof

1. Again, this follows from symmetry
2. Recall that the moments of Z can be computed by integrating powers of the quantile function. Hence

$$\mathbb{E}(Z^n) = \int_0^1 \left[\ln \left(\frac{p}{1-p} \right) \right]^n dp \quad (5.29.11)$$

This integral evaluates to the expression above involving the Bernoulli numbers.

In particular, we have the mean and variance.

The mean and variance of Z are

1. $\mathbb{E}(Z) = 0$
2. $\text{var}(Z) = \frac{\pi^2}{3}$

Proof

1. Again, $\mathbb{E}(Z) = 0$ by symmetry.
2. The second Bernoulli number is $\beta_2 = \frac{1}{6}$. Hence $\text{var}(Z) = \mathbb{E}(Z^2) = (2^2 - 2)\pi^2 \frac{1}{6} = \frac{\pi^2}{3}$.

In the special distribution simulator, select the logistic distribution. Keep the default parameter values and note the shape and location of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = \frac{21}{5}$

Proof

1. Again, $\text{skew}(Z) = 0$ by the symmetry of the distribution.
2. Recall that by symmetry, $\mathbb{E}(Z) = \mathbb{E}(Z^3) = 0$. Also, $|\beta_4| = \frac{1}{30}$, so $\mathbb{E}(Z^4) = (2^4 - 2)\pi^4 \frac{1}{30} = \frac{7\pi^4}{15}$. Hence from the usual computational formula for kurtosis,

$$\text{kurt}(Z) = \frac{\mathbb{E}(Z^4)}{[\text{var}(Z)]^2} = \frac{7\pi^4/15}{\pi^4/9} = \frac{21}{5} \quad (5.29.12)$$

It follows that the *excess kurtosis* of Z is $\text{kurt}(Z) - 3 = \frac{6}{5}$.

Related Distributions

The standard logistic distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) given above. Recall that the *standard uniform distribution* is the continuous uniform distribution on the interval $(0, 1)$.

Connections with the standard uniform distribution.

1. If Z has the standard logistic distribution then

$$U = G(Z) = \frac{e^Z}{1 + e^Z} \quad (5.29.13)$$

has the standard uniform distribution.

2. If U has the standard uniform distribution then

$$Z = G^{-1}(U) = \ln\left(\frac{U}{1-U}\right) = \ln(U) - \ln(1-U) \quad (5.29.14)$$

has the standard logistic distribution.

Since the quantile function has a simple closed form, we can use the usual random quantile method to simulate the standard logistic distribution.

Open the random quantile experiment and select the logistic distribution. Keep the default parameter values and note the shape of the probability density and distribution functions. Run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The standard logistic distribution also has several simple connections with the standard exponential distribution (the exponential distribution with rate parameter 1).

Connections with the standard exponential distribution:

1. If Z has the standard logistic distribution, then $Y = \ln(e^Z + 1)$ has the standard exponential distribution.
2. If Y has the standard exponential distribution then $Z = \ln(e^Y - 1)$ has the standard logistic distribution.

Proof

These results follow from the standard change of variables formula. The transformations, inverses of each other of course, are $y = \ln(e^z + 1)$ and $z = \ln(e^y - 1)$ for $z \in \mathbb{R}$ and $y \in (0, \infty)$. Let g and h denote the PDFs of Z and Y respectively.

1. By definition, $g(z) = e^z / (1 + e^z)^2$ for $z \in \mathbb{R}$ so

$$h(y) = g(z) \frac{dz}{dy} = \frac{\exp[\ln(e^y - 1)]}{(1 + \exp[\ln(e^y - 1)])^2} \frac{e^y}{e^y - 1} = e^{-y}, \quad y \in (0, \infty) \quad (5.29.15)$$

which is the PDF of the standard exponential distribution.

2. By definition, $g(y) = e^{-y}$ for $y \in (0, \infty)$ so

$$g(z) = h(y) \frac{dy}{dz} = \exp[-\ln(e^z + 1)] \frac{e^z}{e^z + 1} = \frac{e^z}{(e^z + 1)^2}, \quad z \in \mathbb{R} \quad (5.29.16)$$

which is the PDF of the standard logistic distribution.

Suppose that X and Y are independent random variables, each with the standard exponential distribution. Then $Z = \ln(X/Y)$ has the standard logistic distribution.

Proof

For $z \in \mathbb{R}$,

$$\mathbb{P}(Z \leq z) = \mathbb{P}[\ln(X/Y) \leq z] = \mathbb{P}(X/Y \leq e^z) = \mathbb{P}(Y \geq e^{-z}X) \quad (5.29.17)$$

Recall that $\mathbb{P}(Y \geq y) = e^{-y}$ for $y \in (0, \infty)$ and X has PDF $x \mapsto e^{-x}$ on $(0, \infty)$. We condition on X :

$$\mathbb{P}(Z \leq z) = \mathbb{E}[\mathbb{P}(Y \geq e^{-z}X | X)] = \int_0^\infty e^{-e^{-z}x} e^{-x} dx = \int_0^\infty e^{(e^{-z}+1)x} dx = \frac{1}{e^{-z}+1} = \frac{e^z}{1+e^z} \quad (5.29.18)$$

As a function of z , this is the distribution function of the standard logistic distribution.

There are also simple connections between the standard logistic distribution and the Pareto distribution.

Connections with the Pareto distribution:

1. If Z has the standard logistic distribution, then $Y = e^Z + 1$ has the Pareto distribution with shape parameter 1.
2. If Y has the Pareto distribution with shape parameter 1, then $Z = \ln(Y - 1)$ has the standard logistic distribution.

Proof

These results follow from the basic change of variables theorem. The transformation, inverses of one another of course, are $y = e^z + 1$, $z = \ln(y - 1)$ for $z \in \mathbb{R}$ and $y \in (1, \infty)$. Let g and h denote PDFs of Z and Y respectively.

1. By definition, $g(z) = e^z / (1 + e^z)^2$ for $z \in \mathbb{R}$. Hence

$$h(y) = g(z) \frac{dz}{dy} = \frac{\exp[\ln(y - 1)]}{(1 + \exp[\ln(y - 1)])^2} \frac{1}{y - 1} = \frac{1}{y^2}, \quad y \in (1, \infty) \quad (5.29.19)$$

which is the PDF of the Pareto distribution with shape parameter 1.

2. By definition, $h(y) = 1/y^2$ for $y \in (1, \infty)$. Hence

$$g(z) = h(y) \frac{dy}{dz} = \frac{1}{(e^z + 1)^2} e^z, \quad z \in \mathbb{R} \quad (5.29.20)$$

which is the PDF of the standard logistic distribution.

Finally, there are simple connections to the extreme value distribution.

If X and Y are independent and each has the standard Gumbel distribution, then $Z = Y - X$ has the standard logistic distribution.

Proof

The distribution function of Y is $G(y) = \exp(-e^{-y})$ for $y \in \mathbb{R}$ and the density function of X is $g(x) = e^{-x} \exp(-e^{-x})$ for $x \in \mathbb{R}$. For $z \in \mathbb{R}$, conditioning on X gives

$$\mathbb{P}(Z \leq z) = \mathbb{P}(Y \leq X + z) = \mathbb{E}[\mathbb{P}(Y \leq X + z | X)] = \int_{-\infty}^{\infty} \exp(-e^{-(x+z)}) e^{-x} \exp(-e^{-x}) dx \quad (5.29.21)$$

Substituting $u = -e^{-(x+z)}$ gives

$$\mathbb{P}(Z \leq z) = \int_{-\infty}^0 e^u \exp(e^z u) e^z du = e^z \int_{-\infty}^0 \exp[u(1 + e^z)] du = \frac{e^z}{1 + e^z}, \quad z \in \mathbb{R} \quad (5.29.22)$$

As a function of z , this is the standard logistic distribution function.

The General Logistic Distribution

The *general logistic distribution* is the location-scale family associated with the standard logistic distribution.

Suppose that Z has the standard logistic distribution. For $a \in \mathbb{R}$ and $b \in (0, \infty)$, random variable $X = a + bZ$ has the *logistic distribution* with *location parameter* a and *scale parameter* b .

Distribution Functions

Analogies of the results above for the general logistic distribution follow easily from basic properties of the location-scale transformation. Suppose that X has the logistic distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$.

The probability density function f of X is given by

$$f(x) = \frac{\exp(\frac{x-a}{b})}{b [1 + \exp(\frac{x-a}{b})]^2}, \quad x \in \mathbb{R} \quad (5.29.23)$$

1. f is symmetric about $x = a$.
2. f increases and then decreases, with mode $x = a$.
3. f is concave upward, then downward, then upward again, with inflection points at $x = a \pm \ln(2 + \sqrt{3})b$.

Proof

Recall that

$$f(x) = \frac{1}{b} g\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.29.24)$$

where g is the [standard logistic PDF](#).

In the special distribution simulator, select the logistic distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function F of X is given by

$$F(x) = \frac{\exp(\frac{x-a}{b})}{1 + \exp(\frac{x-a}{b})}, \quad x \in \mathbb{R} \quad (5.29.25)$$

Proof

Recall that

$$F(x) = G\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.29.26)$$

where G is the [standard logistic CDF](#).

The quantile function F^{-1} of X is given by

$$F^{-1}(p) = a + b \ln\left(\frac{p}{1-p}\right), \quad p \in (0, 1) \quad (5.29.27)$$

1. The first quartile is $a - b \ln 3$.
2. The median is a .
3. The third quartile is $a + b \ln 3$

Proof

Recall that $F^{-1}(p) = a + bG^{-1}(p)$ for $p \in (0, 1)$, where G^{-1} is the [standard logistic quantile function](#).

In the special distribution calculator, select the logistic distribution. Vary the parameters and note the shape and location of the probability density function and the distribution function. For selected values of the parameters, find the quantiles of order 0.1 and 0.9.

Moments

Suppose again that X has the logistic distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. Recall that B denotes the beta function and Γ the gamma function.

The moment generating function M of X is given by

$$M(t) = e^{at} B(1+bt, 1-bt) = e^{at} \Gamma(1+bt) \Gamma(1-bt), \quad t \in (-1, 1) \quad (5.29.28)$$

Proof

Recall that $M(t) = e^{at} m(bt)$ where m is the [standard logistic MGF](#).

The mean and variance of X are

1. $\mathbb{E}(X) = a$
2. $\text{var}(X) = b^2 \frac{\pi^2}{3}$

Proof

By [definition](#) we can assume $X = a + bZ$ where Z has the standard logistic distribution. Using the [mean and variance of \$Z\$](#) we have

1. $\mathbb{E}(X) = a + b\mathbb{E}(Z) = a$
2. $\text{var}(X) = b^2 \text{var}(Z) = b^2 \frac{\pi^2}{3}$

In the special distribution simulator, select the logistic distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = \frac{21}{5}$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are invariant under location-scale transformations. So the skewness and kurtosis of Z are the same as the [skewness and kurtosis of \$Z\$](#) .

Once again, it follows that the excess kurtosis of X is $\text{kurt}(X) - 3 = \frac{6}{5}$. The central moments of X can be given in terms of the Bernoulli numbers. As before, let β_n denote the Bernoulli number of order $n \in \mathbb{N}$.

Let $n \in \mathbb{N}$.

1. If n is odd then $\mathbb{E}[(X - a)^n] = 0$.
2. If n is even then $\mathbb{E}[(X - a)^n] = (2^n - 2)\pi^n b^n |\beta_n|$

Proof

Again by [definition](#) we can take $X = a + bZ$ where Z has the standard logistic distribution. Then $\mathbb{E}[(X - a)^n] = b^n \mathbb{E}(Z^n)$ so the results follow from the [moments of \$Z\$](#) .

Related Distributions

The general logistic distribution is a location-scale family, so it is trivially closed under location-scale transformations.

Suppose that X has the logistic distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, and that $c \in \mathbb{R}$ and $d \in (0, \infty)$. Then $Y = c + dX$ has the logistic distribution with location parameter $c + ad$ and scale parameter bd .

Proof

Again by [definition](#) we can take $X = a + bZ$ where Z has the standard logistic distribution. Then $Y = c + dX = (c + ad) + (bd)Z$.

This page titled [5.29: The Logistic Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.30: The Extreme Value Distribution

Extreme value distributions arise as limiting distributions for maximums or minimums (*extreme values*) of a sample of independent, identically distributed random variables, as the sample size increases. Thus, these distributions are important in probability and mathematical statistics.

The Standard Distribution for Maximums

Distribution Functions

The *standard extreme value distribution* (for maximums) is a continuous distribution on \mathbb{R} with distribution function G given by

$$G(v) = \exp(-e^{-v}), \quad v \in \mathbb{R} \quad (5.30.1)$$

Proof

Note that G is continuous, increasing, and satisfies $G(v) \rightarrow 0$ as $v \rightarrow -\infty$ and $G(v) \rightarrow 1$ as $v \rightarrow \infty$.

The distribution is also known as the *standard Gumbel distribution* in honor of Emil Gumbel. As we will show [below](#), it arises as the limit of the maximum of n independent random variables, each with the standard exponential distribution (when this maximum is appropriately centered). This fact is the main reason that the distribution is *special*, and is the reason for the name. For the remainder of this discussion, suppose that random variable V has the standard Gumbel distribution.

The probability density function g of V is given by

$$g(v) = e^{-v} \exp(-e^{-v}) = \exp[-(e^{-v} + v)], \quad v \in \mathbb{R} \quad (5.30.2)$$

1. g increases and then decreases with mode $v = 0$
2. g is concave upward, then downward, then upward again, with inflection points at $v = \ln[(3 \pm \sqrt{5})/2] \approx \pm 0.9264$.

Proof

These results follow from standard calculus. The PDF is $g = G'$.

1. The first derivative of g satisfies $g'(v) = g(v)(e^{-v} - 1)$ for $v \in \mathbb{R}$.
2. The second derivative of g satisfies $g''(v) = g(v)(e^{-2v} - 3e^{-v} + 1)$ for $v \in \mathbb{R}$.

In the special distribution simulator, select the extreme value distribution. Keep the default parameter values and note the shape and location of the probability density function. In particular, note the lack of symmetry. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function G^{-1} of V is given by

$$G^{-1}(p) = -\ln[-\ln(p)], \quad p \in (0, 1) \quad (5.30.3)$$

1. The first quartile is $-\ln(-\ln 4) \approx -0.3266$.
2. The median is $-\ln(-\ln 2) \approx 0.3665$
3. The third quartile is $-\ln(\ln 4 - \ln 3) \approx 1.2459$

Proof

The formula for G^{-1} follows from solving $p = G(v)$ for v in terms of p .

In the special distribution calculator, select the extreme value distribution. Keep the default parameter values and note the shape and location of the probability density and distribution functions. Compute the quantiles of order 0.1, 0.3, 0.6, and 0.9

Moments

Suppose again that V has the standard Gumbel distribution. The moment generating function of V has a simple expression in terms of the gamma function Γ .

The moment generating function m of V is given by

$$m(t) = \mathbb{E}(e^{tV}) = \Gamma(1-t), \quad t \in (-\infty, 1) \quad (5.30.4)$$

Proof

Note that

$$m(t) = \int_{-\infty}^{\infty} e^{tv} \exp(-e^{-v}) e^{-v} dv \quad (5.30.5)$$

The substitution $x = e^{-v}$, $dx = -e^{-v} dv$ gives $m(t) = \int_0^{\infty} x^{-t} e^{-x} dx = \Gamma(1-t)$ for $t \in (-\infty, 1)$.

Next we give the mean and variance. First, recall that the *Euler constant*, named for Leonhard Euler is defined by

$$\gamma = -\Gamma'(1) = -\int_0^{\infty} e^{-x} \ln x dx \approx 0.5772156649 \quad (5.30.6)$$

The mean and variance of V are

1. $\mathbb{E}(V) = \gamma$
2. $\text{var}(V) = \frac{\pi^2}{6}$

Proof

These results follow from the moment generating function.

1. $m'(t) = -\Gamma'(1-t)$ and so $\mathbb{E}(V) = m'(0) = -\Gamma'(1) = \gamma$.
2. $m''(t) = \Gamma''(1-t)$ and

$$\mathbb{E}(V^2) = m''(0) = \Gamma''(1) = \int_0^{\infty} (\ln x)^2 e^{-x} dx = \gamma^2 + \frac{\pi^2}{6} \quad (5.30.7)$$

$$\text{Hence } \text{var}(V) = \mathbb{E}(V^2) - [\mathbb{E}(V)]^2 = \frac{\pi^2}{6}$$

In the special distribution simulator, select the extreme value distribution and keep the default parameter values. Note the shape and location of the mean \pm standard deviation bar. Run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Next we give the skewness and kurtosis of V . The skewness involves a value of the *Riemann zeta function* ζ , named of course for Georg Riemann. Recall that ζ is defined by

$$\zeta(n) = \sum_{k=1}^{\infty} \frac{1}{k^n}, \quad n > 1 \quad (5.30.8)$$

The skewness and kurtosis of V are

1. $\text{skew}(V) = 12\sqrt{6}\zeta(3)/\pi^3 \approx 1.13955$
2. $\text{kurt}(V) = \frac{27}{5}$

The particular value of the zeta function, $\zeta(3)$, is known as *Apéry's constant*. From (b), it follows that the *excess kurtosis* is $\text{kurt}(V) - 3 = \frac{12}{5}$.

Related Distributions

The standard Gumbel distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) given above. Recall that the standard uniform distribution is the continuous uniform distribution on

the interval $(0, 1)$.

The standard Gumbel and standard uniform distributions are related as follows:

1. If U has the standard uniform distribution then $V = G^{-1}(U) = -\ln(-\ln U)$ has the standard Gumbel distribution.
2. If V has the standard Gumbel distribution then $U = G(V) = \exp(e^{-V})$ has the standard uniform distribution.

So we can simulate the standard Gumbel distribution using the usual random quantile method.

Open the random quantile experiment and select the extreme value distribution. Keep the default parameter values and note again the shape and location of the probability density and distribution functions. Run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The standard Gumbel distribution also has simple connections with the standard exponential distribution (the exponential distribution with rate parameter 1).

The standard Gumbel and standard exponential distributions are related as follows:

1. If X has the standard exponential distribution then $V = -\ln X$ has the standard Gumbel distribution.
2. If V has the standard Gumbel distribution then $X = e^{-V}$ has the standard exponential distribution.

Proof

These results follow from the usual change of variables theorem. The transformations are $v = -\ln x$ and $x = e^{-v}$ for $x \in (0, \infty)$ and $v \in \mathbb{R}$, and these are inverses of each other. Let f and g denote PDFs of X and V respectively.

1. We start with $f(x) = e^{-x}$ for $x \in (0, \infty)$ and then

$$g(v) = f(x) \left| \frac{dx}{dv} \right| = \exp(-e^{-v}) e^{-v}, \quad v \in \mathbb{R} \quad (5.30.9)$$

so V has the standard Gumbel distribution.

2. We start with $g(v) = \exp(-e^{-v}) e^{-v}$ for $v \in \mathbb{R}$ and then

$$f(x) = g(v) \left| \frac{dv}{dx} \right| = \exp[-\exp(\ln x)] \exp(\ln x) \frac{1}{x} = e^{-x}, \quad x \in (0, \infty) \quad (5.30.10)$$

so X has the standard exponential distribution.

As noted in the introduction, the following theorem provides the motivation for the name *extreme value distribution*.

Suppose that (X_1, X_2, \dots) is a sequence of independent random variables, each with the standard exponential distribution. The distribution of $Y_n = \max\{X_1, X_2, \dots, X_n\} - \ln n$ converges to the standard Gumbel distribution as $n \rightarrow \infty$.

Proof

Let $X_{(n)} = \max\{X_1, X_2, \dots, X_n\}$, so that $X_{(n)}$ is the n th order statistics of the random sample (X_1, X_2, \dots, X_n) . Let G denote the standard exponential CDF, so that $G(x) = 1 - e^{-x}$ for $x \in [0, \infty)$. Note that $X_{(n)}$ has CDF G^n . Let F_n denote the CDF of Y_n . For $x \in \mathbb{R}$

$$F_n(x) = \mathbb{P}(Y_n \leq x) = \mathbb{P}[X_{(n)} \leq x + \ln n] = G^n(x + \ln n) = [1 - e^{-(x + \ln n)}]^n = \left(1 - \frac{e^{-x}}{n}\right)^n \quad (5.30.11)$$

By a famous limit from calculus, $F_n(x) \rightarrow e^{-e^{-x}}$ as $n \rightarrow \infty$.

The General Extreme Value Distribution

As with many other distributions we have studied, the standard extreme value distribution can be generalized by applying a linear transformation to the standard variable. First, if V has the standard Gumbel distribution (the standard extreme value distribution for maximums), then $-V$ has the standard extreme value distribution for minimums. Here is the general definition.

Suppose that V has the standard Gumbel distribution, and that $a, b \in \mathbb{R}$ with $b \neq 0$. Then $X = a + bV$ has the *extreme value distribution* with location parameter a and scale parameter $|b|$.

1. If $b > 0$, then the distribution corresponds to maximums.
2. If $b < 0$, then the distribution corresponds to minimums.

So the family of distributions with $a \in \mathbb{R}$ and $b \in (0, \infty)$ is a location-scale family associated with the standard distribution for maximums, and the family of distributions with $a \in \mathbb{R}$ and $b \in (-\infty, 0)$ is the location-scale family associated with the standard distribution for minimums.. The distributions are also referred to more simply as *Gumbel distributions* rather than extreme value distributions. The web apps in this project use only the extreme value distributions for maximums. As you will see below, the differences in the distribution for maximums and the distribution for minimums are minor. For the remainder of this discussion, suppose that X has the form given in the [definition](#).

Distribution Functions

Let F denote the distribution function of X .

1. If $b > 0$ then

$$F(x) = \exp \left[-\exp \left(-\frac{x-a}{b} \right) \right], \quad x \in \mathbb{R} \quad (5.30.12)$$

2. If $b < 0$ then

$$F(x) = 1 - \exp \left[-\exp \left(-\frac{x-a}{b} \right) \right], \quad x \in \mathbb{R} \quad (5.30.13)$$

Proof

Let G denote the [CDF of \$V\$](#) . Then

1. $F(x) = G \left(\frac{x-a}{b} \right)$ for $x \in \mathbb{R}$
2. $F(x) = 1 - G \left(\frac{x-a}{b} \right)$ for $x \in \mathbb{R}$

Let f denote the probability density function of X . Then

$$f(x) = \frac{1}{|b|} \exp \left(-\frac{x-a}{b} \right) \exp \left[-\exp \left(-\frac{x-a}{b} \right) \right], \quad x \in \mathbb{R} \quad (5.30.14)$$

Proof

Let g denote the [PDF of \$V\$](#) . By the change of variables formula,

$$f(x) = \frac{1}{|b|} g \left(\frac{x-a}{b} \right), \quad x \in \mathbb{R} \quad (5.30.15)$$

Open the special distribution simulator and select the extreme value distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function F^{-1} of X is given as follows

1. If $b > 0$ then $F^{-1}(p) = a - b \ln(-\ln p)$ for $p \in (0, 1)$.
2. If $b < 0$ then $F^{-1}(p) = a - b \ln[-\ln(1-p)]$, for $p \in (0, 1)$

Proof

Let G^{-1} denote the [quantile function of \$V\$](#) . Then

1. $F^{-1}(p) = a + bG^{-1}(p)$ for $p \in (0, 1)$.
2. $F^{-1}(p) = a - bG^{-1}(1-p)$ for $p \in (0, 1)$.

Open the special distribution calculator and select the extreme value distribution. Vary the parameters and note the shape and location of the probability density and distribution functions. For selected values of the parameters, compute a few values of the quantile function and the distribution function.

Moments

Suppose again that $X = a + bV$ where V has the standard Gumbel distribution, and that $a, b \in \mathbb{R}$ with $b \neq 0$.

The moment generating function M of X is given by $M(t) = e^{at}\Gamma(1 - bt)$.

1. With domain $t \in (-\infty, 1/b)$ if $b > 0$
2. With domain $t \in (1/b, \infty)$ if $b < 0$

Proof

Let m denote the [MGF of \$V\$](#) . Then $M(t) = e^{at}m(bt)$ for $bt < 1$

The mean and variance of X are

1. $\mathbb{E}(X) = a + b\gamma$
2. $\text{var}(X) = b^2 \frac{\pi^2}{6}$

Proof

These results follow from the [mean and variance of \$V\$](#) and basic properties of expected value and variance.

1. $\mathbb{E}(X) = a + b\mathbb{E}(V)$
2. $\text{var}(X) = b^2\text{var}(V)$

Open the special distribution simulator and select the extreme value distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness of X is

1. $\text{skew}(X) = 12\sqrt{6}\zeta(3)/\pi^3 \approx 1.13955$ if $b > 0$.
2. $\text{skew}(X) = -12\sqrt{6}\zeta(3)/\pi^3 \approx -1.13955$ if $b < 0$

Proof

Recall that skewness is defined in terms of the standard score, and hence is invariant under linear transformations with positive slope. A linear transformation with negative slope changes the sign of the skewness. Hence these results follow from the [skewness of \$V\$](#) .

The kurtosis of X is $\text{kurt}(X) = \frac{27}{5}$

Proof

Recall that kurtosis is defined in terms of the standard score and is invariant under linear transformations with nonzero slope. Hence this result follows from the [kurtosis of \$V\$](#) .

Once again, the excess kurtosis is $\text{kurt}(X) - 3 = \frac{12}{5}$.

Related Distributions

Since the general extreme value distributions are location-scale families, they are trivially closed under linear transformations of the underlying variables (with nonzero slope).

Suppose that X has the extreme value distribution with parameters a, b with $b \neq 0$ and that $c, d \in \mathbb{R}$ with $d \neq 0$. Then $Y = c + dX$ has the extreme value distribution with parameters $ad + c$ and bd .

Proof

By definition, we can write $X = a + bV$ where V has the standard Gumbel distribution. Hence $Y = c + dX = (ad + c) + (bd)V$.

Note if $d > 0$ then X and Y have the same association (max, max) or (min, min). If $d < 0$ then X and Y have opposite associations (max, min) or (min, max).

As with the standard Gumbel distribution, the general Gumbel distribution has the usual connections with the standard uniform distribution by means of the distribution and quantile functions. Since the quantile function has a simple closed form, the latter connection leads to the usual random quantile method of simulation. We state the result for maximums.

Suppose that $a, b \in \mathbb{R}$ with $b \neq 0$. Let F denote [distribution function](#) and let F^{-1} denote the [quantile function](#) above

1. If U has the standard uniform distribution then $X = F^{-1}(U)$ has the extreme value distribution with parameters a and b .
2. If X has the extreme value distribution with parameters a and b then $U = F(X)$ has the standard uniform distribution.

Open the random quantile experiment and select the extreme value distribution. Vary the parameters and note again the shape and location of the probability density and distribution functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The extreme value distribution for maximums has a simple connection to the Weibull distribution, and this generalizes the [connection](#) between the standard Gumbel and exponential distributions above. There is a similar result for the extreme value distribution for minimums.

The extreme value and Weibull distributions are related as follows:

1. If X has the extreme value distribution with parameters $a \in \mathbb{R}$ and $b \in (0, \infty)$, then $Y = e^{-X}$ has the Weibull distribution with shape parameter $\frac{1}{b}$ and scale parameter e^{-a} .
2. If Y has the Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ then $X = -\ln Y$ has the extreme value distribution with parameters $-\ln b$ and $\frac{1}{k}$.

Proof

As before, these results can be obtained using the change of variables theorem for probability density functions. We give an alternate proof using special forms of the random variables.

1. We can write $X = a + bV$ where V has the standard Gumbel distribution. Hence

$$Y = e^{-X} = e^{-a} (e^{-V})^b \quad (5.30.16)$$

As shown in [above](#), e^{-V} has the standard exponential distribution and therefore Y has the Weibull distribution with shape parameter $1/b$ and scale parameter e^{-a} .

2. We can write $Y = bU^{1/k}$ where U has the standard exponential distribution. Hence

$$X = -\ln Y = -\ln b + \frac{1}{k}(-\ln U) \quad (5.30.17)$$

As shown in [above](#), $-\ln U$ has the standard Gumbel distribution and hence X has the Gumbel distribution with location parameter $-\ln b$ and scale parameter $\frac{1}{k}$.

This page titled [5.30: The Extreme Value Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.31: The Hyperbolic Secant Distribution

The *hyperbolic secant distribution* is a location-scale family with a number of interesting parallels to the normal distribution. As the name suggests, the hyperbolic secant function plays an important role in the distribution, so we should first review some definitions

The hyperbolic trig functions \sinh , \cosh , \tanh , and sech are defined as follows, for $x \in \mathbb{R}$

$$\sinh x = \frac{1}{2}(e^x - e^{-x}) \quad (5.31.1)$$

$$\cosh x = \frac{1}{2}(e^x + e^{-x}) \quad (5.31.2)$$

$$\tanh x = \frac{\sinh x}{\cosh x} = \frac{e^x - e^{-x}}{e^x + e^{-x}} \quad (5.31.3)$$

$$\operatorname{sech} x = \frac{1}{\cosh x} = \frac{2}{e^x + e^{-x}} \quad (5.31.4)$$

The Standard Hyperbolic Secant Distribution

Distribution Functions

The *standard hyperbolic secant distribution* is a continuous distribution on \mathbb{R} with probability density function g given by

$$g(z) = \frac{1}{2} \operatorname{sech} \left(\frac{\pi}{2} z \right), \quad z \in \mathbb{R} \quad (5.31.5)$$

1. g is symmetric about 0.
2. g increases and then decreases with mode $z = 0$.
3. g is concave upward then downward then upward again, with inflection points at $z = \pm \frac{2}{\pi} \ln(\sqrt{2} + 1) \approx \pm 0.561$.

Proof

If we multiply numerator and denominator of $\operatorname{sech}(x)$ by e^x and then use the simple substitution $u = e^x$ we see that

$$\int \operatorname{sech}(x) dx = \int 2 \frac{e^x}{e^{2x} + 1} dx = 2 \int \frac{1}{u^2 + 1} du = 2 \arctan(u) = 2 \arctan(e^x) \quad (5.31.6)$$

It follows that

$$\int_{-\infty}^{\infty} g(z) dz = \frac{2}{\pi} \arctan \left(\frac{\pi}{2} e^z \right) \Big|_{-\infty}^{\infty} = 1 \quad (5.31.7)$$

The properties of g result follow from standard calculus. Recall that $\operatorname{sech}' = -\tanh \operatorname{sech}$ and $\tanh' = \operatorname{sech}^2$.

So g has the classic unimodal shape. Recall that the inflection points in the standard normal probability density function are ± 1 . Compared to the standard normal distribution, the hyperbolic secant distribution is more peaked at the mode 0 but has fatter tails.

Open the special distribution simulator and select the hyperbolic secant distribution. Keep the default parameter settings and note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function G of the standard hyperbolic secant distribution is given by

$$G(z) = \frac{2}{\pi} \arctan \left[\exp \left(\frac{\pi}{2} z \right) \right], \quad z \in \mathbb{R} \quad (5.31.8)$$

Proof

Of course, $G(z) = \int_{-\infty}^z g(x) dx$. The form of G follows from the same integration methods used for the [PDF](#).

The quantile function G^{-1} of the standard hyperbolic secant distribution is given by

$$G^{-1}(p) = \frac{2}{\pi} \ln \left[\tan \left(\frac{\pi p}{2} \right) \right], \quad p \in (0, 1) \quad (5.31.9)$$

1. The first quartile is $G^{-1} \left(\frac{1}{4} \right) = -\frac{2}{\pi} \ln(1 + \sqrt{2}) \approx -0.561$
2. The median is $G^{-1} \left(\frac{1}{2} \right) = 0$
3. The third quartile is $G^{-1} \left(\frac{3}{4} \right) = \frac{2}{\pi} \ln(1 + \sqrt{2}) \approx 0.561$

Proof

The formula for G^{-1} follows by solving $G(z) = p$ for z in terms of p . For the quartiles, note that $\tan(\pi/8) = \sqrt{2} - 1 = 1/(\sqrt{2} + 1)$ and $\tan(3\pi/8) = \sqrt{2} + 1$.

Of course, the fact that the median is 0 also follows from the symmetry of the distribution, as does the relationship between the first and third quartiles. In general, $G^{-1}(1 - p) = -G^{-1}(p)$ for $p \in (0, 1)$. Note that the first and third quartiles coincide with the inflection points, whereas in the normal distribution, the inflection points are at ± 1 and coincide with the standard deviation.

Open the sepcial distribution calculator and select the hyperbolic secant distribution. Keep the default values of the parameters and note the shape of the distribution and probability density functions. Compute a few values of the distribution and quantile functions.

Moments

Suppose that Z has the standard hyperbolic secant distribution. The moments of Z are easiest to compute from the generating functions.

The characteristic function χ of Z is the hyperbolic secant function:

$$\chi(t) = \operatorname{sech}(t), \quad t \in \mathbb{R} \quad (5.31.10)$$

Proof

The charateristic function is

$$\chi(t) = \mathbb{E}(e^{itZ}) = \int_{-\infty}^{\infty} \frac{e^{itz}}{e^{\pi z/2} + e^{-\pi z/2}} dz \quad (5.31.11)$$

The evaluation of this integral to $\operatorname{sech}(t)$ is complicated, but the details can be found in the book [Continuous Univariate Distributions](#) by Johnson, Kotz, and Balakrishnan.

Note that the probability density function can be obtained from the characteristic function by a scale transformation: $g(z) = \frac{1}{2} \chi \left(\frac{\pi}{2} z \right)$ for $z \in \mathbb{R}$. This is another curious similarity to the normal distribution: the probability density function ϕ and characteristic function χ of the standard normal distribution are related by $\phi(z) = \frac{1}{\sqrt{2\pi}} \chi(z)$.

The moment generating function m of Z is the secant function:

$$m(t) = \mathbb{E}(e^{tZ}) = \sec(t), \quad t \in \left(-\frac{\pi}{2}, \frac{\pi}{2} \right) \quad (5.31.12)$$

Proof

This follows from the [characteristic function](#) since $m(t) = \chi(-it)$.

It follows that Z has moments of all orders, and then by symmetry, that the odd order moments are all 0.

The mean and variance of Z are

1. $\mathbb{E}(Z) = 0$
2. $\operatorname{var}(Z) = 1$

Proof

As noted, the mean is 0 by symmetry. Hence also $\text{var}(Z) = \mathbb{E}(Z^2) = m''(0)$. But $m''(t) = \sec(t) \tan^2(t) + \sec^3(t)$, so $\text{var}(Z) = 1$.

Thus, the standard hyperbolic secant distribution has mean 0 and variance 1, just like the standard normal distribution.

Open the special distribution simulator and select the hyperbolic secant distribution. Keep the default parameters and note the size and location of the mean \pm standard deviation bar. Run the simulation 1000 times compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of Z are

1. $\text{skew}(Z) = 0$
2. $\text{kurt}(Z) = 5$

Proof

The skewness is 0 by the symmetry of the distribution. Also, since the mean is 0 and the variance 1, $\text{kurt}(Z) = \mathbb{E}(Z^4) = m^{(4)}(0)$. But by standard calculus,

$$m^{(4)}(t) = \sec(t) \tan^4(t) + 18 \sec^3(t) \tan^2(t) + 5 \sec^5(t) \quad (5.31.13)$$

and hence $m^{(4)}(0) = 5$.

Recall that the kurtosis of the standard normal distribution is 3, so the *excess kurtosis* of the standard hyperbolic secant distribution is $\text{kurt}(Z) - 3 = 2$. This distribution is more sharply peaked at the mean 0 and has fatter tails, compared with the normal.

Related Distributions

The standard hyperbolic secant distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) computed above.

The standard hyperbolic secant distribution is related to the standard uniform distribution as follows:

1. If Z has the standard hyperbolic secant distribution then

$$U = G(Z) = \frac{2}{\pi} \arctan \left[\exp \left(\frac{\pi}{2} Z \right) \right] \quad (5.31.14)$$

has the standard uniform distribution.

2. If U has the standard uniform distribution then

$$Z = G^{-1}(U) = \frac{2}{\pi} \ln \left[\tan \left(\frac{\pi}{2} U \right) \right] \quad (5.31.15)$$

has the standard hyperbolic secant distribution.

Since the quantile function has a simple closed form, the standard hyperbolic secant distribution can be easily simulated by means of the random quantile method.

Open the random quantile experiment and select the hyperbolic secant distribution. Keep the default parameter values and note again the shape of the probability density and distribution functions. Run the experiment 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The General Hyperbolic Secant Distribution

The standard hyperbolic secant distribution is generalized by adding location and scale parameters.

Suppose that Z has the standard hyperbolic secant distribution and that $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$. Then $X = \mu + \sigma Z$ has the *hyperbolic secant distribution* with location parameter μ and scale parameter σ .

Distribution Functions

Suppose that X has the hyperbolic secant distribution with location parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$.

The probability density function f of X is given by

$$f(x) = \frac{1}{2\sigma} \operatorname{sech} \left[\frac{\pi}{2} \left(\frac{x - \mu}{\sigma} \right) \right], \quad x \in \mathbb{R} \quad (5.31.16)$$

1. f is symmetric about μ .
2. f increases and then decreases with mode $x = \mu$.
3. f is concave upward then downward then upward again, with inflection points at $x = \mu \pm \frac{2}{\pi} \ln(\sqrt{2} + 1) \sigma \approx \mu \pm 0.561 \sigma$.

Proof

Recall that $f(x) = \frac{1}{\sigma} g\left(\frac{x - \mu}{\sigma}\right)$ for $x \in \mathbb{R}$ where g is the [standard hyperbolic secant PDF](#).

Open the special distribution simulator and select the hyperbolic secant distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function F of X is given by

$$F(x) = \frac{2}{\pi} \arctan \left\{ \exp \left[\frac{\pi}{2} \left(\frac{x - \mu}{\sigma} \right) \right] \right\}, \quad x \in \mathbb{R} \quad (5.31.17)$$

Proof

Recall that $F(x) = G\left(\frac{x - \mu}{\sigma}\right)$ for $x \in \mathbb{R}$ where G is the [standard hyperbolic secant CDF](#).

The quantile function F^{-1} of X is given by

$$F^{-1}(p) = \mu + \sigma \frac{2}{\pi} \ln \left[\tan \left(\frac{\pi}{2} p \right) \right], \quad p \in (0, 1) \quad (5.31.18)$$

1. The first quartile is $F^{-1}\left(\frac{1}{4}\right) = \mu - \frac{2}{\pi} \ln(1 + \sqrt{2}) \sigma \approx \mu - 0.561 \sigma$
2. The median is $F^{-1}\left(\frac{1}{2}\right) = \mu$
3. The third quartile is $F^{-1}\left(\frac{3}{4}\right) = \mu + \frac{2}{\pi} \ln(1 + \sqrt{2}) \sigma \approx \mu + 0.561 \sigma$

Proof

Recall that $F^{-1}(p) = \mu + \sigma G^{-1}(p)$ where G^{-1} is the [standard quantile function](#).

Open the special distribution calculator and select the hyperbolic secant distribution. Vary the parameters and note the shape of the distribution and density functions. For various values of the parameters, compute a few values of the distribution and quantile functions.

Moments

Suppose again that X has the hyperbolic secant distribution with location parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$.

The moment generating function M of X is given by

$$M(t) = e^{\mu t} \operatorname{sech}(\sigma t), \quad t \in \left(-\frac{\pi}{2\sigma}, \frac{\pi}{2\sigma} \right) \quad (5.31.19)$$

Proof

Recall that $M(t) = e^{\mu t} m(\sigma t)$ where m is the [standard hyperbolic secant MGF](#).

Just as in the normal distribution, the location and scale parameters are the mean and standard deviation, respectively.

The mean and variance of X are

1. $\mathbb{E}(X) = \mu$
2. $\text{var}(X) = \sigma^2$

Proof

These results follow from the representation $X = \mu + \sigma Z$ where Z has the standard hyperbolic secant distribution, basic properties of expected value and variance, and the [mean and variance of \$Z\$](#) :

1. $\mathbb{E}(X) = \mu + \sigma \mathbb{E}(Z) = \mu$
2. $\text{var}(X) = \sigma^2 \text{var}(Z) = \sigma^2$

Open the special distribution simulator and select the hyperbolic secant distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 0$
2. $\text{kurt}(X) = 5$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are invariant under location-scale transformations. Thus, the skewness and kurtosis of X are the same as the [skewness and kurtosis of the standard distribution](#).

Once again, the excess kurtosis is $\text{kurt}(X) - 3 = 2$

Related Distributions

Since the hyperbolic secant distribution is a location-scale family, it is trivially closed under location-scale transformations.

Suppose that X has the hyperbolic secant distribution with location parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$, and that $a \in \mathbb{R}$ and $b \in (0, \infty)$. Then $Y = a + bX$ has the hyperbolic secant distribution with location parameter $a + b\mu$ and scale parameter $b\sigma$.

Proof

By [definition](#), we can take $X = \mu + \sigma Z$ where Z has the standard hyperbolic secant distribution. Hence $Y = a + bX = (a + b\mu) + (b\sigma)Z$.

The hyperbolic secant distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) computed above.

Suppose that $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$.

1. If X has the hyperbolic secant distribution with location parameter μ and scale parameter σ then

$$U = F(X) = \frac{2}{\pi} \arctan \left\{ \exp \left[\frac{\pi}{2} \left(\frac{X - \mu}{\sigma} \right) \right] \right\} \quad (5.31.20)$$

has the standard uniform distribution.

2. If U has the standard uniform distribution then

$$X = F^{-1}(U) = \mu + \sigma \frac{2}{\pi} \ln \left[\tan \left(\frac{\pi}{2} U \right) \right] \quad (5.31.21)$$

has the hyperbolic secant distribution with location parameter μ and scale parameter σ .

Since the quantile function has a simple closed form, the hyperbolic secant distribution can be easily simulated by means of the random quantile method.

Open the random quantile experiment and select the hyperbolic secant distribution. Vary the parameters and note again the shape of the probability density and distribution functions. Run the experiment 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

This page titled [5.31: The Hyperbolic Secant Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.32: The Cauchy Distribution

The *Cauchy distribution*, named of course for the ubiquitous Augustin Cauchy, is interesting for a couple of reasons. First, it is a simple family of distributions for which the expected value (and other moments) do not exist. Second, the family is closed under the formation of sums of independent variables, and hence is an infinitely divisible family of distributions.

The Standard Cauchy Distribution

Distribution Functions

The *standard Cauchy distribution* is a continuous distribution on \mathbb{R} with probability density function g given by

$$g(x) = \frac{1}{\pi(1+x^2)}, \quad x \in \mathbb{R} \quad (5.32.1)$$

1. g is symmetric about $x = 0$
2. g increases and then decreases, with mode $x = 0$.
3. g is concave upward, then downward, and then upward again, with inflection points at $x = \pm \frac{1}{\sqrt{3}}$.
4. $g(x) \rightarrow 0$ as $x \rightarrow \infty$ and as $x \rightarrow -\infty$

Proof

Note that

$$\int_{-\infty}^{\infty} \frac{1}{1+x^2} dx = \arctan x \Big|_{-\infty}^{\infty} = \frac{\pi}{2} - \left(-\frac{\pi}{2}\right) = \pi \quad (5.32.2)$$

and hence g is a PDF. Parts (a)–(d) follow from basic calculus.

Thus, the graph of g has a simple, symmetric, unimodal shape that is qualitatively (but certainly not quantitatively) like the standard normal probability density function. The probability density function g is obtained by normalizing the function

$$x \mapsto \frac{1}{1+x^2}, \quad x \in \mathbb{R} \quad (5.32.3)$$

The graph of this function is known as the *witch of Agnesi*, named for the Italian mathematician Maria Agnesi.

Open the special distribution simulator and select the Cauchy distribution. Keep the default parameter values to get the standard Cauchy distribution and note the shape and location of the probability density function. Run the simulation 1000 times and compare the empirical density function to the probability density function.

The standard Cauchy distribution function G given by $G(x) = \frac{1}{2} + \frac{1}{\pi} \arctan x$ for $x \in \mathbb{R}$

Proof

For $x \in \mathbb{R}$,

$$G(x) = \int_{-\infty}^x g(t) dt = \frac{1}{\pi} \arctan t \Big|_{-\infty}^x = \frac{1}{\pi} \arctan x + \frac{1}{2} \quad (5.32.4)$$

The standard Cauchy quantile function G^{-1} is given by $G^{-1}(p) = \tan\left[\pi\left(p - \frac{1}{2}\right)\right]$ for $p \in (0, 1)$. In particular,

1. The first quartile is $G^{-1}\left(\frac{1}{4}\right) = -1$
2. The median is $G^{-1}\left(\frac{1}{2}\right) = 0$
3. The third quartile is $G^{-1}\left(\frac{3}{4}\right) = 1$

Proof

As usual, G^{-1} is computed from the CDF G by solving $G(x) = p$ for x in terms of p .

Of course, the fact that the median is 0 also follows from the symmetry of the distribution, as does the fact that $G^{-1}(1-p) = -G^{-1}(p)$ for $p \in (0, 1)$.

Open the special distribution calculator and select the Cauchy distribution. Keep the default parameter values and note the shape of the distribution and probability density functions. Compute a few quantiles.

Moments

Suppose that random variable X has the standard Cauchy distribution. As we noted in the introduction, part of the fame of this distribution comes from the fact that the expected value does not exist.

$\mathbb{E}(X)$ does not exist.

Proof

By definition, $\mathbb{E}(X) = \int_{-\infty}^{\infty} xg(x) dx$. For the improper integral to exist, even as an extended real number, at least one of the integrals $\int_{-\infty}^a xg(x) dx$ and $\int_a^{\infty} xg(x) dx$ must be finite, for some (and hence every) $a \in \mathbb{R}$. But by a simple substitution,

$$\int_a^{\infty} xg(x) dx = \int_a^{\infty} x \frac{1}{\pi(1+x^2)} dx = \frac{1}{2\pi} \ln(1+x^2) \Big|_a^{\infty} = \infty \quad (5.32.5)$$

and similarly, $\int_{-\infty}^a xg(x) dx = -\infty$.

By symmetry, if the expected value *did* exist, it would have to be 0, just like the median and the mode, but alas the mean does not exist. Moreover, this is not just an artifact of how mathematicians define improper integrals, but has real consequences. Recall that if we think of the probability distribution as a mass distribution, then the mean is center of mass, the balance point, the point where the moment (in the sense of physics) to the right is balanced by the moment to the left. But as the proof of the last result shows, the moments to the right and to the left at any point $a \in \mathbb{R}$ are infinite. In this sense, 0 is no more important than any other $a \in \mathbb{R}$. Finally, if you are not convinced by the argument from physics, the next exercise may convince you that the law of large numbers fails as well.

Open the special distribution simulator and select the Cauchy distribution. Keep the default parameter values, which give the standard Cauchy distribution. Run the simulation 1000 times and note the behavior of the sample mean.

Earlier we noted some superficial similarities between the standard Cauchy distribution and the standard normal distribution (unimodal, symmetric about 0). But clearly there are huge quantitative differences. The Cauchy distribution is a *heavy tailed* distribution because the probability density function $g(x)$ decreases at a *polynomial rate* as $x \rightarrow \infty$ and $x \rightarrow -\infty$, as opposed to an *exponential rate*. This is yet another way to understand why the expected value does not exist.

In terms of the higher moments, $\mathbb{E}(X^n)$ does not exist if n is odd, and is ∞ if n is even. It follows that the moment generating function $m(t) = \mathbb{E}(e^{tX})$ cannot be finite in an interval about 0. In fact, $m(t) = \infty$ for every $t \neq 0$, so this generating function is of no use to us. But *every* distribution on \mathbb{R} has a characteristic function, and for the Cauchy distribution, this generating function will be quite useful.

X has characteristic function χ_0 given by $\chi_0(t) = \exp(-|t|)$ for $t \in \mathbb{R}$.

Proof

By definition,

$$\chi_0(t) = \mathbb{E}(e^{itX}) = \int_{-\infty}^{\infty} e^{itx} \frac{1}{\pi(1+x^2)} dx \quad (5.32.6)$$

We will compute this integral by evaluating a related contour integral in the complex plane using, appropriately enough, *Cauchy's integral formula* (named for you know who).

Suppose first that $t \geq 0$. For $r > 1$, let Γ_r denote the curve in the complex plane consisting of the line segment L_r on the x -axis from $-r$ to r and the upper half circle C_r of radius r centered at the origin. We give Γ_r the usual counter-clockwise orientation. On the one hand we have

$$\int_{\Gamma_r} \frac{e^{itz}}{\pi(1+z^2)} dz = \int_{L_r} \frac{e^{itz}}{\pi(1+z^2)} dz + \int_{C_r} \frac{e^{itz}}{\pi(1+z^2)} dz \quad (5.32.7)$$

On L_r , $z = x$ and $dz = dx$ so

$$\int_{L_r} \frac{e^{itz}}{\pi(1+z^2)} dz = \int_{-r}^r \frac{e^{itx}}{\pi(1+x^2)} dx \quad (5.32.8)$$

On C_r , let $z = x + iy$. Then $e^{itz} = e^{-ty+itx} = e^{-ty} [\cos(tx) + i \sin(tx)]$. Since $y \geq 0$ on C_r and $t \geq 0$, we have $|e^{itz}| \leq 1$. Also, $\left| \frac{1}{1+z^2} \right| \leq \frac{1}{r^2-1}$ on C_r . It follows that

$$\left| \int_{C_r} \frac{e^{itz}}{\pi(1+z^2)} dz \right| \leq \frac{1}{\pi(r^2-1)} \pi r = \frac{r}{r^2-1} \rightarrow 0 \text{ as } r \rightarrow \infty \quad (5.32.9)$$

On the other hand, $e^{itz}/[\pi(1+z^2)]$ has one singularity inside Γ_r , at i . The residue is

$$\lim_{z \rightarrow i} (z-i) \frac{e^{itz}}{\pi(1+z^2)} = \lim_{z \rightarrow i} \frac{e^{itz}}{\pi(z+i)} = \frac{e^{-t}}{2\pi i} \quad (5.32.10)$$

Hence by Cauchy's integral formula,

$$\int_{\Gamma_r} \frac{e^{itz}}{\pi(1+z^2)} dz = 2\pi i \frac{e^{-t}}{2\pi i} = e^{-t} \quad (5.32.11)$$

. Putting the pieces together we have

$$e^{-t} = \int_{-r}^r \frac{e^{itx}}{\pi(1+x^2)} dx + \int_{C_r} \frac{e^{itz}}{\pi(1+z^2)} dz \quad (5.32.12)$$

Letting $r \rightarrow \infty$ gives

$$\int_{-\infty}^{\infty} \frac{e^{itx}}{\pi(1+x^2)} dx = e^{-t} \quad (5.32.13)$$

For $t < 0$, we can use the substitution $u = -x$ and our previous result to get

$$\int_{-\infty}^{\infty} \frac{e^{itx}}{\pi(1+x^2)} dx = \int_{-\infty}^{\infty} \frac{e^{i(-t)u}}{\pi(1+u^2)} du = e^t \quad (5.32.14)$$

Related Distributions

The standard Cauchy distribution a member of the Student t family of distributions.

The standard Cauchy distribution is the Student t distribution with one degree of freedom.

Proof

The Student t distribution with one degree of freedom has PDF g given by

$$g(t) = \frac{\Gamma(1)}{\sqrt{\pi}\Gamma(1/2)} (1+t^2)^{-1} = \frac{1}{\pi(1+t^2)}, \quad t \in \mathbb{R} \quad (5.32.15)$$

which is the standard Cauchy PDF.

The standard Cauchy distribution also arises naturally as the ratio of independent standard normal variables.

Suppose that Z and W are independent random variables, each with the standard normal distribution. Then $X = Z/W$ has the standard Cauchy distribution. Equivalently, the standard Cauchy distribution is the Student t distribution with 1 degree of freedom.

Proof

By definition, W^2 has the chi-square distribution with 1 degree of freedom, and is independent of Z . Hence, also by definition, $X = Z/\sqrt{W^2} = Z/W$ has the Student t distribution with 1 degree of freedom, so the theorem follows from the [previous result](#).

If X has the standard Cauchy distribution, then so does $Y = 1/X$

Proof

This is a corollary of the [previous result](#). Suppose that Z and W are independent variables, each with the standard normal distribution. Then $X = Z/W$ has the standard Cauchy distribution. But then $1/X = W/Z$ also has the standard Cauchy distribution.

The standard Cauchy distribution has the usual connections to the standard uniform distribution via the [distribution function](#) and the [quantile function](#) computed above.

The standard Cauchy distribution and the standard uniform distribution are related as follows:

1. If U has the standard uniform distribution then $X = G^{-1}(U) = \tan\left[\pi\left(U - \frac{1}{2}\right)\right]$ has the standard Cauchy distribution.
2. If X has the standard Cauchy distribution then $U = G(X) = \frac{1}{2} + \frac{1}{\pi} \arctan(X)$ has the standard uniform distribution.

Proof

Recall that if U has the standard uniform distribution, then $G^{-1}(U)$ has distribution function G . Conversely, if X has distribution function G , then since G is strictly increasing, $G(X)$ has the standard uniform distribution.

Since the quantile function has a simple, closed form, it's easy to simulate the standard Cauchy distribution using the random quantile method.

Open the random quantile experiment and select the Cauchy distribution. Keep the default parameter values and note again the shape and location of the distribution and probability density functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function. Note the behavior of the empirical mean and standard deviation.

For the Cauchy distribution, the random quantile method has a nice physical interpretation. Suppose that a light source is 1 unit away from position 0 of an infinite, straight wall. We shine the light at the wall at an angle Θ (to the perpendicular) that is uniformly distributed on the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$. Then the position $X = \tan \Theta$ of the light beam on the wall has the standard Cauchy distribution. Note that this follows since Θ has the same distribution as $\pi(U - \frac{1}{2})$ where U has the standard uniform distribution.

Open the Cauchy experiment and keep the default parameter values.

1. Run the experiment in single-step mode a few times, to make sure that you understand the experiment.
2. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function. Note the behavior of the empirical mean and standard deviation.

The General Cauchy Distribution

Like so many other “standard” distributions, the Cauchy distribution is generalized by adding location and scale parameters. Most of the results in this subsection follow immediately from results for the standard Cauchy distribution above and general results for location scale families.

Suppose that Z has the standard Cauchy distribution and that $a \in \mathbb{R}$ and $b \in (0, \infty)$. Then $X = a + bZ$ has the *Cauchy distribution* with *location parameter* a and *scale parameter* b .

Distribution Functions

Suppose that X has the Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$.

X has probability density function f given by

$$f(x) = \frac{b}{\pi[b^2 + (x-a)^2]}, \quad x \in \mathbb{R} \quad (5.32.16)$$

1. f is symmetric about $x = a$.
2. f increases and then decreases, with mode $x = a$.
3. f is concave upward, then downward, then upward again, with inflection points at $x = a \pm \frac{1}{\sqrt{3}}b$.
4. $f(x) \rightarrow 0$ as $x \rightarrow \infty$ and as $x \rightarrow -\infty$.

Proof

Recall that

$$f(x) = \frac{1}{b}g\left(\frac{x-a}{b}\right) \quad (5.32.17)$$

where g is the [standard Cauchy PDF](#).

Open the special distribution simulator and select the Cauchy distribution. Vary the parameters and note the location and shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has distribution function F given by

$$F(x) = \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{x-a}{b}\right), \quad x \in \mathbb{R} \quad (5.32.18)$$

Proof

Recall that

$$F(x) = G\left(\frac{x-a}{b}\right) \quad (5.32.19)$$

where G is the [standard Cauchy CDF](#).

X has quantile function F^{-1} given by

$$F^{-1}(p) = a + b \tan\left[\pi\left(p - \frac{1}{2}\right)\right], \quad p \in (0, 1) \quad (5.32.20)$$

In particular,

1. The first quartile is $F^{-1}\left(\frac{1}{4}\right) = a - b$.
2. The median is $F^{-1}\left(\frac{1}{2}\right) = a$.
3. The third quartile is $F^{-1}\left(\frac{3}{4}\right) = a + b$.

Proof

Recall that $F^{-1}(p) = a + bG^{-1}(p)$ where G^{-1} is the [standard Cauchy quantile function](#).

Open the special distribution calculator and select the Cauchy distribution. Vary the parameters and note the shape and location of the distribution and probability density functions. Compute a few values of the distribution and quantile functions.

Moments

Suppose again that X has the Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. Since the mean and other moments of the standard Cauchy distribution do not exist, they don't exist for the general Cauchy distribution either.

Open the special distribution simulator and select the Cauchy distribution. For selected values of the parameters, run the simulation 1000 times and note the behavior of the sample mean.

But of course the characteristic function of the Cauchy distribution exists and is easy to obtain from the characteristic function of the standard distribution.

X has characteristic function χ given by $\chi(t) = \exp(ait - b|t|)$ for $t \in \mathbb{R}$.

Proof

Recall that $\chi(t) = e^{ita} \chi_0(bt)$ where χ_0 is the [standard Cauchy characteristic function](#).

Related Distributions

Like all location-scale families, the general Cauchy distribution is closed under location-scale transformations.

Suppose that X has the Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, and that $c \in \mathbb{R}$ and $d \in (0, \infty)$. Then $Y = c + dX$ has the Cauchy distribution with location parameter $c + da$ and scale parameter bd .

Proof

Once again, we give the standard proof. By [definition](#) we can take $X = a + bZ$ where Z has the standard Cauchy distribution. But then $Y = c + dX = (c + ad) + (bd)Z$.

Much more interesting is the fact that the Cauchy family is closed under sums of independent variables. In fact, this is the main reason that the generalization to a location-scale family is justified.

Suppose that X_i has the Cauchy distribution with location parameter $a_i \in \mathbb{R}$ and scale parameter $b_i \in (0, \infty)$ for $i \in \{1, 2\}$, and that X_1 and X_2 are independent. Then $Y = X_1 + X_2$ has the Cauchy distribution with location parameter $a_1 + a_2$ and scale parameter $b_1 + b_2$.

Proof

This follows easily from the [characteristic function](#). Let χ_i denote the characteristic function of X_i for $i = 1, 2$ and χ the characteristic function of Y . Then

$$\chi(t) = \chi_1(t)\chi_2(t) = \exp(a_1 it - b_1 |t|) \exp(a_2 it - b_2 |t|) = \exp[(a_1 + a_2) it - (b_1 + b_2) |t|] \quad (5.32.21)$$

As a corollary, the Cauchy distribution is stable, with index $\alpha = 1$:

If $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent variables, each with the Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$, then $X_1 + X_2 + \dots + X_n$ has the Cauchy distribution with location parameter na and scale parameter nb .

Another corollary is the strange property that the sample mean of a random sample from a Cauchy distribution has that same Cauchy distribution. No wonder the expected value does not exist!

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent random variables, each with the Cauchy distribution with location parameter $a \in \mathbb{R}$ and scale parameter $b \in (0, \infty)$. (That is, \mathbf{X} is a random sample of size n from the Cauchy distribution.) Then the sample mean $M = \frac{1}{n} \sum_{i=1}^n X_i$ also has the Cauchy distribution with location parameter a and scale parameter b .

Proof

From the previous [stability result](#), $Y = \sum_{i=1}^n X_i$ has the Cauchy distribution with location parameter na and scale parameter nb . But then by the [scaling result](#), $M = Y/n$ has the Cauchy distribution with location parameter a and scale parameter b .

The next result shows explicitly that the Cauchy distribution is infinitely divisible. But of course, infinite divisibility is also a consequence of stability.

Suppose that $a \in \mathbb{R}$ and $b \in (0, \infty)$. For every $n \in \mathbb{N}_+$ the Cauchy distribution with location parameter a and scale parameter b is the distribution of the sum of n independent variables, each of which has the Cauchy distribution with location parameters a/n and scale parameter b/n .

Our next result is a very slight generalization of the [reciprocal result](#) above for the standard Cauchy distribution.

Suppose that X has the Cauchy distribution with location parameter 0 and scale parameter $b \in (0, \infty)$. Then $Y = 1/X$ has the Cauchy distribution with location parameter 0 and scale parameter $1/b$.

Proof

X has the same distribution as bZ where Z has the standard Cauchy distribution. Hence $\frac{1}{X}$ has the same distribution as $\frac{1}{b} \frac{1}{Z}$. But by the [result above](#), $\frac{1}{Z}$ also has the standard Cauchy distribution, so $\frac{1}{b} \frac{1}{Z}$ has the Cauchy distribution with location parameter 0 and scale parameter $1/b$.

As with its standard cousin, the general Cauchy distribution has simple connections with the standard uniform distribution via the [distribution function](#) and [quantile function](#) computed above, and in particular, can be simulated via the random quantile method.

Suppose that $a \in \mathbb{R}$ and $b \in (0, \infty)$.

1. If U has the standard uniform distribution, then $X = F^{-1}(U) = a + b \tan\left[\pi \left(U - \frac{1}{2}\right)\right]$ has the Cauchy distribution with location parameter a and scale parameter b .
2. If X has the Cauchy distribution with location parameter a and scale parameter b , then $U = F(X) = \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{X-a}{b}\right)$ has the standard uniform distribution.

Open the random quantile experiment and select the Cauchy distribution. Vary the parameters and note again the shape and location of the distribution and probability density functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function. Note the behavior of the empirical mean and standard deviation.

As before, the random quantile method has a nice physical interpretation. Suppose that a light source is b units away from position a of an infinite, straight wall. We shine the light at the wall at an angle Θ (to the perpendicular) that is uniformly distributed on the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$. Then the position $X = a + b \tan \Theta$ of the light beam on the wall has the Cauchy distribution with location parameter a and scale parameter b .

Open the Cauchy experiment. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function. Note the behavior of the empirical mean and standard deviation.

This page titled [5.32: The Cauchy Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.33: The Exponential-Logarithmic Distribution

The *exponential-logarithmic distribution* arises when the rate parameter of the exponential distribution is randomized by the logarithmic distribution. The exponential-logarithmic distribution has applications in reliability theory in the context of devices or organisms that improve with age, due to hardening or immunity.

The Standard Exponential-Logarithmic Distribution

Distribution Functions

The *standard exponential-logarithmic distribution* with shape parameter $p \in (0, 1)$ is a continuous distribution on $[0, \infty)$ with probability density function g given by

$$g(x) = -\frac{(1-p)e^{-x}}{\ln(p)[1-(1-p)e^{-x}]}, \quad x \in [0, \infty) \quad (5.33.1)$$

1. g is decreasing on $[0, \infty)$ with mode $x = 0$.
2. g is concave upward on $[0, \infty)$.

Proof

Substituting $u = (1-p)e^{-x}$, $du = -(1-p)e^{-x}dx$ gives

$$\int_0^\infty \frac{(1-p)e^{-x}}{1-(1-p)e^{-x}} dx = \int_0^{1-p} \frac{du}{1-u} = -\ln(p) \quad (5.33.2)$$

so it follows that g is a PDF. For the shape of the graph of g note that

$$g'(x) = \frac{(1-p)e^{-x}}{\ln(p)[1-(1-p)e^{-x}]^2}, \quad x \in [0, \infty) \quad (5.33.3)$$

$$g''(x) = -\frac{(1-p)e^{-x}[1+(1-p)e^{-x}]}{\ln(p)[1-(1-p)e^{-x}]^3}, \quad x \in [0, \infty) \quad (5.33.4)$$

Open the special distribution simulator and select the exponential-logarithmic distribution. Vary the shape parameter and note the shape of the probability density function. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function G is given by

$$G(x) = 1 - \frac{\ln[1-(1-p)e^{-x}]}{\ln(p)}, \quad x \in [0, \infty) \quad (5.33.5)$$

Proof

This follows from the same integral substitution used in the [previous proof](#).

The quantile function G^{-1} is given by

$$G^{-1}(u) = \ln\left(\frac{1-p}{1-p^{1-u}}\right) = \ln(1-p) - \ln(1-p^{1-u}), \quad u \in [0, 1) \quad (5.33.6)$$

1. The first quartile is $q_1 = \ln(1-p) - \ln(1-p^{3/4})$.
2. The median is $q_2 = \ln(1-p) - \ln(1-p^{1/2}) = \ln(1+\sqrt{p})$.
3. The third quartile is $q_3 = \ln(1-p) - \ln(1-p^{1/4})$.

Proof

The formula for G^{-1} follows from the [distribution function](#) by solving $u = G(x)$ for x in terms of u .

Open the special distribution calculator and select the exponential-logarithmic distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the shape parameter, computer a few values of the distribution function and the quantile function.

The reliability function G^c given by

$$G^c(x) = \frac{\ln[1 - (1-p)e^{-x}]}{\ln(p)}, \quad x \in [0, \infty) \quad (5.33.7)$$

Proof

This follows trivially from the [distribution function](#) since $G^c = 1 - G$.

The standard exponential-logarithmic distribution has decreasing failure rate.

The failure rate function r is given by

$$r(x) = -\frac{(1-p)e^{-x}}{[1 - (1-p)e^{-x}] \ln[1 - (1-p)e^{-x}]}, \quad x \in (0, \infty) \quad (5.33.8)$$

1. r is decreasing on $[0, \infty)$.
2. r is concave upward on $[0, \infty)$.

Proof

Recall that $r(x) = g(x)/G^c(x)$ so the formula follows from the [probability density function](#) and the [distribution function](#) given above.

The Polylogarithm

The moments of the standard exponential-logarithmic distribution cannot be expressed in terms of the usual elementary functions, but can be expressed in terms of a special function known as the *polylogarithm*.

The *polylogarithm* of order $s \in \mathbb{R}$ is defined by

$$\text{Li}_s(x) = \sum_{k=1}^{\infty} \frac{x^k}{k^s}, \quad x \in (-1, 1) \quad (5.33.9)$$

The polylogarithm is a power series in x with radius of convergence is 1 for each $s \in \mathbb{R}$.

Proof

To show that the radius of convergence is 1, we use the *ratio test* from calculus. For $s \in \mathbb{R}$,

$$\frac{|x|^{k+1}/(k+1)^s}{|x|^k/k^s} = |x| \left(\frac{k}{k+1} \right)^s \rightarrow |x| \text{ as } k \rightarrow \infty \quad (5.33.10)$$

Hence the series converges absolutely for $|x| < 1$ and diverges for $|x| > 1$.

In this section, we are only interested in nonnegative integer orders, but the polylogarithm will show up again, for non-integer orders, in the study of the zeta distribution.

The polylogarithm functions of orders 0, 1, 2, and 3.

1. The polylogarithm of order 0 is

$$\text{Li}_0(x) = \sum_{k=1}^{\infty} x^k = \frac{x}{1-x}, \quad x \in (-1, 1) \quad (5.33.11)$$

2. The polylogarithm of order 1 is

$$\text{Li}_1(x) = \sum_{k=1}^{\infty} \frac{x^k}{k} = -\ln(1-x), \quad x \in (-1, 1) \quad (5.33.12)$$

3. The polylogarithm of order 2 is known as the *dilogarithm*
4. The polylogarithm of order 3 is known as the *trilogarithm*.

Thus, the polylogarithm of order 0 is a simple geometric series, and the polylogarithm of order 1 is the standard power series for the natural logarithm. Note that the probability density function of X can be written in terms of the polylogarithms of orders 0 and 1:

$$g(x) = -\frac{\text{Li}_0[(1-p)e^{-x}]}{\ln(p)} = \frac{\text{Li}_0[(1-p)e^{-x}]}{\text{Li}_1(1-p)}, \quad x \in [0, \infty) \quad (5.33.13)$$

The most important property of the polylogarithm is given in the following theorem:

The polylogarithm satisfies the following recursive integral formula:

$$\text{Li}_{s+1}(x) = \int_0^x \frac{\text{Li}_s(t)}{t} dt; \quad s \in \mathbb{R}, \quad x \in (-1, 1) \quad (5.33.14)$$

Equivalently, $x \text{Li}'_{s+1}(x) = \text{Li}_s(x)$ for $x \in (-1, 1)$ and $s \in \mathbb{R}$.

Proof

Recall that a power series may be integrated term by term, and the integrated series has the same radius of convergence. Hence for $s \in \mathbb{R}$,

$$\int_0^x \frac{\text{Li}_s(t)}{t} dt = \sum_{k=1}^{\infty} \int_0^x \frac{t^{k-1}}{k^s} dt = \sum_{k=1}^{\infty} \frac{x^k}{k^{s+1}} = \text{Li}_{s+1}(x), \quad x \in (-1, 1) \quad (5.33.15)$$

When $s > 1$, the polylogarithm series converges at $x = 1$ also, and

$$\text{Li}_s(1) = \zeta(s) = \sum_{k=1}^{\infty} \frac{1}{k^s} \quad (5.33.16)$$

where ζ is the *Riemann zeta function*, named for Georg Riemann. The polylogarithm can be extended to complex orders and defined for complex z with $|z| < 1$, but the simpler version suffices for our work here.

Moments

We assume that X has the standard exponential-logarithmic distribution with shape parameter $p \in (0, 1)$.

The moments of X (about 0) are

$$\mathbb{E}(X^n) = -n! \frac{\text{Li}_{n+1}(1-p)}{\ln(p)} = n! \frac{\text{Li}_{n+1}(1-p)}{\text{Li}_1(1-p)}, \quad n \in \mathbb{N} \quad (5.33.17)$$

1. $\mathbb{E}(X^n) \rightarrow 0$ as $p \downarrow 0$
2. $\mathbb{E}(X^n) \rightarrow n!$ as $p \uparrow 1$

Proof

As noted earlier in the discussion of the polylogarithm, the PDF of X can be written as

$$g(x) = -\frac{1}{\ln(p)} \sum_{k=1}^{\infty} (1-p)^k e^{-kx}, \quad x \in [0, \infty) \quad (5.33.18)$$

Hence

$$\mathbb{E}(X^n) = -\frac{1}{\ln(p)} \int_0^{\infty} \sum_{k=1}^{\infty} (1-p)^k x^n e^{-kx} dx = -\frac{1}{\ln(p)} \sum_{k=1}^{\infty} (1-p)^k \int_0^{\infty} x^n e^{-kx} dx \quad (5.33.19)$$

But $\int_0^{\infty} x^n e^{-kx} dx = n! / k^{n+1}$ and hence

$$\mathbb{E}(X^n) = -\frac{1}{\ln(p)} n! \sum_{k=1}^{\infty} \frac{(1-p)^k}{k^{n+1}} = -n! \frac{\text{Li}_{n+1}(1-p)}{\ln(p)} \quad (5.33.20)$$

1. As $p \downarrow 0$, the numerator in the last expression for $\mathbb{E}(X^n)$ converges to $n! \zeta(n+1)$ while the denominator diverges to ∞ .
2. As $p \uparrow 1$, the expression for $\mathbb{E}(X^n)$ has the indeterminate form $\frac{0}{0}$. An application of L'Hospital's rule and the [derivative rule](#) above gives

$$\lim_{p \uparrow 1} \mathbb{E}(X^n) = \lim_{p \uparrow 1} n! p \frac{\text{Li}_n(1-p)}{1-p} \quad (5.33.21)$$

But from the [series definition of the polylogarithm](#), $\text{Li}_n(x)/x \rightarrow 1$ as $x \rightarrow 0$.

We will get some additional insight into the asymptotics below when we consider the [limiting distribution](#) as $p \downarrow 0$ and $p \uparrow 1$. The mean and variance of the standard exponential logarithmic distribution follow easily from the general moment formula.

The mean and variance of X are

1. $\mathbb{E}(X) = -\text{Li}_2(1-p)/\ln(p)$
2. $\text{var}(X) = -2\text{Li}_3(1-p)/\ln(p) - [\text{Li}_2(1-p)/\ln(p)]^2$

From the asymptotics of the [general moments](#), note that $\mathbb{E}(X) \rightarrow 0$ and $\text{var}(X) \rightarrow 0$ as $p \downarrow 0$, and $\mathbb{E}(X) \rightarrow 1$ and $\text{var}(X) \rightarrow 1$ as $p \uparrow 1$.

Open the special distribution simulator and select the exponential-logarithmic distribution. Vary the shape parameter and note the size and location of the mean \pm standard deviation bar. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

The standard exponential-logarithmic distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) computed above.

Suppose that $p \in (0, 1)$.

1. If U has the standard uniform distribution then

$$X = \ln\left(\frac{1-p}{1-p^U}\right) = \ln(1-p) - \ln(1-p^U) \quad (5.33.22)$$

has the standard exponential-logarithmic distribution with shape parameter p .

2. If X has the standard exponential-logarithmic distribution with shape parameter p then

$$U = \frac{\ln[1 - (1-p)e^{-X}]}{\ln(p)} \quad (5.33.23)$$

has the standard uniform distribution.

Proof

1. Recall that if U has the standard uniform distribution, then $G^{-1}(U)$ has the exponential-logarithmic distribution with shape parameter p . But $1 - U$ also has the standard uniform distribution and hence $X = G^{-1}(1 - U)$ also has the exponential-logarithmic distribution with shape parameter p .
2. Similarly, if X has the exponential-logarithmic distribution with shape parameter p then $G(X)$ has the standard uniform distribution. Hence $U = 1 - G(X)$ also has the standard uniform distribution.

Since the quantile function of the basic exponential-logarithmic distribution has a simple closed form, the distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the exponential-logarithmic distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

As the name suggests, the basic exponential-logarithmic distribution arises from the exponential distribution and the logarithmic distribution via a certain type of randomization.

Suppose that $\mathbf{T} = (T_1, T_2, \dots)$ is a sequence of independent random variables, each with the standard exponential distribution. Suppose also that N has the logarithmic distribution with parameter $1 - p \in (0, 1)$ and is independent of \mathbf{T} . Then $X = \min\{T_1, T_2, \dots, T_N\}$ has the basic exponential-logarithmic distribution with shape parameter p .

Proof

It's best to work with reliability functions. For $n \in \mathbb{N}_+$, $\min\{T_1, T_2, \dots, T_n\}$ has the exponential distribution with rate parameter n , and hence $\mathbb{P}(\min\{T_1, T_2, \dots, T_n\} > x) = e^{-nx}$ for $x \in [0, \infty)$. Recall also that

$$\mathbb{P}(N = n) = -\frac{(1-p)^n}{n \ln(p)}, \quad n \in \mathbb{N}_+ \quad (5.33.24)$$

Hence, using the [polylogarithm of order 1](#) (the standard power series for the logarithm),

$$\mathbb{P}(X > x) = \mathbb{E}[\mathbb{P}(X > x \mid N)] = -\frac{1}{\ln(p)} \sum_{n=1}^{\infty} e^{-nx} \frac{(1-p)^n}{n} = -\frac{1}{\ln(p)} \sum_{n=1}^{\infty} \frac{[e^{-x}(1-p)]^n}{n} = \frac{\ln[1 - e^{-x}(1-p)]}{\ln(p)} \quad (5.33.25)$$

As a function of x , this is the reliability function of the exponential-logarithmic distribution with shape parameter p .

Also of interest, of course, are the limiting distributions of the standard exponential-logarithmic distribution as $p \rightarrow 0$ and as $p \rightarrow 1$.

The standard exponential-logarithmic distribution with shape parameter $p \in (0, 1)$ converges to

1. Point mass at 0 as $p \rightarrow 0$.
2. The standard exponential distribution as $p \rightarrow 1$.

Proof

It's slightly easier to work with the [reliability function](#) G^c rather than the ordinary (left) distribution function G .

1. Note that $G^c(0) = 1$ for every $p \in (0, 1)$. On the other hand, if $x > 0$ then $G^c(x) \rightarrow 0$ as $p \rightarrow 0$.
2. $G^c(x)$ has the indeterminate form $\frac{0}{0}$ as $p \rightarrow 1$. An application of L'Hospital's rule shows that

$$\lim_{p \rightarrow 1} G^c(x) = \lim_{p \rightarrow 1} \frac{pe^{-x}}{1 - (1-p)e^{-x}} = e^{-x}, \quad x \in [0, \infty) \quad (5.33.26)$$

As a function of x , this is the reliability function of the standard exponential distribution.

The General Exponential-Logarithmic Distribution

The standard exponential-logarithmic distribution is generalized, like so many distributions on $[0, \infty)$, by adding a scale parameter.

Suppose that Z has the standard exponential-logarithmic distribution with shape parameter $p \in (0, 1)$. If $b \in (0, \infty)$, then $X = bZ$ has the *exponential-logarithmic distribution* with shape parameter p and scale parameter b .

Using the same terminology as the exponential distribution, $1/b$ is called the *rate parameter*.

Distribution Functions

Suppose that X has the exponential-logarithmic distribution with shape parameter $p \in (0, 1)$ and scale parameter $b \in (0, \infty)$.

X has probability density function f given by

$$f(x) = -\frac{(1-p)e^{-x/b}}{b \ln(p)[1 - (1-p)e^{-x/b}]}, \quad x \in [0, \infty) \quad (5.33.27)$$

1. f is decreasing on $[0, \infty)$ with mode $x = 0$.
2. f is concave upward on $[0, \infty)$.

Proof

Recall that $f(x) = \frac{1}{b}g\left(\frac{x}{b}\right)$ for $x \in [0, \infty)$ where g is the [PDF of the standard distribution](#).

Open the special distribution simulator and select the exponential-logarithmic distribution. Vary the shape and scale parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has distribution function F given by

$$F(x) = 1 - \frac{\ln[1 - (1-p)e^{-x/b}]}{\ln(p)}, \quad x \in [0, \infty) \quad (5.33.28)$$

Proof

Recall that $F(x) = G(x/b)$ for $x \in [0, \infty)$ where G is the [CDF of the standard distribution](#).

X has quantile function F^{-1} given by

$$F^{-1}(u) = b \ln\left(\frac{1-p}{1-p^{1-u}}\right) = b [\ln(1-p) - \ln(1-p^{1-u})], \quad u \in [0, 1) \quad (5.33.29)$$

1. The first quartile is $q_1 = b [\ln(1-p) - \ln(1-p^{3/4})]$.
2. The median is $q_2 = b [\ln(1-p) - \ln(1-p^{1/2})] = b \ln(1 + \sqrt{p})$.
3. The third quartile is $q_3 = b [\ln(1-p) - \ln(1-p^{1/4})]$.

Proof

Recall that $F^{-1}(u) = bG^{-1}(u)$ where G^{-1} is the [quantile function of the standard distribution](#).

Open the special distribution calculator and select the exponential-logarithmic distribution. Vary the shape and scale parameter and note the shape and location of the probability density and distribution functions. For selected values of the parameters, compute a few values of the distribution function and the quantile function.

X has reliability function F^c given by

$$F^c(x) = \frac{\ln[1 - (1-p)e^{-x/b}]}{\ln(p)}, \quad x \in [0, \infty) \quad (5.33.30)$$

Proof

This follows trivially from the [distribution function](#) since $F^c = 1 - F$.

The exponential-logarithmic distribution has decreasing failure rate.

The failure rate function R of X is given by.

$$R(x) = -\frac{(1-p)e^{-x/b}}{b[1 - (1-p)e^{-x/b}] \ln[1 - (1-p)e^{-x/b}]}, \quad x \in [0, \infty) \quad (5.33.31)$$

1. R is decreasing on $[0, \infty)$.
2. R is concave upward on $[0, \infty)$.

Proof

Recall that $R(x) = \frac{1}{b}r\left(\frac{x}{b}\right)$ for $x \in [0, \infty)$, where r is the [failure rate function of the standard distribution](#). Alternately, $R(x) = f(x)/F^c(x)$.

Moments

Suppose again that X has the exponential-logarithmic distribution with shape parameter $p \in (0, 1)$ and scale parameter $b \in (0, \infty)$. The moments of X can be computed easily from the representation $X = bZ$ where Z has the basic exponential-logarithmic distribution.

The moments of X (about 0) are

$$\mathbb{E}(X^n) = -b^n n! \frac{\text{Li}_{n+1}(1-p)}{\ln(p)}, \quad n \in \mathbb{N} \quad (5.33.32)$$

1. $\mathbb{E}(X^n) \rightarrow 0$ as $p \downarrow 0$
2. $\mathbb{E}(X^n) \rightarrow b^n n!$ as $p \uparrow 1$

Proof

These results follow from basic properties of expected value and the corresponding [results for the standard distribution](#). We can write $X = bZ$ where Z has the standard exponential-logarithmic distribution with shape parameter p . Hence $\mathbb{E}(X^n) = b^n \mathbb{E}(Z^n)$.

The mean and variance of X are

1. $\mathbb{E}(X) = -b \text{Li}_2(1-p) / \ln(p)$
2. $\text{var}(X) = b^2 \left(-2 \text{Li}_3(1-p) / \ln(p) - [\text{Li}_2(1-p) / \ln(p)]^2 \right)$

From the [general moment results](#), note that $\mathbb{E}(X) \rightarrow 0$ and $\text{var}(X) \rightarrow 0$ as $p \downarrow 0$, while $\mathbb{E}(X) \rightarrow b$ and $\text{var}(X) \rightarrow b^2$ as $p \uparrow 1$.

Open the special distribution simulator and select the exponential-logarithmic distribution. Vary the shape and scale parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

Since the exponential-logarithmic distribution is a scale family for each value of the shape parameter, it is trivially closed under scale transformations.

Suppose that X has the exponential-logarithmic distribution with shape parameter $p \in (0, 1)$ and scale parameter $b \in (0, \infty)$. If $c \in (0, \infty)$, then $Y = cX$ has the exponential-logarithmic distribution with shape parameter p and scale parameter bc .

Proof

By [definition](#), we can take $X = bZ$ where Z has the standard exponential-logarithmic distribution with shape parameter p . But then $Y = cX = (bc)Z$.

Once again, the exponential-logarithmic distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) computed above.

Suppose that $p \in (0, 1)$ and $b \in (0, \infty)$.

1. If U has the standard exponential distribution then

$$X = b \left[\ln \left(\frac{1-p}{1-p^U} \right) \right] = b [\ln(1-p) - \ln(1-p^U)] \quad (5.33.33)$$

has the exponential-logarithmic distribution with shape parameter p and scale parameter b .

2. If X has the exponential-logarithmic distribution with shape parameter p and scale parameter b , then

$$U = \frac{\ln[1 - (1-p)e^{-X/b}]}{\ln(p)} \quad (5.33.34)$$

has the standard uniform distribution.

Proof

These results follow from the representation $X = bZ$, where Z has the standard exponential-logarithmic distribution with shape parameter p , and the corresponding [result for \$Z\$](#) .

Again, since the quantile function of the exponential-logarithmic distribution has a simple closed form, the distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the exponential-logarithmic distribution. Vary the shape and scale parameters and note the shape and location of the distribution and probability density functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Suppose that $\mathbf{T} = (T_1, T_2, \dots)$ is a sequence of independent random variables, each with the exponential distribution with scale parameter $b \in (0, \infty)$. Suppose also that N has the logarithmic distribution with parameter $1-p \in (0, 1)$ and is independent of \mathbf{T} . Then $X = \min\{T_1, T_2, \dots, T_N\}$ has the exponential-logarithmic distribution with shape parameter p and scale parameter b .

Proof

Note that $V_i = T_i/b$ has the standard exponential distribution. Hence by the [corresponding result above](#), $Z = \min\{V_1, V_2, \dots, V_N\}$ has the basic exponential-logarithmic distribution with shape parameter p . Hence $X = bZ$ has the exponential-logarithmic distribution with shape parameter p and scale parameter b .

The limiting distributions as $p \downarrow 0$ and as $p \uparrow 1$ also follow easily from the corresponding results for the standard case.

For fixed $b \in (0, \infty)$, the exponential-logarithmic distribution with shape parameter $p \in (0, 1)$ and scale parameter b converges to

1. Point mass at 0 as $p \downarrow 0$.
2. The exponential distribution with scale parameter b as $p \uparrow 1$.

Proof

Suppose that X has the exponential-logarithmic distribution with shape parameter p and scale parameter b , so that $X = bZ$ where Z has the standard exponential-logarithmic distribution with shape parameter p . Using the [corresponding result above](#),

1. The distribution of Z converges to point mass at 0 as $p \downarrow 0$ and hence so does the distribution of X .
2. The distribution of Z converges to the standard exponential distribution as $p \uparrow 1$ and hence the the distribution of X converges to the exponential distribution with scale parameter b .

This page titled [5.33: The Exponential-Logarithmic Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.34: The Gompertz Distribution

The *Gompertz distribution*, named for Benjamin Gompertz, is a continuous probability distribution on $[0, \infty)$ that has exponentially increasing failure rate. Unfortunately, the death rate of adult humans increases exponentially, so the Gompertz distribution is widely used in actuarial science.

The Basic Gompertz Distribution

Distribution Functions

We will start by giving the reliability function, since most applications of the Gompertz distribution deal with mortality.

The *basic Gompertz distribution* with shape parameter $a \in (0, \infty)$ is a continuous distribution on $[0, \infty)$ with reliability function G^c given by

$$G^c(x) = \exp[-a(e^x - 1)], \quad x \in [0, \infty) \quad (5.34.1)$$

The special case $a = 1$ gives the *standard Gompertz distribution*.

Proof

Note that G^c is continuous and decreasing on $[0, \infty)$ with $G^c(0) = 1$ and $G^c(x) \rightarrow 0$ as $x \rightarrow \infty$.

The distribution function G is given by

$$G(x) = 1 - \exp[-a(e^x - 1)], \quad x \in [0, \infty) \quad (5.34.2)$$

Proof

This follows trivially from the reliability function [reliability function](#), since $G = 1 - G^c$.

The quantile function G^{-1} is given by

$$G^{-1}(p) = \ln \left[1 - \frac{1}{a} \ln(1 - p) \right], \quad p \in [0, 1) \quad (5.34.3)$$

1. The first quartile is $q_1 = \ln[1 + (\ln 4 - \ln 3)/a]$.
2. The median is $q_2 = \ln(1 + \ln 2/a)$.
3. The third quartile is $q_3 = \ln(1 + \ln 4/a)$.

Proof

The formula for G^{-1} follows from the [distribution function](#) by solving $p = G(x)$ for x in terms of p .

For the standard Gompertz distribution ($a = 1$), the first quartile is $q_1 = \ln[1 + (\ln 4 - \ln 3)] \approx 0.2529$, the median is $q_2 = \ln(1 + \ln 2) \approx 0.5266$, and the third quartile is $q_3 = \ln(1 + \ln 4) \approx 0.8697$.

Open the special distribution calculator and select the Gompertz distribution. Vary the shape parameter and note the shape of the distribution function. For selected values of the shape parameter, compute a few values of the distribution function and the quantile function.

The probability density function g is given by

$$g(x) = ae^x \exp[-a(e^x - 1)], \quad x \in [0, \infty) \quad (5.34.4)$$

1. If $a < 1$ then g is increasing and then decreasing with mode $x = -\ln(a)$.
2. If $a \geq 1$ then g is decreasing with mode $x = 0$.
3. If $a < (3 - \sqrt{5})/2 \approx 0.382$ then g is concave up and then down then up again, with inflection points at $x = \ln[(3 \pm \sqrt{5})/2a]$.
4. If $(3 - \sqrt{5})/2 \leq a < (3 + \sqrt{5})/2 \approx 2.618$ then g is concave down and then up, with inflection point at $x = \ln[(3 + \sqrt{5})/2a]$.

5. If $a \geq (3 + \sqrt{5})/2$ then g is concave up.

Proof

The formula for g follows from the [distribution function](#) since $g = G'$ Parts (a)–(d) follow from

$$g'(x) = ae^x(1 - ae^x) \exp[-a(e^x - 1)] \quad (5.34.5)$$

$$g''(x) = ae^x(1 - 3ae^x + a^2e^{2x}) \exp[-a(e^x - 1)] \quad (5.34.6)$$

So for the standard Gompertz distribution ($a = 1$), the inflection point is $x = \ln(3 + \sqrt{5}) \approx 1.6555$.

Open the special distribution simulator and select the Gompertz distribution. Vary the shape parameter and note the shape of the probability density function. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

Finally, as promised, the Gompertz distribution has exponentially increasing failure rate.

The failure rate function r is given by $r(x) = ae^x$ for $x \in [0, \infty)$

Proof

Recall that the is $r(x) = g(x)/G^c(x)$ so the result follows from the [distribution function](#) and the [probability density function](#).

Moments

The moments of the basic Gompertz distribution cannot be given in simple closed form, but the mean and moment generating function can at least be expressed in terms of a special function known as the *exponential integral*. There are many variations on the exponential integral, but for our purposes, the following version is best:

The *exponential integral* with parameter $a \in (0, \infty)$ is the function $E_a : \mathbb{R} \rightarrow (0, \infty)$ defined by

$$E_a(t) = \int_1^\infty u^t e^{-au} du, \quad t \in \mathbb{R} \quad (5.34.7)$$

For the remainder of this discussion, we assume that X has the basic Gompertz distribution with shape parameter $a \in (0, \infty)$.

X has moment generating function m given by

$$m(t) = \mathbb{E}(e^{tX}) = ae^a E_a(t), \quad t \in \mathbb{R} \quad (5.34.8)$$

Proof

Using the substitution $u = e^x$ we have

$$m(t) = \int_0^\infty e^{tx} ae^x e^a \exp(-ae^x) dx = ae^a \int_1^\infty u^t e^{-au} du = ae^a E_a(t) \quad (5.34.9)$$

It follows that X has moments of all orders. Here is the mean:

X has mean $\mathbb{E}(X) = e^a E_a(-1)$.

Proof

First we use the substitution $y = e^x$ to get

$$\mathbb{E}(X) = \int_0^\infty x ae^x e^a \exp(-ae^x) dx = ae^a \int_1^\infty \ln(y) e^{-ay} dy \quad (5.34.10)$$

Next, integration by parts with $u = \ln y$, $dv = e^{-ay} dy$ gives

$$\mathbb{E}(X) = e^a \int_1^\infty \frac{1}{y} e^{-ay} dy = e^a E_a(-1) \quad (5.34.11)$$

If X has the standard Gompertz distribution, $\mathbb{E}(X) \approx 0.5963$.

Open the special distribution simulator and select the Gompertz distribution. Vary the shape parameter and note the size and location of the mean \pm standard deviation bar. For selected values of the parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

The basic Gompertz distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) computed above.

Suppose that $a \in (0, \infty)$.

1. If U has the standard uniform distribution then $X = \ln(1 - \frac{1}{a} \ln U)$ has the basic Gompertz distribution with shape parameter a .
2. If X has the basic Gompertz distribution with shape parameter a then $U = \exp[-a(e^X - 1)]$ has the standard uniform distribution.

Proof

1. Recall that if U has the standard uniform distribution, then $1 - U$ also has the standard uniform distribution, and hence $X = G^{-1}(1 - U)$ has the basic Gompertz distribution with shape parameter a .
2. If X has the basic Gompertz distribution with shape parameter a then $G(X)$ has the standard uniform distribution, and hence so does $U = 1 - G(X)$.

Since the quantile function of the basic Gompertz distribution has a simple closed form, the distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the Gompertz distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The basic Gompertz distribution also has simple connections to the exponential distribution.

Suppose that $a \in (0, \infty)$.

1. If X has the basic Gompertz distribution with shape parameter a , then $Y = e^X - 1$ has the exponential distribution with rate parameter a .
2. If Y has the exponential distribution with rate parameter a , then $X = \ln(Y + 1)$ has the Gompertz distribution with shape parameter a .

Proof

These results follow from the standard change of variables formula. The transformations, which are inverses of each other, are $y = e^x - 1$ and $x = \ln(y + 1)$ for $x, y \in [0, \infty)$. Let g and h denote PDFs of X and Y respectively.

1. We start with $g(x) = ae^x \exp[-a(e^x - 1)]$ for $x \in [0, \infty)$ and then

$$h(y) = g(x) \frac{dx}{dy} = a \exp[\ln(y + 1)] \exp\{-a[\exp(\ln(y + 1)) - 1]\} \frac{1}{y + 1} = ae^{-ay}, \quad y \in [0, \infty) \quad (5.34.12)$$

which is the PDF of the exponential distribution with rate parameter a .

2. We start with $h(y) = ae^{-ay}$ for $y \in [0, \infty)$ and then

$$g(x) = h(y) \frac{dy}{dx} = a \exp[-a(e^x - 1)] e^x, \quad x \in [0, \infty) \quad (5.34.13)$$

which is the PDF of the Gompertz distribution with shape parameter a .

In particular, if Y has the standard exponential distribution (rate parameter 1), then $X = \ln(Y + 1)$ has the standard Gompertz distribution (shape parameter 1). Since the exponential distribution is a scale family (the scale parameter is the reciprocal of the rate

parameter), we can construct an arbitrary basic Gompertz variable from a standard exponential variable. Specifically, if Y has the standard exponential variable and $a \in (0, \infty)$, then

$$X = \ln\left(\frac{1}{a}Y + 1\right) \quad (5.34.14)$$

has the Gompertz distribution with shape parameter a .

The extreme value distribution (Gumbel distribution) is also related to the Gompertz distribution.

If X has the standard extreme value distribution for minimums, then the conditional distribution of X given $X \geq 0$ is the standard Gompertz distribution.

Proof

By definition, X has PDF f given by $f(x) = e^x \exp(-e^x)$ for $x \in \mathbb{R}$. The conditional PDF of X given $X \geq 0$ is

$$g(x) = \frac{f(x)}{\mathbb{P}(X \geq 0)} = \frac{e^x \exp(-e^x)}{e^{-1}} = e^x \exp[-(e^x - 1)], \quad x \in [0, \infty) \quad (5.34.15)$$

which is the PDF of the standard Gompertz distribution.

The General Gompertz Distribution

The basic Gompertz distribution is generalized, like so many distributions on $[0, \infty)$, by adding a scale parameter. Recall that scale transformations often correspond to a change of units (minutes to hours, for example) and thus are fundamental.

If Z has the basic Gompertz distribution with shape parameter $a \in (0, \infty)$ and $b \in (0, \infty)$ then $X = bZ$ has the *Gompertz distribution* with shape parameter a and scale parameter b .

Distribution Functions

Suppose that X has the Gompertz distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

X has reliability function F^c given by

$$F^c(x) = \mathbb{P}(X > x) = \exp\left[-a\left(e^{x/b} - 1\right)\right], \quad x \in [0, \infty) \quad (5.34.16)$$

Proof

Recall that $F^c(x) = G^c(x/b)$ where G^c is the [reliability function of the corresponding basic distribution](#).

X has distribution function F given by

$$F(x) = \mathbb{P}(X \leq x) = 1 - \exp\left[-a\left(e^{x/b} - 1\right)\right], \quad x \in [0, \infty) \quad (5.34.17)$$

Proof

As before, $F = 1 - F^c$. Also, $F(x) = G(x/b)$ where G is the [CDF of the corresponding basic distribution](#).

X has quantile function F^{-1} given by

$$F^{-1}(p) = b \ln\left[1 - \frac{1}{a} \ln(1 - p)\right], \quad p \in [0, 1) \quad (5.34.18)$$

1. The first quartile is $q_1 = b \ln\left[1 + (\ln 4 - \ln 3)/a\right]$.
2. The median is $q_2 = b \ln(1 + \ln 2/a)$.
3. The third quartile is $q_3 = b \ln(1 + \ln 4/a)$.

Proof

Recall that $F^{-1}(p) = bG^{-1}(p)$ where G^{-1} is the [quantile function of the corresponding basic distribution](#).

Open the special distribution calculator and select the Gompertz distribution. Vary the shape and scale parameters and note the shape and location of the distribution function. For selected values of the parameters, compute a few values of the distribution function and the quantile function.

X has probability density function f given by

$$f(x) = \frac{a}{b} e^{x/b} \exp\left[-a \left(e^{x/b} - 1\right)\right], \quad x \in [0, \infty) \quad (5.34.19)$$

1. If $a < 1$ then f is increasing and then decreasing with mode $x = -b \ln(a)$.
2. If $a \geq 1$ then f is decreasing with mode $x = 0$.
3. If $a < (3 - \sqrt{5})/2 \approx 0.382$ then f is concave up and then down then up again, with inflection points at $x = b \ln[(3 \pm \sqrt{5})/2a]$.
4. If $(3 - \sqrt{5})/2 \leq a < (3 + \sqrt{5})/2 \approx 2.618$ then f has is concave down and then up, with inflection point at $x = b \ln[(3 + \sqrt{5})/2a]$.
5. If $a \geq (3 + \sqrt{5})/2$ then f is concave up.

Proof

Recall that $f(x) = \frac{1}{b} g\left(\frac{x}{b}\right)$ where g is the [PDF of the corresponding basic distribution](#).

Open the special distribution simulator and select the Gompertz distribution. Vary the shape and scale parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Once again, X has exponentially increasing failure rate.

X has failure rate function R given by

$$R(x) = \frac{a}{b} e^{x/b}, \quad x \in [0, \infty) \quad (5.34.20)$$

Proof

Recall that $R(x) = f(x)/F^c(x)$. Also, $R(x) = \frac{1}{b} r\left(\frac{x}{b}\right)$ where r is the [failure rate function of the corresponding basic distribution](#).

Moments

As with the basic distribution, the moment generating function and mean of the general Gompertz distribution can be expressed in terms of the [exponential integral](#). Suppose again that X has the Gompertz distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

X has moment generating function M given by

$$M(t) = \mathbb{E}(e^{tX}) = ae^a E_a(bt), \quad t \in \mathbb{R} \quad (5.34.21)$$

Proof

Recall that $M(t) = m(bt)$ where m is the [MGF of the corresponding basic distribution](#).

X has mean $\mathbb{E}(X) = be^a E_a(-1)$.

Proof

This follows from the [mean of the corresponding basic distribution](#), and the standard property $\mathbb{E}(X) = b\mathbb{E}(Z)$.

Open the special distribution simulator and select the Gompertz distribution. Vary the shape and scale parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

Since the Gompertz distribution is a scale family for each value of the shape parameter, it is trivially closed under scale transformations.

Suppose that X has the Gompertz distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. If $c \in (0, \infty)$ then $Y = cX$ has the Gompertz distribution with shape parameter a and scale parameter bc .

Proof

By [definition](#), we can take $X = bZ$ where Z has the standard Gompertz distribution with shape parameter a . But then $Y = cX = (bc)Z$.

As with the basic distribution, the Gompertz distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) computed above.

Suppose that $a, b \in (0, \infty)$.

1. If U has the standard uniform distribution then $X = b \ln(1 - \frac{1}{a} \ln U)$ has the Gompertz distribution with shape parameter a and scale parameter b .
2. If X has the Gompertz distribution with shape parameter a and scale parameter b , then $U = \exp[-a(e^{X/b} - 1)]$ has the standard uniform distribution.

Proof

This follows from the [corresponding result for the basic distribution](#) and the [definition](#) of the general Gompertz variable as $X = bZ$ where Z has the basic Gompertz distribution with shape parameter a .

Again, since the quantile function of the Gompertz distribution has a simple closed form, the distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the Gompertz distribution. Vary the shape and scale parameters and note the shape and location of the distribution and probability density functions. For selected values of the parameters, run the simulation 1000 times and note the agreement between the empirical density function and the probability density function.

The following result is a slight generalization of the [connection above](#) between the basic Gompertz distribution and the extreme value distribution.

If X has the extreme value distribution for minimums with scale parameter $b > 0$, then the conditional distribution of X given $X \geq 0$ is the Gompertz distribution with shape parameter 1 and scale parameter b .

Proof

We can take $X = bV$ where V has the standard extreme value distribution for minimums. Note that $X \geq 0$ if and only if $V \geq 0$. Hence the conditional distribution of X given $X \geq 0$ is the same as the conditional distribution of bV given $V \geq 0$. But by the [result above](#) the conditional distribution of V given $V \geq 0$ has the standard Gompertz distribution.

Finally, we give a slight generalization of the [connection above](#) between the Gompertz distribution and the exponential distribution.

Suppose that $a, b \in (0, \infty)$.

1. If X has the Gompertz distribution with shape parameter a and scale parameter b , then $Y = e^{X/b} - 1$ has the exponential distribution with rate parameter a .
2. If Y has the exponential distribution with rate parameter a , then $X = b \ln(Y + 1)$ has the Gompertz distribution with shape parameter a and scale parameter b .

Proof

These results follow from the [corresponding result for the basic distribution](#).

1. If X has the Gompertz distribution with shape parameter a and scale parameter b , then X/b has the basic Gompertz distribution with shape parameter a . Hence $Y = e^{X/b} - 1$ has the exponential distribution with rate parameter a .

2. If Y has the exponential distribution with rate parameter a then $\ln(Y + 1)$ has the basic Gompertz distribution with shape parameter a and hence $X = b \ln(Y + 1)$ has the Gompertz distribution with shape parameter a and scale parameter b . (

As a corollary, we can construct a general Gompertz variable from a standard exponential variable. Specifically, if Y has the standard exponential distribution and if $a, b \in (0, \infty)$ then

$$X = b \ln\left(\frac{1}{a}Y + 1\right) \quad (5.34.22)$$

has the Gompertz distribution with shape parameter a and scale parameter b .

This page titled [5.34: The Gompertz Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.35: The Log-Logistic Distribution

As the name suggests, the *log-logistic distribution* is the distribution of a variable whose logarithm has the logistic distribution. The log-logistic distribution is often used to model random lifetimes, and hence has applications in reliability.

The Basic Log-Logistic Distribution

Distribution Functions

The *basic log-logistic distribution* with *shape parameter* $k \in (0, \infty)$ is a continuous distribution on $[0, \infty)$ with distribution function G given by

$$G(z) = \frac{z^k}{1 + z^k}, \quad z \in [0, \infty) \quad (5.35.1)$$

In the special case that $k = 1$, the distribution is the *standard log-logistic distribution*.

Proof

Note that G is continuous on $[0, \infty)$ with $G(0) = 0$ and $G(z) \rightarrow 1$ as $z \rightarrow \infty$. Moreover,

$$g(z) = G'(z) = \frac{kz^{k-1}}{(1 + z^k)^2} > 0, \quad z \in (0, \infty) \quad (5.35.2)$$

so G is strictly increasing on $[0, \infty)$.

The probability density function g is given by

$$g(z) = \frac{kz^{k-1}}{(1 + z^k)^2}, \quad z \in (0, \infty) \quad (5.35.3)$$

1. If $0 < k < 1$, g is decreasing with $g(z) \rightarrow \infty$ as $z \downarrow 0$.
2. If $k = 1$, g is decreasing with mode $z = 0$.
3. If $k > 1$, g increases and then decreases with mode $z = \left(\frac{k-1}{k+1}\right)^{1/k}$.
4. If $k \leq 1$, g is concave upward.
5. If $1 < k \leq 2$, g is concave downward and then upward, with inflection point at

$$z = \left[\frac{2(k^2 - 1) + 2k\sqrt{3(k^2 - 1)}}{(k+1)(k+2)} \right]^{1/k} \quad (5.35.4)$$

6. If $k > 2$, g is concave upward then downward then upward again, with inflection points at

$$z = \left[\frac{2(k^2 - 1) \pm 2k\sqrt{3(k^2 - 1)}}{(k+1)(k+2)} \right]^{1/k} \quad (5.35.5)$$

Proof

The PDF $g = G'$ was computed in the proof of the [CDF result](#). The rest follows from

$$g'(z) = \frac{kz^{k-2}[(k-1) - (k+1)z^k]}{(1 + z^k)^3}, \quad z \in (0, \infty) \quad (5.35.6)$$

$$g''(z) = \frac{kz^{k-3}[(k-1)(k-2) - 4(k^2-1)z^k + (k+1)(k+2)z^{2k}]}{(1 + z^k)^4}, \quad z \in (0, \infty) \quad (5.35.7)$$

So g has a rich variety of shapes, and is unimodal if $k > 1$. When $k \geq 1$, g is defined at 0 as well.

Open the special distribution simulator and select the log-logistic distribution. Vary the shape parameter and note the shape of the probability density function. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function G^{-1} is given by

$$G^{-1}(p) = \left(\frac{p}{1-p} \right)^{1/k}, \quad p \in [0, 1) \quad (5.35.8)$$

1. The first quartile is $q_1 = (1/3)^{1/k}$.
2. The median is $q_2 = 1$.
3. The third quartile is $q_3 = 3^{1/k}$.

Proof

The formula for G^{-1} follows from the [distribution function](#) by solving $p = G(z)$ for z in terms of p .

Recall that $p/(1-p)$ is the *odds ratio* associated with probability $p \in (0, 1)$. Thus, the quantile function of the basic log-logistic distribution with shape parameter k is the k th root of the odds ratio function. In particular, the quantile function of the standard log-logistic distribution is the odds ratio function itself. Also of interest is that the median is 1 for every value of the shape parameter.

Open the special distribution calculator and select the log-logistic distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the shape parameter, compute a few values of the distribution function and the quantile function.

The reliability function G^c is given by

$$G^c(z) = \frac{1}{1+z^k}, \quad z \in [0, \infty) \quad (5.35.9)$$

Proof

This follows trivially from the [distribution function](#) since $G^c = 1 - G$.

The basic log-logistic distribution has either decreasing failure rate, or mixed decreasing-increasing failure rate, depending on the shape parameter.

The failure rate function r is given by

$$r(z) = \frac{kz^{k-1}}{1+z^k}, \quad z \in (0, \infty) \quad (5.35.10)$$

1. If $0 < k \leq 1$, r is decreasing.
2. If $k > 1$, r decreases and then increases with minimum at $z = (k-1)^{1/k}$.

Proof

Recall that the is $r(z) = g(z)/G^c(z)$ for $z \in (0, \infty)$ so the formula follows from the [PDF](#) and the [reliability function](#) above. Parts (a) and (b) follow from

$$r'(z) = \frac{kz^{k-1}[(k-1) - z^k]}{(1+z^k)^2}, \quad z \in (0, \infty) \quad (5.35.11)$$

If $k \geq 1$, r is defined at 0 also.

Moments

Suppose that Z has the basic log-logistic distribution with shape parameter $k \in (0, \infty)$. The moments (about 0) of the Z have an interesting expression in terms of the beta function B and in terms of the sine function. The simplest representation is in terms of a new special function constructed from the sine function.

The (normalized) *cardinal sine function* sinc is defined by

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}, \quad x \in \mathbb{R} \quad (5.35.12)$$

where it is understood that $\text{sinc}(0) = 1$ (the limiting value).




Figure 5.35.1: The graph of the sinc function on the interval $[-10, 10]$

If $n \geq k$ then $\mathbb{E}(Z^n) = \infty$. If $0 \leq n < k$ then

$$\mathbb{E}(Z^n) = B\left(1 - \frac{n}{k}, 1 + \frac{n}{k}\right) = \frac{1}{\text{sinc}(n/k)} \quad (5.35.13)$$

Proof

Using the PDF,

$$\mathbb{E}(Z^n) = \int_0^\infty z^n \frac{kz^{k-1}}{(1+z^k)^2} dz \quad (5.35.14)$$

The substitution $u = 1/(1+z^k)$, $du = -kz^{k-1}/(1+z^k)^2$ gives

$$\mathbb{E}(Z^n) = \int_0^1 (1/u - 1)^{n/k} du = \int_0^1 u^{-n/k} (1-u)^{n/k} du \quad (5.35.15)$$

The result now follows from the definition of the beta function.

In particular, we can give the mean and variance.

If $k > 1$ then

$$\mathbb{E}(Z) = \frac{1}{\text{sinc}(1/k)} \quad (5.35.16)$$

If $k > 2$ then

$$\text{var}(Z) = \frac{1}{\text{sinc}(2/k)} - \frac{1}{\text{sinc}^2(1/k)} \quad (5.35.17)$$

Open the special distribution simulator and select the log-logistic distribution. Vary the shape parameter and note the size and location of the mean \pm standard deviation bar. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

The basic log-logistic distribution is preserved under power transformations.

If Z has the basic log-logistic distribution with shape parameter $k \in (0, \infty)$ and if $n \in (0, \infty)$, then $W = Z^n$ has the basic log-logistic distribution with shape parameter k/n .

Proof

For $w \in [0, \infty)$,

$$\mathbb{P}(W \leq w) = \mathbb{P}(Z \leq w^{1/n}) = G(w^{1/n}) = \frac{w^{k/n}}{1 + w^{k/n}} \quad (5.35.18)$$

As a function of w , this is the CDF of the basic log-logistic distribution with shape parameter k/n .

In particular, it follows that if V has the standard log-logistic distribution and $k \in (0, \infty)$, then $Z = V^{1/k}$ has the basic log-logistic distribution with shape parameter k .

The log-logistic distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) given above.

Suppose that $k \in (0, \infty)$.

1. If U has the standard uniform distribution then $Z = G^{-1}(U) = [U/(1-U)]^{1/k}$ has the basic log-logistic distribution with shape parameter k .
2. If Z has the basic log-logistic distribution with shape parameter k then $U = G(Z) = Z^k/(1+Z^k)$ has the standard uniform distribution.

Since the quantile function of the basic log-logistic distribution has a simple closed form, the distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the log-logistic distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts..

Of course, as mentioned in the introduction, the log-logistic distribution is related to the logistic distribution.

Suppose that $k, b \in (0, \infty)$.

1. If Z has the basic log-logistic distribution with shape parameter k then $Y = \ln Z$ has the logistic distribution with location parameter 0 and scale parameter $1/k$.
2. If Y has the logistic distribution with location parameter 0 and scale parameter b then $Z = e^Y$ has the basic log-logistic distribution with shape parameter $1/b$.

Proof

1. Suppose first that Z has the standard log-logistic distribution. Then

$$\mathbb{P}(Y \leq y) = \mathbb{P}(Z \leq e^y) = \frac{e^y}{1 + e^y}, \quad y \in \mathbb{R} \quad (5.35.19)$$

and as a function of y , this is the CDF of the standard logistic distribution. Suppose now that Z has the basic log-logistic distribution with shape parameter k . From the [power result](#), we can take $Z = W^{1/k}$ where W has the standard log-logistic distribution. Then $Y = \ln Z = \frac{1}{k} \ln W$. But $\ln(W)$ has the standard logistic distribution, and hence $\frac{1}{k} \ln W$ has the logistic distribution with location parameter 0 and scale parameter $1/k$.

2. Suppose first that Y has the standard logistic distribution. Then

$$\mathbb{P}(Z \leq z) = \mathbb{P}[Y \leq \ln(z)] = \frac{e^{\ln z}}{1 + e^{\ln z}} = \frac{z}{1 + z}, \quad z \in (0, \infty) \quad (5.35.20)$$

and as a function of z , this is the CDF of the standard log-logistic distribution. Suppose now that Y has the logistic distribution with location parameter 0 and scale parameter b . We can take $Y = bV$ where V has the standard logistic distribution. Hence $Z = e^Y = e^{bV} = (e^V)^b$. But e^V has the standard log-logistic distribution, and again by the [power result](#) $(e^V)^b$ has the log-logistic distribution with shape parameter $1/b$.

As a special case, (and as noted in the proof), if Z has the standard log-logistic distribution, then $Y = \ln Z$ has the standard logistic distribution, and if Y has the standard logistic distribution, then $Z = e^Y$ has the standard log-logistic distribution.

The standard log-logistic distribution is the same as the standard beta prime distribution.

Proof

The PDF of the standard log-logistic distribution is $g(z) = 1/(1+z)^2$ for $z \in [0, \infty)$, which is the same as the PDF of the standard beta prime distribution.

Of course, limiting distributions with respect to parameters are always interesting.

The basic log-logistic distribution with shape parameter $k \in (0, \infty)$ converges to point mass at 1 as $k \rightarrow \infty$.

Proof from the definition

Note that the [distribution function](#) satisfies $G(z) \rightarrow 0$ as $k \rightarrow \infty$ for $0 \leq z < 1$, $G(1) = \frac{1}{2}$ for all $k > 1$, and $G(z) \rightarrow 1$ as $k \rightarrow \infty$ for $z > 1$. Except for the point of discontinuity $z = 1$, the limiting values are the distribution function of point mass at 1.

Random variable proof

Suppose that V has the standard log-logistic distribution, and for $k \in (0, \infty)$, let $Z_k = V^{1/k}$, so that Z_k has the basic log-logistic distribution with shape parameter k . The event $\{V > 0\}$ has probability 1, and on this event, $Z_k \rightarrow 1$ as $k \rightarrow \infty$. But convergence with probability 1 implies convergence in distribution.

The General Log-Logistic Distribution

The basic log-logistic distribution is generalized, like so many distributions on $[0, \infty)$, by adding a scale parameter. Recall that a scale transformation often corresponds to a change of units (gallons into liters, for example), and so such transformations are of basic importance.

If Z has the basic log-logistic distribution with shape parameter $k \in (0, \infty)$ and if $b \in (0, \infty)$ then $X = bZ$ has the *log-logistic distribution* with *shape parameter* k and *scale parameter* b .

Distribution Functions

Suppose that X has the log-logistic distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

X has distribution function F given by

$$F(x) = \frac{x^k}{b^k + x^k}, \quad x \in [0, \infty) \quad (5.35.21)$$

Proof

Recall that $F(x) = G(x/b)$ where G is the [distribution function](#) of the basic log-logistic distribution with shape parameter k .

X has probability density function f given by

$$f(x) = \frac{b^k k x^{k-1}}{(b^k + x^k)^2}, \quad x \in (0, \infty) \quad (5.35.22)$$

When $k \geq 1$, f is defined at 0 also. f satisfies the following properties:

1. If $0 < k < 1$, f is decreasing with $f(x) \rightarrow \infty$ as $x \downarrow 0$.
2. If $k = 1$, f is decreasing with mode $x = 0$.
3. If $k > 1$, f increases and then decreases with mode $x = b \left(\frac{k-1}{k+1} \right)^{1/k}$.
4. If $k \leq 1$, f is concave upward.
5. If $1 < k \leq 2$, f is concave downward and then upward, with inflection point at

$$x = b \left[\frac{2(k^2 - 1) + 2k\sqrt{3(k^2 - 1)}}{(k+1)(k+2)} \right]^{1/k} \quad (5.35.23)$$

6. If $k > 2$, f is concave upward then downward then upward again, with inflection points at

$$x = b \left[\frac{2(k^2 - 1) \pm 2k\sqrt{3(k^2 - 1)}}{(k+1)(k+2)} \right]^{1/k} \quad (5.35.24)$$

Proof

Recall that $f(x) = \frac{1}{b}g\left(\frac{x}{b}\right)$ where g is the [probability density function](#) of the basic log-logistic distribution with shape parameter k . Also of course, $f = F'$.

Open the special distribution simulator and select the log-logistic distribution. Vary the shape and scale parameters and note the shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has quantile function F^{-1} given by

$$F^{-1}(p) = b \left(\frac{p}{1-p} \right)^{1/k}, \quad p \in [0, 1) \quad (5.35.25)$$

1. The first quartile is $q_1 = b(1/3)^{1/k}$.
2. The median is $q_2 = b$.
3. The third quartile is $q_3 = b3^{1/k}$.

Proof

Recall that $F^{-1}(p) = bG^{-1}(p)$ for $p \in [0, 1)$ where G^{-1} is the [quantile function](#) of the basic log-logistic distribution with shape parameter k .

Open the special distribution calculator and select the log-logistic distribution. Vary the shape and scale parameters and note the shape of the distribution and probability density functions. For selected values of the parameters, compute a few values of the distribution function and the quantile function.

X has reliability function F^c given by

$$F^c(x) = \frac{b^k}{b^k + x^k}, \quad x \in [0, \infty) \quad (5.35.26)$$

Proof

This follows trivially from the [distribution function](#), since $F^c = 1 - F$.

The log-logistic distribution has either decreasing failure rate, or mixed decreasing-increasing failure rate, depending on the shape parameter.

X has failure rate function R given by

$$R(x) = \frac{kx^{k-1}}{b^k + x^k}, \quad x \in (0, \infty) \quad (5.35.27)$$

1. If $0 < k \leq 1$, R is decreasing.
2. If $k > 1$, R decreases and then increases with minimum at $x = b(k-1)^{1/k}$.

Proof

Recall that $R(x) = \frac{1}{b}r\left(\frac{x}{b}\right)$ where r is the [failure rate function](#) of the basic log-logistic distribution with shape parameter k . Also, $R = f/F^c$ where f is the [PDF](#) and F^c is the [reliability function](#).

Moments

Suppose again that X has the log-logistic distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. The moments of X can be computed easily from the representation $X = bZ$ where Z has the basic log-logistic distribution with shape parameter k . Again, the expressions are simplest in terms of the beta function B and in terms of the normalized cardinal sine function [sinc](#).

If $n \geq k$ then $\mathbb{E}(X^n) = \infty$. If $0 \leq n < k$ then

$$\mathbb{E}(X^n) = b^n B\left(1 - \frac{n}{k}, 1 + \frac{n}{k}\right) = \frac{b^n}{\text{sinc}(n/k)} \quad (5.35.28)$$

If $k > 1$ then

$$\mathbb{E}(X) = \frac{b}{\text{sinc}(1/k)} \quad (5.35.29)$$

If $k > 2$ then

$$\text{var}(X) = b^2 \left[\frac{1}{\text{sinc}(2/k)} - \frac{1}{\text{sinc}^2(1/k)} \right] \quad (5.35.30)$$

Open the special distribution simulator and select the log-logistic distribution. Vary the shape and scale parameters and note the size and location of the mean/standard deviation bar. For selected values of the parameters, run the simulation 1000 times compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

Since the log-logistic distribution is a scale family for each value of the shape parameter, it is trivially closed under scale transformations.

If X has the log-logistic distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$, and if $c \in (0, \infty)$, then $Y = cX$ has the log-logistic distribution with shape parameter k and scale parameter bc .

Proof

By [definition](#) we can take $X = bZ$ where Z has the basic log-logistic distribution with shape parameter k . But then $Y = cX = (bc)Z$.

The log-logistic distribution is preserved under power transformations.

If X has the log-logistic distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$, and if $n \in (0, \infty)$, then $Y = X^n$ has the log-logistic distribution with shape parameter k/n and scale parameter b^n .

Proof

Again we can take $X = bZ$ where Z has the basic log-logistic distribution with shape parameter k . Then $X^n = b^n Z^n$. But by the [power result](#) for the standard distribution, Z^n has the basic log-logistic distribution with shape parameter k/n and hence X has the log-logistic distribution with shape parameter k/n and scale parameter b^n .

In particular, if V has the standard log-logistic distribution, then $X = bV^{1/k}$ has the log-logistic distribution with shape parameter k and scale parameter b .

As before, the log-logistic distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) computed above.

Suppose that $k, b \in (0, \infty)$.

1. If U has the standard uniform distribution then $X = F^{-1}(U) = b[U/(1-U)]^{1/k}$ has the log-logistic distribution with shape parameter k and scale parameter b .

2. If X has the log-logistic distribution with shape parameter k and scale parameter b , then $U = F(X) = X^k / (b^k + X^k)$ has the standard uniform distribution.

Again, since the quantile function of the log-logistic distribution has a simple closed form, the distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the log-logistic distribution. Vary the shape and scale parameters and note the shape and location of the distribution and probability density functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function, mean and standard deviation to their distributional counterparts.

Again, the logarithm of a log-logistic variable has the logistic distribution.

Suppose that $k, b, c \in (0, \infty)$ and $a \in \mathbb{R}$.

1. If X has the log-logistic distribution with shape parameter k and scale parameter b then $Y = \ln X$ has the logistic distribution with location parameter $\ln b$ and scale parameter $1/k$.
2. If Y has the logistic distribution with location parameter a and scale parameter c then $X = e^Y$ has the log-logistic distribution with shape parameter $1/c$ and scale parameter e^a .

Proof

1. As noted [above](#), we can take $X = bV^{1/k}$ where V has the standard log-logistic distribution. Then $Y = \ln X = \ln b + \frac{1}{k} \ln V$. But by the [corresponding result](#) for the basic distribution, $\ln V$ has the standard logistic distribution, so Y has the logistic distribution with location parameter $\ln b$ and scale parameter $1/k$.
2. We can take $Y = a + cU$ where U has the standard logistic distribution. Hence $X = e^Y = e^a e^{cU} = e^a (e^U)^c$. But by the [result corresponding result](#) for the standard distribution, e^U has the standard log-logistic distribution so X has the log-logistic distribution with shape parameter $1/c$ and scale parameter e^a .

Once again, the limiting distribution is also of interest.

For fixed $b \in (0, \infty)$, the log-logistic distribution with shape parameter $k \in (0, \infty)$ and scale parameter b converges to point mass at b as $k \rightarrow \infty$.

Proof

If X has the log-logistic distribution with shape parameter k and scale parameter b , then as usual, we can write $X = bZ$ where Z has the basic log-logistic distribution with shape parameter k . From the [limit result](#) for the basic distribution, we know that the distribution of Z converges to point mass at 1 as $k \rightarrow \infty$, so it follows by the continuity theorem that the distribution of X converges to point mass at b as $k \rightarrow \infty$.

This page titled [5.35: The Log-Logistic Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.36: The Pareto Distribution

The Pareto distribution is a skewed, heavy-tailed distribution that is sometimes used to model the distribution of incomes and other financial variables.

The Basic Pareto Distribution

Distribution Functions

The basic *Pareto distribution* with *shape parameter* $a \in (0, \infty)$ is a continuous distribution on $[1, \infty)$ with distribution function G given by

$$G(z) = 1 - \frac{1}{z^a}, \quad z \in [1, \infty) \quad (5.36.1)$$

The special case $a = 1$ gives the *standard Pareto distribution*.

Proof

Clearly G is increasing and continuous on $[1, \infty)$, with $G(1) = 0$ and $G(z) \rightarrow 1$ as $z \rightarrow \infty$.

The Pareto distribution is named for the economist Vilfredo Pareto.

The probability density function g is given by

$$g(z) = \frac{a}{z^{a+1}}, \quad z \in [1, \infty) \quad (5.36.2)$$

1. g is decreasing with mode $z = 1$
2. g is concave upward.

Proof

Recall that $g = G'$. Parts (a) and (b) follow from standard calculus.

The reason that the Pareto distribution is *heavy-tailed* is that the g decreases at a *power rate* rather than an *exponential rate*.

Open the special distribution simulator and select the Pareto distribution. Vary the shape parameter and note the shape of the probability density function. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function G^{-1} is given by

$$G^{-1}(p) = \frac{1}{(1-p)^{1/a}}, \quad p \in [0, 1) \quad (5.36.3)$$

1. The first quartile is $q_1 = \left(\frac{4}{3}\right)^{1/a}$.
2. The median is $q_2 = 2^{1/a}$.
3. The third quartile is $q_3 = 4^{1/a}$.

Proof

The formula for $G^{-1}(p)$ comes from solving $G(z) = p$ for z in terms of p .

Open the special distribution calculator and select the Pareto distribution. Vary the shape parameter and note the shape of the probability density and distribution functions. For selected values of the parameters, compute a few values of the distribution and quantile functions.

Moments

Suppose that random variable Z has the basic Pareto distribution with shape parameter $a \in (0, \infty)$. Because the distribution is heavy-tailed, the mean, variance, and other moments of Z are finite only if the shape parameter a is sufficiently large.

The moments of Z (about 0) are

1. $\mathbb{E}(Z^n) = \frac{a}{a-n}$ if $0 < n < a$
2. $\mathbb{E}(Z^n) = \infty$ if $n \geq a$

Proof

Note that

$$E(Z^n) = \int_1^\infty z^n \frac{a}{z^{a+1}} dz = \int_1^\infty a z^{-(a+1-n)} dz \quad (5.36.4)$$

The integral diverges to ∞ if $a+1-n \leq 1$ and evaluates to $\frac{a}{a-n}$ if $a+1-n > 1$.

It follows that the moment generating function of Z cannot be finite on any interval about 0.

In particular, the mean and variance of Z are

1. $\mathbb{E}(Z) = \frac{a}{a-1}$ if $a > 1$
2. $\text{var}(Z) = \frac{a}{(a-1)^2(a-2)}$ if $a > 2$

Proof

This results follow from the [general moment formula](#) above and the computational formula $\text{var}(Z) = \mathbb{E}(Z^2) - [E(Z)]^2$.

In the special distribution simulator, select the Pareto distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For each of the following parameter values, run the simulation 1000 times and note the behavior of the empirical moments:

1. $a = 1$
2. $a = 2$
3. $a = 3$

The skewness and kurtosis of Z are as follows:

1. If $a > 3$,

$$\text{skew}(Z) = \frac{2(1+a)}{a-3} \sqrt{1 - \frac{2}{a}} \quad (5.36.5)$$

2. If $a > 4$,

$$\text{kurt}(Z) = \frac{3(a-2)(3a^2+a+2)}{a(a-3)(a-4)} \quad (5.36.6)$$

Proof

These results follow from the standard computational formulas for skewness and kurtosis, and the first 4 [moments of \$Z\$](#) given above.

So the distribution is positively skewed and $\text{skew}(Z) \rightarrow 2$ as $a \rightarrow \infty$ while $\text{skew}(Z) \rightarrow \infty$ as $a \downarrow 3$. Similarly, $\text{kurt}(Z) \rightarrow 9$ as $a \rightarrow \infty$ and $\text{kurt}(Z) \rightarrow \infty$ as $a \downarrow 4$. Recall that the *excess kurtosis* of Z is

$$\text{kurt}(Z) - 3 = \frac{3(a-2)(3a^2+a+2)}{a(a-3)(a-4)} - 3 = \frac{6(a^3+a^2-6a-1)}{a(a-3)(a-4)} \quad (5.36.7)$$

Related Distributions

The basic Pareto distribution is invariant under positive powers of the underlying variable.

Suppose that Z has the basic Pareto distribution with shape parameter $a \in (0, \infty)$ and that $n \in (0, \infty)$. Then $W = Z^n$ has the basic Pareto distribution with shape parameter a/n .

Proof

We use the [CDF of \$Z\$](#) given above.

$$\mathbb{P}(W \leq w) = \mathbb{P}\left(Z \leq w^{1/n}\right) = 1 - \frac{1}{w^{a/n}}, \quad w \in [1, \infty) \quad (5.36.8)$$

As a function of w , this is the Pareto CDF with shape parameter a/n .

In particular, if Z has the standard Pareto distribution and $a \in (0, \infty)$, then $Z^{1/a}$ has the basic Pareto distribution with shape parameter a . Thus, all basic Pareto variables can be constructed from the standard one.

The basic Pareto distribution has a reciprocal relationship with the beta distribution.

Suppose that $a \in (0, \infty)$.

1. If Z has the basic Pareto distribution with shape parameter a then $V = 1/Z$ has the beta distribution with left parameter a and right parameter 1.
2. If V has the beta distribution with left parameter a and right parameter 1, then $Z = 1/V$ has the basic Pareto distribution with shape parameter a .

Proof

We will use the standard change of variables theorem. The transformations are $v = 1/z$ and $z = 1/v$ for $z \in [1, \infty)$ and $v \in (0, 1]$. These are inverses of each other. Let g and h denote PDFs of Z and V respectively.

1. We start with $g(z) = a/z^{a+1}$ for $z \in [1, \infty)$, the [PDF of \$Z\$](#) given above. Then

$$h(v) = g(z) \left| \frac{dz}{dv} \right| = \frac{a}{(1/v)^{a+1}} \frac{1}{v^2} = av^{a-1}, \quad v \in (0, 1] \quad (5.36.9)$$

which is the PDF of the beta distribution with left parameter a and right parameter 1.

2. We start with $h(v) = av^{a-1}$ for $v \in (0, 1]$. Then

$$g(z) = h(v) \left| \frac{dv}{dz} \right| = a \left(\frac{1}{z} \right)^{a-1} \frac{1}{z^2} = \frac{a}{z^{a+1}}, \quad z \in [1, \infty) \quad (5.36.10)$$

which is the PDF of the basic Pareto distribution with shape parameter a .

The basic Pareto distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) computed above.

Suppose that $a \in (0, \infty)$.

1. If U has the standard uniform distribution then $Z = 1/U^{1/a}$ has the basic Pareto distribution with shape parameter a .
2. If Z has the basic Pareto distribution with shape parameter a then $U = 1/Z^a$ has the standard uniform distribution.

Proof

1. If U has the standard uniform distribution, then so does $1 - U$. Hence $Z = G^{-1}(1 - U) = 1/U^{1/a}$ has the basic Pareto distribution with shape parameter a .
2. If Z has the basic Pareto distribution with shape parameter a , then $G(Z)$ has the standard uniform distribution. But then $U = 1 - G(Z) = 1/Z^a$ also has the standard uniform distribution.

Since the quantile function has a simple closed form, the basic Pareto distribution can be simulated using the random quantile method.

Open the random quantile experiment and selected the Pareto distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, run the experiment 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The basic Pareto distribution also has simple connections to the exponential distribution.

Suppose that $a \in (0, \infty)$.

1. If Z has the basic Pareto distribution with shape parameter a , then $T = \ln Z$ has the exponential distribution with rate parameter a .
2. If T has the exponential distribution with rate parameter a , then $Z = e^T$ has the basic Pareto distribution with shape parameter a .

Proof

We use the [Pareto CDF](#) given above and the CDF of the exponential distribution.

1. If $t \in [0, \infty)$ then

$$\mathbb{P}(T \leq t) = \mathbb{P}(Z \leq e^t) = 1 - \frac{1}{(e^t)^a} = 1 - e^{-at} \quad (5.36.11)$$

which is the CDF of the exponential distribution with rate parameter a .

2. If $z \in [1, \infty)$ then

$$\mathbb{P}(Z \leq z) = \mathbb{P}(T \leq \ln z) = 1 - \exp(-a \ln z) = 1 - \frac{1}{z^a} \quad (5.36.12)$$

which is the CDF of the basic Pareto distribution with shape parameter a .

The General Pareto Distribution

As with many other distributions that govern positive variables, the Pareto distribution is often generalized by adding a scale parameter. Recall that a scale transformation often corresponds to a change of units (dollars into Euros, for example) and thus such transformations are of basic importance.

Suppose that Z has the basic Pareto distribution with shape parameter $a \in (0, \infty)$ and that $b \in (0, \infty)$. Random variable $X = bZ$ has the *Pareto distribution* with shape parameter a and scale parameter b .

Note that X has a continuous distribution on the interval $[b, \infty)$.

Distribution Functions

Suppose again that X has the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

X has distribution function F given by

$$F(x) = 1 - \left(\frac{b}{x}\right)^a, \quad x \in [b, \infty) \quad (5.36.13)$$

Proof

Recall that $F(x) = G\left(\frac{x}{b}\right)$ for $x \in [b, \infty)$ where G is the [CDF of the basic distribution](#) with shape parameter a .

X has probability density function f given by

$$f(x) = \frac{ab^a}{x^{a+1}}, \quad x \in [b, \infty) \quad (5.36.14)$$

Proof

Recall that $f(x) = \frac{1}{b}g\left(\frac{x}{b}\right)$ for $x \in [b, \infty)$ where g is the [PDF of the basic distribution](#) with shape parameter a .

Open the special distribution simulator and select the Pareto distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has quantile function F^{-1} given by

$$F^{-1}(p) = \frac{b}{(1-p)^{1/a}}, \quad p \in [0, 1) \quad (5.36.15)$$

1. The first quartile is $q_1 = b\left(\frac{4}{3}\right)^{1/a}$.
2. The median is $q_2 = b2^{1/a}$.
3. The third quartile is $q_3 = b4^{1/a}$.

Proof

Recall that $F^{-1}(p) = bG^{-1}(p)$ for $p \in [0, 1)$ where G^{-1} is the [quantile function](#) of the basic distribution with shape parameter a .

Open the special distribution calculator and select the Pareto distribution. Vary the parameters and note the shape and location of the probability density and distribution functions. For selected values of the parameters, compute a few values of the distribution and quantile functions.

Moments

Suppose again that X has the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$

The moments of X are given by

1. $\mathbb{E}(X^n) = b^n \frac{a}{a-n}$ if $0 < n < a$
2. $\mathbb{E}(X^n) = \infty$ if $n \geq a$

Proof

By [definition](#) we can take $X = bZ$ where Z has the basic Pareto distribution with shape parameter a . By the linearity of expected value, $\mathbb{E}(X^n) = b^n \mathbb{E}(Z^n)$, so the result follows from the [moments of \$Z\$](#) given above.

The mean and variance of X are

1. $\mathbb{E}(X) = b \frac{a}{a-1}$ if $a > 1$
2. $\text{var}(X) = b^2 \frac{a}{(a-1)^2(a-2)}$ if $a > 2$

Open the special distribution simulator and select the Pareto distribution. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are as follows:

1. If $a > 3$,

$$\text{skew}(X) = \frac{2(1+a)}{a-3} \sqrt{1 - \frac{2}{a}} \quad (5.36.16)$$

2. If $a > 4$,

$$\text{kurt}(X) = \frac{3(a-2)(3a^2+a+2)}{a(a-3)(a-4)} \quad (5.36.17)$$

Proof

Recall that skewness and kurtosis are defined in terms of the standard score, and hence are invariant under scale transformations. Thus the skewness and kurtosis of X are the same as the [skewness and kurtosis](#) of $Z = X/b$ given above.

Related Distributions

Since the Pareto distribution is a scale family for fixed values of the shape parameter, it is trivially closed under scale transformations.

Suppose that X has the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. If $c \in (0, \infty)$ then $Y = cX$ has the Pareto distribution with shape parameter a and scale parameter bc .

Proof

By [definition](#) we can take $X = bZ$ where Z has the basic Pareto distribution with shape parameter a . But then $Y = cX = (bc)Z$.

The Pareto distribution is closed under positive powers of the underlying variable.

Suppose that X has the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. If $n \in (0, \infty)$ then $Y = X^n$ has the Pareto distribution with shape parameter a/n and scale parameter b^n .

Proof

Again we can write $X = bZ$ where Z has the basic Pareto distribution with shape parameter a . Then from the [power result](#) above Z^n has the basic Pareto distribution with shape parameter a/n and hence $Y = X^n = b^n Z^n$ has the Pareto distribution with shape parameter a/n and scale parameter b^n .

All Pareto variables can be constructed from the standard one. If Z has the standard Pareto distribution and $a, b \in (0, \infty)$ then $X = bZ^{1/a}$ has the Pareto distribution with shape parameter a and scale parameter b .

As before, the Pareto distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) given above.

Suppose that $a, b \in (0, \infty)$.

1. If U has the standard uniform distribution then $X = b/U^{1/a}$ has the Pareto distribution with shape parameter a and scale parameter b .
2. If X has the Pareto distribution with shape parameter a and scale parameter b , then $U = (b/X)^a$ has the standard uniform distribution.

Proof

1. If U has the standard uniform distribution, then so does $1 - U$. Hence $X = F^{-1}(1 - U) = b/U^{1/a}$ has the Pareto distribution with shape parameter a and scale parameter b .
2. If X has the Pareto distribution with shape parameter a and scale parameter b , then $F(X)$ has the standard uniform distribution. But then $U = 1 - F(X) = (b/X)^a$ also has the standard uniform distribution.

Again, since the quantile function has a simple closed form, the basic Pareto distribution can be simulated using the random quantile method.

Open the random quantile experiment and selected the Pareto distribution. Vary the parameters and note the shape of the distribution and probability density functions. For selected values of the parameters, run the experiment 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

The Pareto distribution is closed with respect to conditioning on a right-tail event.

Suppose that X has the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. For $c \in [b, \infty)$, the conditional distribution of X given $X \geq c$ is Pareto with shape parameter a and scale parameter c .

Proof

Not surprisingly, it's best to use right-tail distribution functions. Recall that this is the function $F^c = 1 - F$ where F is the ordinary CDF given above. If $x \geq c$, then

$$\mathbb{P}(X > x \mid X > c) = \frac{\mathbb{P}(X > x)}{\mathbb{P}(X > c)} = \frac{(b/x)^a}{(b/c)^a} = (c/x)^a \quad (5.36.18)$$

Finally, the Pareto distribution is a general exponential distribution with respect to the shape parameter, for a fixed value of the scale parameter.

Suppose that X has the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. For fixed b , the distribution of X is a general exponential distribution with natural parameter $-(a+1)$ and natural statistic $\ln X$.

Proof

This follows from the definition of the general exponential family, since the pdf above can be written in the form

$$f(x) = ab^a \exp[-(a+1) \ln x], \quad x \in [b, \infty) \quad (5.36.19)$$

Computational Exercises

Suppose that the income of a certain population has the Pareto distribution with shape parameter 3 and scale parameter 1000. Find each of the following:

1. The proportion of the population with incomes between 2000 and 4000.
2. The median income.
3. The first and third quartiles and the interquartile range.
4. The mean income.
5. The standard deviation of income.
6. The 90th percentile.

Answer

1. $\mathbb{P}(2000 < X < 4000) = 0.1637$ so the proportion is 16.37%
2. $Q_2 = 1259.92$
3. $Q_1 = 1100.64$, $Q_3 = 1587.40$, $Q_3 - Q_1 = 486.76$
4. $\mathbb{E}(X) = 1500$
5. $\text{sd}(X) = 866.03$
6. $F^{-1}(0.9) = 2154.43$

This page titled [5.36: The Pareto Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.37: The Wald Distribution

The *Wald distribution*, named for Abraham Wald, is important in the study of Brownian motion. Specifically, the distribution governs the first time that a Brownian motion with positive drift hits a fixed, positive value. In Brownian motion, the distribution of the *random position* at a *fixed time* has a normal (Gaussian) distribution, and thus the Wald distribution, which governs the *random time* at a *fixed position*, is sometimes called the *inverse Gaussian distribution*.

The Basic Wald Distribution

Distribution Functions

As usual, let Φ denote the standard normal distribution function.

The *basic Wald distribution* with *shape parameter* $\lambda \in (0, \infty)$ is a continuous distribution on $(0, \infty)$ with distribution function G given by

$$G(u) = \Phi \left[\sqrt{\frac{\lambda}{u}}(u-1) \right] + e^{2\lambda} \Phi \left[-\sqrt{\frac{\lambda}{u}}(u+1) \right], \quad u \in (0, \infty) \quad (5.37.1)$$

The special case $\lambda = 1$ gives the *standard Wald distribution*.

Proof

Note that as $u \rightarrow \infty$, $\sqrt{\frac{\lambda}{u}}(u-1) \rightarrow \infty$ and $-\sqrt{\frac{\lambda}{u}}(u+1) \rightarrow -\infty$, and hence $G(u) \rightarrow 1$. As $u \downarrow 0$, $\sqrt{\frac{\lambda}{u}}(u-1) \rightarrow -\infty$ and $-\sqrt{\frac{\lambda}{u}}(u+1) \rightarrow -\infty$, and hence $G(u) \rightarrow 0$. Of course, G is clearly continuous on $(0, \infty)$, so it remains to show that G is increasing on this interval. Differentiating gives

$$G'(u) = \phi \left[\sqrt{\frac{\lambda}{u}}(u-1) \right] \left[\frac{\sqrt{\lambda}}{2} (u^{-1/2} + u^{-3/2}) \right] + e^{2\lambda} \phi \left[-\sqrt{\frac{\lambda}{u}}(u+1) \right] \left[-\frac{\sqrt{\lambda}}{2} (u^{-1/2} - u^{-3/2}) \right] \quad (5.37.2)$$

where $\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$ is the standard normal PDF. Simple algebra shows that

$$\phi \left[\sqrt{\frac{\lambda}{u}}(u-1) \right] = e^{2\lambda} \phi \left[-\sqrt{\frac{\lambda}{u}}(u+1) \right] = \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{\lambda}{2u} (u-1)^2 \right] \quad (5.37.3)$$

so simplifying further gives

$$G'(u) = \sqrt{\frac{\lambda}{2\pi}} u^{-3/2} \exp \left[-\frac{\lambda}{2u} (u-1)^2 \right] > 0, \quad u \in (0, \infty) \quad (5.37.4)$$

The probability density function g is given by

$$g(u) = \sqrt{\frac{\lambda}{2\pi u^3}} \exp \left[-\frac{\lambda}{2u} (u-1)^2 \right], \quad u \in (0, \infty) \quad (5.37.5)$$

1. g increases and then decreases with mode $u_0 = \sqrt{1 + \left(\frac{3}{2\lambda}\right)^2} - \frac{3}{2\lambda}$
2. g is concave upward then downward then upward again.

Proof

The formula for the PDF follows immediately from the proof of the CDF above, since $g = G'$. The first order properties come from

$$g'(u) = \sqrt{\frac{\lambda}{8\pi u^7}} \exp \left[-\frac{\lambda}{2u} (u-1)^2 \right] [\lambda(1-u^2) - 3u], \quad u \in (0, \infty) \quad (5.37.6)$$

and the second order properties from

$$g''(u) = \sqrt{\frac{\lambda}{32\pi u^{11}}} \exp\left[-\frac{\lambda}{2u}(u-1)^2\right] [15u^2 + \lambda^2(u^2-1)^2 + 2\lambda u(3u^2-5)], \quad u \in (0, \infty) \quad (5.37.7)$$

So g has the classic unimodal shape, but the inflection points are very complicated functions of λ . For the mode, note that $u_0 \downarrow 0$ as $\lambda \downarrow 0$ and $u_0 \uparrow 1$ as $\lambda \uparrow \infty$. The probability density function of the standard Wald distribution is

$$g(u) = \sqrt{\frac{1}{2\pi u^3}} \exp\left[-\frac{1}{2u}(u-1)^2\right], \quad u \in (0, \infty) \quad (5.37.8)$$

Open the special distribution simulator and select the Wald distribution. Vary the shape parameter and note the shape of the probability density function. For various values of the parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function of the standard Wald distribution does not have a simple closed form, so the median and other quantiles must be approximated.

Open the special distribution calculator and select the Wald distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, compute approximate values of the first quartile, the median, and the third quartile.

Moments

Suppose that random variable U has the standard Wald distribution with shape parameter $\lambda \in (0, \infty)$.

U has moment generating function m given by

$$m(t) = \mathbb{E}(e^{tU}) = \exp\left[\lambda\left(1 - \sqrt{1 - \frac{2t}{\lambda}}\right)\right], \quad t < \frac{\lambda}{2} \quad (5.37.9)$$

Proof

The proof requires some facts about the *modified Bessel function of the second kind*, denoted K_α where the parameter $\alpha \in \mathbb{R}$. This function is one of the two linearly independent solutions of the differential equation

$$x^2 \frac{d^2 y}{dx^2} + x \frac{dy}{dx} - (x^2 + \alpha^2)y = 0 \quad (5.37.10)$$

The other solution, appropriately enough, is the *modified Bessel function of the first kind*. The function of the second kind, the one that we care about here, is the solution that decays exponentially as $x \rightarrow \infty$. The first fact we need is that

$$K_{-1/2}(x) = \frac{1}{\sqrt{x}} e^{-x}, \quad x \in (0, \infty) \quad (5.37.11)$$

which you can verify by direct substitution into the differential equation. The second fact that we need is the identity

$$\int_0^\infty x^{p-1} \exp\left[-\frac{1}{2}\left(ax + \frac{b}{x}\right)\right] dx = \frac{2K_p(\sqrt{ab})}{(a/b)^{p/2}}, \quad a, b \in (0, \infty); p \in \mathbb{R} \quad (5.37.12)$$

Now, for the moment generating function of U we have

$$m(t) = \int_0^\infty e^{tx} \sqrt{\frac{\lambda}{2\pi x^3}} \exp\left[-\frac{\lambda}{2x}(x-1)^2\right] dx \quad (5.37.13)$$

Combining the exponentials and doing some algebra, we can rewrite this as

$$m(t) = \sqrt{\frac{\lambda}{2\pi}} e^\lambda \int_0^\infty x^{-3/2} \exp\left[-\frac{1}{2}(\lambda - 2t)x - \frac{1}{2}\frac{\lambda}{x}\right] dx \quad (5.37.14)$$

The integral now has the form of the identity given above with $p = -1/2$, $a = \lambda - 2t$, and $b = \lambda$. Hence we have

$$m(t) = \sqrt{\frac{\lambda}{2\pi}} e^{\lambda \frac{{}_2K_{-1/2}[\sqrt{\lambda(\lambda-2t)}]}{[(\lambda-2t)/\lambda]^{-1/4}}} \quad (5.37.15)$$

Using the explicit form of $K_{-1/2}$ given above and doing more algebra we get

$$m(t) = \exp\left[\lambda - \sqrt{\lambda(\lambda-2t)}\right] = \exp\left[\lambda \left(1 - \sqrt{1 - \frac{2t}{\lambda}}\right)\right] \quad (5.37.16)$$

Since the moment generating function is finite in an interval containing 0, the basic Wald distribution has moments of all orders.

The mean and variance of U are

1. $\mathbb{E}(U) = 1$
2. $\text{var}(U) = \frac{1}{\lambda}$

Proof

Differentiating gives

$$m'(t) = m(t) \left(1 - \frac{2t}{\lambda}\right)^{-1/2} \quad (5.37.17)$$

$$m''(t) = m(t) \left[\left(1 - \frac{2t}{\lambda}\right)^{-1} + \frac{1}{\lambda} \left(1 - \frac{2t}{\lambda}\right)^{-3/2} \right] \quad (5.37.18)$$

and hence $\mathbb{E}(U) = m'(0) = 1$ and $\mathbb{E}(U^2) = m''(0) = 1 + \frac{1}{\lambda}$.

So interestingly, the mean is 1 for all values of the shape parameter, while $\text{var}(U) \rightarrow \infty$ as $\lambda \downarrow 0$ and $\text{var}(U) \rightarrow 0$ as $\lambda \rightarrow \infty$.

Open the special distribution simulator and select the Wald distribution. Vary the shape parameter and note the size and location of the mean \pm standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of U are

1. $\text{skew}(U) = 3/\sqrt{\lambda}$
2. $\text{kurt}(U) = 3 + 15/\lambda$

Proof

The main tool is the differential equation for the [moment generating function](#) that we used in computing the [mean and variance](#):

$$m'(t) = m(t) \left(1 - \frac{2t}{\lambda}\right)^{-1/2} \quad (5.37.19)$$

Using this recursively, we can find the first four moments of U . We already know the first two: $m'(0) = \mathbb{E}(U) = 1$, $m''(0) = \mathbb{E}(U^2) = 1 + 1/\lambda$. The third and fourth are

$$m^{(3)}(0) = E(U^3) = 1 + 3/\lambda + 3/\lambda^2 \quad (5.37.20)$$

$$m^{(4)}(0) = E(U^4) = 1 + 6/\lambda + 15/\lambda^2 + 15/\lambda^3 \quad (5.37.21)$$

The results then follow from the standard computational formulas for the skewness and kurtosis in terms of the moments.

It follows that the *excess kurtosis* is $\text{kurt}(U) - 3 = 15/\lambda$. Note that $\text{skew}(U) \rightarrow \infty$ as $\lambda \rightarrow 0$ and $\text{skew}(U) \rightarrow 0$ as $\lambda \rightarrow \infty$. Similarly, $\text{kurt}(U) \rightarrow \infty$ as $\lambda \rightarrow 0$ and $\text{kurt}(U) \rightarrow 3$ as $\lambda \rightarrow \infty$.

The General Wald Distribution

The basic Wald distribution is generalized into a scale family. Scale parameters often correspond to a change of units, and so are of basic importance.

Suppose that $\lambda, \mu \in (0, \infty)$ and that U has the basic Wald distribution with shape parameter λ/μ . Then $X = \mu U$ has the Wald distribution with shape parameter λ and mean μ .

Justification for the name of the parameter μ as the *mean* is given [below](#). Note that the generalization is consistent—when $\mu = 1$ we have the basic Wald distribution with shape parameter λ .

Distribution Functions

Suppose that X has the Wald distribution with shape parameter $\lambda \in (0, \infty)$ and mean $\mu \in (0, \infty)$. Again, we let Φ denote the standard normal distribution function.

X has distribution function F given by

$$F(x) = \Phi \left[\sqrt{\frac{\lambda}{x}} \left(\frac{x}{\mu} - 1 \right) \right] + \exp \left(\frac{2\lambda}{\mu} \right) \Phi \left[-\sqrt{\frac{\lambda}{x}} \left(\frac{x}{\mu} + 1 \right) \right], \quad x \in (0, \infty) \quad (5.37.22)$$

Proof

Recall that the CDF F of X is related to the CDF G of U by

$$F(x) = G \left(\frac{x}{\mu} \right), \quad x \in (0, \infty) \quad (5.37.23)$$

so the result follows from the [CDF](#) above, with λ replaced by λ/μ , and x with x/μ .

X has probability density function f given by

$$f(x) = \sqrt{\frac{\lambda}{2\pi x^3}} \exp \left[\frac{-\lambda(x - \mu)^2}{2\mu^2 x} \right], \quad x \in (0, \infty) \quad (5.37.24)$$

1. f increases and then decreases with mode $x_0 = \mu \left[\sqrt{1 + \left(\frac{3\mu}{2\lambda} \right)^2} - \frac{3\mu}{2\lambda} \right]$
2. f is concave upward then downward then upward again.

Proof

Recall that the PDF f of X is related to the PDF g of U by

$$f(x) = \frac{1}{\mu} g \left(\frac{x}{\mu} \right), \quad x \in (0, \infty) \quad (5.37.25)$$

Hence the result follows from the [PDF](#) above with λ replaced by λ/μ and x with x/μ .

Once again, the graph of f has the classic unimodal shape, but the inflection points are complicated functions of the parameters.

Open the special distribution simulator and select the Wald distribution. Vary the parameters and note the shape of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Again, the quantile function cannot be expressed in a simple closed form, so the median and other quantiles must be approximated.

Open the special distribution calculator and select the Wald distribution. Vary the parameters and note the shape of the distribution and density functions. For selected values of the parameters, compute approximate values of the first quartile, the median, and the third quartile.

Moments

Suppose again that X has the Wald distribution with shape parameter $\lambda \in (0, \infty)$ and mean $\mu \in (0, \infty)$. By [definition](#), we can take $X = \mu U$ where U has the basic Wald distribution with shape parameter λ/μ .

X has moment generating function M given by

$$M(t) = \exp \left[\frac{\lambda}{\mu} \left(1 - \sqrt{1 - \frac{2\mu^2 t}{\lambda}} \right) \right], \quad t < \frac{\lambda}{2\mu^2} \quad (5.37.26)$$

Proof

Recall that the MGF M of X is related to the MGF m of U by $M(t) = m(t\mu)$. Hence the result follows from the result [MFG](#) above with λ replaced by λ/μ and t with $t\mu$.

As promised, the parameter μ is the mean of Wald distribution.

The mean and variance of X are

1. $\mathbb{E}(X) = \mu$
2. $\text{var}(X) = \mu^3/\lambda$

Proof

From the results for the [mean and variance](#) above and basic properties of expected value and variance, we have $\mathbb{E}(X) = \mu\mathbb{E}(U) = \mu \cdot 1$ and $\text{var}(X) = \mu^2 \text{var}(U) = \mu^2 \frac{\mu}{\lambda}$.

Open the special distribution simulator and select the Wald distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of X are

1. $\text{skew}(X) = 3\sqrt{\mu/\lambda}$
2. $\text{kurt}(U) = 3 + 15\mu/\lambda$

Proof

Skewness and kurtosis are invariant under scale transformations, so $\text{skew}(X) = \text{skew}(U)$ and $\text{kurt}(X) = \text{kurt}(U)$. The results then follow from the [skewness and kurtosis](#) above, with λ replaced by λ/μ .

Related Distribution

As noted earlier, the Wald distribution is a scale family, although neither of the parameters is a scale parameter.

Suppose that X has the Wald distribution with shape parameters $\lambda \in (0, \infty)$ and mean $\mu \in (0, \infty)$ and that $c \in (0, \infty)$. Then $Y = cX$ has the Wald distribution with shape parameter $c\lambda$ and mean $c\mu$.

Proof

By [definition](#), we can take $X = \mu U$ where U has the basic Wald distribution with shape parameter λ/μ . Then $Y = cX = (c\mu)U$. Since U has shape parameter $c\lambda/c\mu$, the result follows from the definition.

For the next result, it's helpful to re-parameterize the Wald distribution with the mean μ and the ratio $r = \lambda/\mu^2 \in (0, \infty)$. This parametrization is clearly equivalent, since we can recover the shape parameter from the mean and ratio as $\lambda = r\mu^2$. Note also that $r = \mathbb{E}(X)/\text{var}(X)$, the ratio of the mean to the variance. Finally, note that the [moment generating function](#) above becomes

$$M(t) = \exp \left[r\mu \left(1 - \sqrt{1 - \frac{2}{r}t} \right) \right], \quad t < \frac{r}{2} \quad (5.37.27)$$

and of course, this function characterizes the Wald distribution with this parametrization. Our next result is that the Wald distribution is closed under convolution (corresponding to sums of independent variables) when the ratio is fixed.

Suppose that X_1 has the Wald distribution with mean $\mu_1 \in (0, \infty)$ and ratio $r \in (0, \infty)$; X_2 has the Wald distribution with mean $\mu_2 \in (0, \infty)$ and ratio r ; and that X_1 and X_2 are independent. Then $Y = X_1 + X_2$ has the Wald distribution with mean $\mu_1 + \mu_2$ and ratio r .

Proof

For $i \in \{1, 2\}$, the MGF of X_i is

$$M_i(t) = \exp \left[r \mu_i \left(1 - \sqrt{1 - \frac{2}{r} t} \right) \right], \quad t < \frac{r}{2} \quad (5.37.28)$$

Hence the MGF of $Y = X_1 + X_2$ is

$$M(t) = M_1(t)M_2(t) = \exp \left[r (\mu_1 + \mu_2) \left(1 - \sqrt{1 - \frac{2}{r} t} \right) \right], \quad t < \frac{r}{2} \quad (5.37.29)$$

Hence Y has the Wald distribution with mean $\mu_1 + \mu_2$ and ratio r .

In the previous result, note that the shape parameter of X_1 is $r\mu_1^2$, the shape parameter of X_2 is $r\mu_2^2$, and the shape parameter of Y is $\lambda = r(\mu_1 + \mu_2)^2$. Also, of course, the result generalizes to a sum of any finite number of independent Wald variables, as long as the ratio is fixed. The next couple of results are simple corollaries.

Suppose that (X_1, X_2, \dots, X_n) is a sequence of independent variables, each with the Wald distribution with shape parameter $\lambda \in (0, \infty)$ and mean $\mu \in (0, \infty)$. Then

1. $Y_n = \sum_{i=1}^n X_i$ has the Wald distribution with shape parameter $n^2 \lambda$ and mean $n\mu$.
2. $M_n = \frac{1}{n} \sum_{i=1}^n X_i$ has the Wald distribution with shape parameter $n\lambda$ and mean μ .

Proof

1. This follows from the [previous result](#) and induction. The mean of Y_n of course is $n\mu$. The common ratio is $r = \lambda/\mu^2$, and hence the shape parameter of Y_n is $(\lambda/\mu^2)(n\mu)^2 = n^2 \lambda$.
2. This follows from (a) and the [scaling result](#) above. The mean of M_n of course is μ and the shape parameter is $(1/n)(n^2 \lambda) = n\lambda$.

In the context of the previous theorem, (X_1, X_2, \dots, X_n) is a random sample of size n from the Wald distribution, and $\frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean. The Wald distribution is infinitely divisible:

Suppose that X has the Wald distribution with shape parameter $\lambda \in (0, \infty)$ and mean $\mu \in (0, \infty)$. For every $n \in \mathbb{N}_+$, X has the same distribution as $\sum_{i=1}^n X_i$ where (X_1, X_2, \dots, X_n) are independent, and each has the Wald distribution with shape parameter λ/n^2 and mean μ/n .

The Lévy distribution is related to the Wald distribution, not surprising since the Lévy distribution governs the first time that a standard Brownian motion hits a fixed positive value.

For fixed $\lambda \in (0, \infty)$, the Wald distribution with shape parameter λ and mean $\mu \in (0, \infty)$ converges to the Lévy distribution with location parameter 0 and scale parameter λ as $\mu \rightarrow \infty$.

Proof

From the formula for the [CDF](#) above, note that

$$F(x) \rightarrow \Phi \left(-\sqrt{\frac{\lambda}{x}} \right) + \Phi \left(-\sqrt{\frac{\lambda}{x}} \right) = 2\Phi \left(-\sqrt{\frac{\lambda}{x}} \right) = 2 \left[1 - \Phi \left(\sqrt{\frac{\lambda}{x}} \right) \right] \text{ as } \mu \rightarrow \infty \quad (5.37.30)$$

But the last expression is the distribution function of the Lévy distribution with location parameter 0 and shape parameter λ .

The other limiting distribution, this time with the mean fixed, is less interesting.

For fixed $\mu \in (0, \infty)$, the Wald distribution with shape parameter $\lambda \in (0, \infty)$ and mean μ converges to point mass at μ and variance 1 as $\lambda \rightarrow \infty$.

Proof

This time, it's better to use M , the [moment generating function](#) above. By rationalizing we see that for fixed $\mu \in (0, \infty)$ and $t \in \mathbb{R}$

$$\frac{\lambda}{\mu} \left(1 - \sqrt{1 - \frac{2\mu^2 t}{\lambda}} \right) = \frac{2\mu t}{1 + \sqrt{1 - 2\mu^2 t/\lambda}} \rightarrow \mu t \text{ as } \lambda \rightarrow \infty \quad (5.37.31)$$

Hence $M(t) \rightarrow e^{\mu t}$ as $\lambda \rightarrow \infty$ and the limit is the MGF of the constant random variable μ .

Finally, the Wald distribution is a member of the general exponential family of distributions.

The Wald distribution is a general exponential distribution with natural parameters $-\lambda/(2\mu^2)$ and $-\lambda/2$, and natural statistics X and $1/X$.

Proof

This follows from the PDF f above. If we expand the square and simplify, we can write f in the form

$$f(x) = \sqrt{\frac{\lambda}{2\pi}} \exp\left(\frac{\lambda}{2\mu}\right) x^{-3/2} \exp\left(-\frac{\lambda}{2\mu^2}x - \frac{\lambda}{2}\frac{1}{x}\right), \quad x \in (0, \infty) \quad (5.37.32)$$

This page titled [5.37: The Wald Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.38: The Weibull Distribution

In this section, we will study a two-parameter family of distributions that has special importance in reliability.

The Basic Weibull Distribution

Distribution Functions

The *basic Weibull distribution* with *shape parameter* $k \in (0, \infty)$ is a continuous distribution on $[0, \infty)$ with distribution function G given by

$$G(t) = 1 - \exp(-t^k), \quad t \in [0, \infty) \quad (5.38.1)$$

The special case $k = 1$ gives the *standard Weibull distribution*.

Proof

Clearly G is continuous and increasing on $[0, \infty)$ with $G(0) = 0$ and $G(t) \rightarrow 1$ as $t \rightarrow \infty$.

The Weibull distribution is named for Waloddi Weibull. Weibull was not the first person to use the distribution, but was the first to study it extensively and recognize its wide use in applications. The standard Weibull distribution is the same as the standard exponential distribution. But as we will see, every Weibull random variable can be obtained from a standard Weibull variable by a simple deterministic transformation, so the terminology is justified.

The probability density function g is given by

$$g(t) = kt^{k-1} \exp(-t^k), \quad t \in (0, \infty) \quad (5.38.2)$$

1. If $0 < k < 1$, g is decreasing and concave upward with $g(t) \rightarrow \infty$ as $t \downarrow 0$.
2. If $k = 1$, g is decreasing and concave upward with mode $t = 0$.
3. If $k > 1$, g increases and then decreases, with mode $t = \left(\frac{k-1}{k}\right)^{1/k}$.
4. If $1 < k \leq 2$, g is concave downward and then upward, with inflection point at $t = \left[\frac{3(k-1) + \sqrt{(5k-1)(k-1)}}{2k}\right]^{1/k}$.
5. If $k > 2$, g is concave upward, then downward, then upward again, with inflection points at $t = \left[\frac{3(k-1) \pm \sqrt{(5k-1)(k-1)}}{2k}\right]^{1/k}$.

Proof

These results follow from basic calculus. The PDF is $g = G'$ where G is the CDF above. The first order properties come from

$$g'(t) = kt^{k-2} \exp(-t^k) [-kt^k + (k-1)] \quad (5.38.3)$$

The second order properties come from

$$g''(t) = kt^{k-3} \exp(-t^k) [k^2 t^{2k} - 3k(k-1)t^k + (k-1)(k-2)] \quad (5.38.4)$$

So the Weibull density function has a rich variety of shapes, depending on the shape parameter, and has the classic unimodal shape when $k > 1$. If $k \geq 1$, g is defined at 0 also.

In the special distribution simulator, select the Weibull distribution. Vary the shape parameter and note the shape of the probability density function. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function G^{-1} is given by

$$G^{-1}(p) = [-\ln(1-p)]^{1/k}, \quad p \in [0, 1) \quad (5.38.5)$$

1. The first quartile is $q_1 = (\ln 4 - \ln 3)^{1/k}$.
2. The median is $q_2 = (\ln 2)^{1/k}$.

3. The third quartile is $q_3 = (\ln 4)^{1/k}$.

Proof

The formula for $G^{-1}(p)$ comes from solving $G(t) = p$ for t in terms of p .

Open the special distribution calculator and select the Weibull distribution. Vary the shape parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, compute the median and the first and third quartiles.

The reliability function G^c is given by

$$G^c(t) = \exp(-t^k), \quad t \in [0, \infty) \quad (5.38.6)$$

Proof

This follows trivially from the [CDF](#) above, since $G^c = 1 - G$.

The failure rate function r is given by

$$r(t) = kt^{k-1}, \quad t \in (0, \infty) \quad (5.38.7)$$

1. If $0 < k < 1$, r is decreasing with $r(t) \rightarrow \infty$ as $t \downarrow 0$ and $r(t) \rightarrow 0$ as $t \rightarrow \infty$.
2. If $k = 1$, r is constant 1.
3. If $k > 1$, r is increasing with $r(0) = 0$ and $r(t) \rightarrow \infty$ as $t \rightarrow \infty$.

Proof

The formula for r follows immediately from the [PDF](#) g and the [reliability function](#) G^c given above, since $r = g/G^c$.

Thus, the Weibull distribution can be used to model devices with decreasing failure rate, constant failure rate, or increasing failure rate. This versatility is one reason for the wide use of the Weibull distribution in reliability. If $k \geq 1$, r is defined at 0 also.

Moments

Suppose that Z has the basic Weibull distribution with shape parameter $k \in (0, \infty)$. The moments of Z , and hence the mean and variance of Z can be expressed in terms of the gamma function Γ

$$\mathbb{E}(Z^n) = \Gamma\left(1 + \frac{n}{k}\right) \text{ for } n \geq 0.$$

Proof

For $n \geq 0$,

$$\mathbb{E}(Z^n) = \int_0^\infty t^n kt^{k-1} \exp(-t^k) dt \quad (5.38.8)$$

Substituting $u = t^k$ gives

$$\mathbb{E}(Z^n) = \int_0^\infty u^{n/k} e^{-u} du = \Gamma\left(1 + \frac{n}{k}\right) \quad (5.38.9)$$

So the Weibull distribution has moments of all orders. The moment generating function, however, does not have a simple, closed expression in terms of the usual elementary functions.

In particular, the mean and variance of Z are

1. $\mathbb{E}(Z) = \Gamma\left(1 + \frac{1}{k}\right)$
2. $\text{var}(Z) = \Gamma\left(1 + \frac{2}{k}\right) - \Gamma^2\left(1 + \frac{1}{k}\right)$

Note that $\mathbb{E}(Z) \rightarrow 1$ and $\text{var}(Z) \rightarrow 0$ as $k \rightarrow \infty$. We will learn more about the [limiting distribution](#) below.

In the special distribution simulator, select the Weibull distribution. Vary the shape parameter and note the size and location of the mean \pm standard deviation bar. For selected values of the shape parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis also follow easily from the [general moment result](#) above, although the formulas are not particularly helpful.

Skewness and kurtosis

1. The skewness of Z is

$$\text{skew}(Z) = \frac{\Gamma(1+3/k) - 3\Gamma(1+1/k)\Gamma(1+2/k) + 2\Gamma^3(1+1/k)}{[\Gamma(1+2/k) - \Gamma^2(1+1/k)]^{3/2}} \quad (5.38.10)$$

2. The kurtosis of Z is

$$\text{kurt}(Z) = \frac{\Gamma(1+4/k) - 4\Gamma(1+1/k)\Gamma(1+3/k) + 6\Gamma^2(1+1/k)\Gamma(1+2/k) - 3\Gamma^4(1+1/k)}{[\Gamma(1+2/k) - \Gamma^2(1+1/k)]^2} \quad (5.38.11)$$

Proof

The results follow directly from the [general moment result](#) and the computational formulas for skewness and kurtosis.

Related Distributions

As noted above, the standard Weibull distribution (shape parameter 1) is the same as the standard exponential distribution. More generally, any basic Weibull variable can be constructed from a standard exponential variable.

Suppose that $k \in (0, \infty)$.

1. If U has the standard exponential distribution then $Z = U^{1/k}$ has the basic Weibull distribution with shape parameter k .
2. If Z has the basic Weibull distribution with shape parameter k then $U = Z^k$ has the standard exponential distribution.

Proof

We use distribution functions. The basic Weibull [CDF](#) is given above; the standard exponential CDF is $u \mapsto 1 - e^{-u}$ on $[0, \infty)$. Note that the inverse transformations $z = u^k$ and $u = z^{1/k}$ are strictly increasing and map $[0, \infty)$ onto $[0, \infty)$.

1. $\mathbb{P}(Z \leq z) = \mathbb{P}(U \leq z^k) = 1 - \exp(-z^k)$ for $z \in [0, \infty)$.
2. $\mathbb{P}(U \leq u) = \mathbb{P}(Z \leq u^{1/k}) = 1 - \exp[-(u^{1/k})^k] = 1 - e^{-u}$ for $u \in [0, \infty)$.

The basic Weibull distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) given above.

Suppose that $k \in (0, \infty)$.

1. If U has the standard uniform distribution then $Z = (-\ln U)^{1/k}$ has the basic Weibull distribution with shape parameter k .
2. If Z has the basic Weibull distribution with shape parameter k then $U = \exp(-Z^k)$ has the standard uniform distribution.

Proof

Let G denote the [CDF](#) of the basic Weibull distribution with shape parameter k and G^{-1} the corresponding [quantile function](#), given above.

1. If U has the standard uniform distribution then so does $1 - U$. Hence $Z = G^{-1}(1 - U) = (-\ln U)^{1/k}$ has the basic Weibull distribution with shape parameter k .
2. If Z has the basic Weibull distribution with shape parameter k then $G(Z)$ has the standard uniform distribution. But then so does $U = 1 - G(Z) = \exp(-Z^k)$.

Since the quantile function has a simple, closed form, the basic Weibull distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the Weibull distribution. Vary the shape parameter and note again the shape of the distribution and density functions. For selected values of the parameter, run the simulation 1000 times and compare the empirical density, mean, and standard deviation to their distributional counterparts.

The limiting distribution with respect to the shape parameter is concentrated at a single point.

The basic Weibull distribution with shape parameter $k \in (0, \infty)$ converges to point mass at 1 as $k \rightarrow \infty$.

Proof

Once again, let G denote the basic Weibull CDF with shape parameter k given above. Note that $G(t) \rightarrow 0$ as $k \rightarrow \infty$ for $0 \leq t < 1$; $G(1) = 1 - e^{-1}$ for all k ; and $G(t) \rightarrow 1$ as $k \rightarrow \infty$ for $t > 1$. Except for the point of discontinuity $t = 1$, the limits are the CDF of point mass at 1.

The General Weibull Distribution

Like most special continuous distributions on $[0, \infty)$, the basic Weibull distribution is generalized by the inclusion of a scale parameter. A scale transformation often corresponds in applications to a change of units, and for the Weibull distribution this usually means a change in time units.

Suppose that Z has the basic Weibull distribution with shape parameter $k \in (0, \infty)$. For $b \in (0, \infty)$, random variable $X = bZ$ has the Weibull distribution with shape parameter k and scale parameter b .

Generalizations of the results given above follow easily from basic properties of the scale transformation.

Distribution Functions

Suppose that X has the Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$.

X distribution function F given by

$$F(t) = 1 - \exp\left[-\left(\frac{t}{b}\right)^k\right], \quad t \in [0, \infty) \quad (5.38.12)$$

Proof

Recall that $F(t) = G\left(\frac{t}{b}\right)$ for $t \in [0, \infty)$ where G is the CDF of the basic Weibull distribution with shape parameter k , given above.

X has probability density function f given by

$$f(t) = \frac{k}{b^k} t^{k-1} \exp\left[-\left(\frac{t}{b}\right)^k\right], \quad t \in (0, \infty) \quad (5.38.13)$$

1. If $0 < k < 1$, f is decreasing and concave upward with $f(t) \rightarrow \infty$ as $t \downarrow 0$.
2. If $k = 1$, f is decreasing and concave upward with mode $t = 0$.
3. If $k > 1$, f increases and then decreases, with mode $t = b\left(\frac{k-1}{k}\right)^{1/k}$.
4. If $1 < k \leq 2$, f is concave downward and then upward, with inflection point at $t = b\left[\frac{3(k-1) + \sqrt{(5k-1)(k-1)}}{2k}\right]^{1/k}$.
5. If $k > 2$, f is concave upward, then downward, then upward again, with inflection points at $t = b\left[\frac{3(k-1) \pm \sqrt{(5k-1)(k-1)}}{2k}\right]^{1/k}$.

Proof

Recall that $f(t) = \frac{1}{b}g\left(\frac{t}{b}\right)$ for $t \in (0, \infty)$ where g is the PDF of the corresponding basic Weibull distribution given above.

Open the special distribution simulator and select the Weibull distribution. Vary the parameters and note the shape of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

X has quantile function F^{-1} given by

$$F^{-1}(p) = b[-\ln(1-p)]^{1/k}, \quad p \in [0, 1) \quad (5.38.14)$$

1. The first quartile is $q_1 = b(\ln 4 - \ln 3)^{1/k}$.
2. The median is $q_2 = b(\ln 2)^{1/k}$.
3. The third quartile is $q_3 = b(\ln 4)^{1/k}$.

Proof

Recall that $F^{-1}(p) = bG^{-1}(p)$ for $p \in [0, 1)$ where G^{-1} is the [quantile function](#) of the corresponding basic Weibull distribution given above.

Open the special distribution calculator and select the Weibull distribution. Vary the parameters and note the shape of the distribution and probability density functions. For selected values of the parameters, compute the median and the first and third quartiles.

X has reliability function F^c given by

$$F^c(t) = \exp\left[-\left(\frac{t}{b}\right)^k\right], \quad t \in [0, \infty) \quad (5.38.15)$$

Proof

This follows trivially from the [CDF](#) F given above, since $F^c = 1 - F$.

As before, the Weibull distribution has decreasing, constant, or increasing failure rates, depending only on the shape parameter.

X has failure rate function R given by

$$R(t) = \frac{kt^{k-1}}{b^k}, \quad t \in (0, \infty) \quad (5.38.16)$$

1. If $0 < k < 1$, R is decreasing with $R(t) \rightarrow \infty$ as $t \downarrow 0$ and $R(t) \rightarrow 0$ as $t \rightarrow \infty$.
2. If $k = 1$, R is constant $\frac{1}{b}$.
3. If $k > 1$, R is increasing with $R(0) = 0$ and $R(t) \rightarrow \infty$ as $t \rightarrow \infty$.

Moments

Suppose again that X has the Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. Recall that by [definition](#), we can take $X = bZ$ where Z has the basic Weibull distribution with shape parameter k .

$$\mathbb{E}(X^n) = b^n \Gamma\left(1 + \frac{n}{k}\right) \text{ for } n \geq 0.$$

Proof

The result then follows from the [moments of \$Z\$](#) above, since $\mathbb{E}(X^n) = b^n \mathbb{E}(Z^n)$.

In particular, the mean and variance of X are

1. $\mathbb{E}(X) = b \Gamma\left(1 + \frac{1}{k}\right)$
2. $\text{var}(X) = b^2 \left[\Gamma\left(1 + \frac{2}{k}\right) - \Gamma^2\left(1 + \frac{1}{k}\right) \right]$

Note that $\mathbb{E}(X) \rightarrow b$ and $\text{var}(X) \rightarrow 0$ as $k \rightarrow \infty$.

Open the special distribution simulator and select the Weibull distribution. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Skewness and kurtosis

1. The skewness of X is

$$\text{skew}(X) = \frac{\Gamma(1 + 3/k) - 3\Gamma(1 + 1/k)\Gamma(1 + 2/k) + 2\Gamma^3(1 + 1/k)}{[\Gamma(1 + 2/k) - \Gamma^2(1 + 1/k)]^{3/2}} \quad (5.38.17)$$

2. The kurtosis of X is

$$\text{kurt}(X) = \frac{\Gamma(1 + 4/k) - 4\Gamma(1 + 1/k)\Gamma(1 + 3/k) + 6\Gamma^2(1 + 1/k)\Gamma(1 + 2/k) - 3\Gamma^4(1 + 1/k)}{[\Gamma(1 + 2/k) - \Gamma^2(1 + 1/k)]^2} \quad (5.38.18)$$

Proof

Skewness and kurtosis depend only on the standard score of the random variable, and hence are invariant under scale transformations. So the results are the same as the [skewness and kurtosis of \$Z\$](#) .

Related Distributions

Since the Weibull distribution is a scale family for each value of the shape parameter, it is trivially closed under scale transformations.

Suppose that X has the Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. If $c \in (0, \infty)$ then $Y = cX$ has the Weibull distribution with shape parameter k and scale parameter bc .

Proof

By [definition](#), we can take $X = bZ$ where Z has the basic Weibull distribution with shape parameter k . But then $Y = cX = (bc)Z$.

The exponential distribution is a special case of the Weibull distribution, the case corresponding to constant failure rate.

The Weibull distribution with shape parameter 1 and scale parameter $b \in (0, \infty)$ is the exponential distribution with scale parameter b .

Proof

When $k = 1$, the [Weibull CDF](#) F is given by $F(t) = 1 - e^{-t/b}$ for $t \in [0, \infty)$. But this is also the CDF of the exponential distribution with scale parameter b .

More generally, any Weibull distributed variable can be constructed from the standard variable. The following result is a simple generalization of the [connection](#) between the basic Weibull distribution and the exponential distribution.

Suppose that $k, b \in (0, \infty)$.

1. If X has the standard exponential distribution (parameter 1), then $Y = bX^{1/k}$ has the Weibull distribution with shape parameter k and scale parameter b .
2. If Y has the Weibull distribution with shape parameter k and scale parameter b , then $X = (Y/b)^k$ has the standard exponential distribution.

Proof

The results are a simple consequence of the [corresponding result](#) above

1. If X has the standard exponential distribution then $X^{1/k}$ has the basic Weibull distribution with shape parameter k , and hence $Y = bX^{1/k}$ has the Weibull distribution with shape parameter k and scale parameter b .
2. If Y has the Weibull distribution with shape parameter k and scale parameter b then Y/b has the basic Weibull distribution with shape parameter k , and hence $X = (Y/b)^k$ has the standard exponential distribution.

The Rayleigh distribution, named for William Strutt, Lord Rayleigh, is also a special case of the Weibull distribution.

The Rayleigh distribution with scale parameter $b \in (0, \infty)$ is the Weibull distribution with shape parameter 2 and scale parameter $\sqrt{2}b$.

Proof

The Rayleigh distribution with scale parameter b has CDF F given by

$$F(x) = 1 - \exp\left(-\frac{x^2}{2b^2}\right), \quad x \in [0, \infty) \quad (5.38.19)$$

But this is also the [Weibull CDF](#) with shape parameter 2 and scale parameter $\sqrt{2}b$.

Recall that the minimum of independent, exponentially distributed variables also has an exponential distribution (and the rate parameter of the minimum is the sum of the rate parameters of the variables). The Weibull distribution has a similar, but more restricted property.

Suppose that (X_1, X_2, \dots, X_n) is an independent sequence of variables, each having the Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. Then $U = \min\{X_1, X_2, \dots, X_n\}$ has the Weibull distribution with shape parameter k and scale parameter $b/n^{1/k}$.

Proof

Recall that the reliability function of the minimum of independent variables is the product of the reliability functions of the variables. It follows that U has reliability function given by

$$\mathbb{P}(U > t) = \left\{ \exp\left[-\left(\frac{t}{b}\right)^k\right] \right\}^n = \exp\left[-n\left(\frac{t}{b}\right)^k\right] = \exp\left[-\left(\frac{t}{b/n^{1/k}}\right)^k\right], \quad t \in [0, \infty) \quad (5.38.20)$$

and so the result follows.

As before, Weibull distribution has the usual connections with the standard uniform distribution by means of the [distribution function](#) and the [quantile function](#) given above..

Suppose that $k, b \in (0, \infty)$.

1. If U has the standard uniform distribution then $X = b(-\ln U)^{1/k}$ has the Weibull distribution with shape parameter k and scale parameter b .
2. If X has the basic Weibull distribution with shape parameter k then $U = \exp[-(X/b)^k]$ has the standard uniform distribution.

Proof

Let F denote the [Weibull CDF](#) with shape parameter k and scale parameter b and so that F^{-1} is the [corresponding quantile function](#).

1. If U has the standard uniform distribution then so does $1 - U$. Hence $X = F^{-1}(1 - U) = b(-\ln U)^{1/k}$ has the Weibull distribution with shape parameter k and scale parameter b .
2. If X has the Weibull distribution with shape parameter k and scale parameter b then $F(X)$ has the standard uniform distribution. But then so does $U = 1 - F(X) = \exp[-(X/b)^k]$.

Again, since the quantile function has a simple, closed form, the Weibull distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the Weibull distribution. Vary the parameters and note again the shape of the distribution and density functions. For selected values of the parameters, run the simulation 1000 times and compare the empirical density, mean, and standard deviation to their distributional counterparts.

The limiting distribution with respect to the shape parameter is concentrated at a single point.

The Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ converges to point mass at b as $k \rightarrow \infty$.

Proof

If X has the Weibull distribution with shape parameter k and scale parameter b , then we can write $X = bZ$ where Z has the basic Weibull distribution with shape parameter k . We showed above that the distribution of Z [converges to point mass at 1](#), so by the continuity theorem for convergence in distribution, the distribution of X converges to point mass at b .

Finally, the Weibull distribution is a member of the family of general exponential distributions if the shape parameter is fixed.

Suppose that X has the Weibull distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. For fixed k , X has a general exponential distribution with respect to b , with natural parameter $k - 1$ and natural statistics $\ln X$.

Proof

This follows from the definition of the general exponential distribution, since the [Weibull PDF](#) can be written in the form

$$f(t) = \frac{k}{b^k} \exp(-t^k) \exp[(k-1) \ln t], \quad t \in (0, \infty) \quad (5.38.21)$$

Computational Exercises

The lifetime T of a device (in hours) has the Weibull distribution with shape parameter $k = 1.2$ and scale parameter $b = 1000$.

1. Find the probability that the device will last at least 1500 hours.
2. Approximate the mean and standard deviation of T .
3. Compute the failure rate function.

Answer

1. $\mathbb{P}(T > 1500) = 0.1966$
2. $\mathbb{E}(T) = 940.7$, $\text{sd}(T) = 787.2$
3. $h(t) = 0.000301t^{0.2}$

This page titled [5.38: The Weibull Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.39: Benford's Law

Benford's law refers to probability distributions that seem to govern the significant digits in real data sets. The law is named for the American physicist and engineer Frank Benford, although the “law” was actually discovered earlier by the astronomer and mathematician Simon Newcomb.

To understand Benford's law, we need some preliminaries. Recall that a positive real number x can be written uniquely in the form $x = y \cdot 10^n$ (sometimes called *scientific notation*) where $y \in [\frac{1}{10}, 1)$ is the *mantissa* and $n \in \mathbb{Z}$ is the *exponent* (both of these terms are *base 10*, of course). Note that

$$\log x = \log y + n \quad (5.39.1)$$

where the logarithm function is the base 10 *common logarithm* instead of the usual base e *natural logarithm*. In the old days BC (before calculators), one would compute the logarithm of a number by looking up the logarithm of the mantissa in a *table of logarithms*, and then adding the exponent. Of course, these remarks apply to any base $b > 1$, not just base 10. Just replace 10 with b and the common logarithm with the base b logarithm.

Distribution of the Mantissa

Distribution Functions

Suppose now that X is a number selected at random from a certain data set of positive numbers. Based on empirical evidence from a number of different types of data, Newcomb, and later Benford, noticed that the mantissa Y of X seemed to have distribution function $F(y) = 1 + \log_b y$ for $y \in [1/b, 1)$. We will generalize this to an arbitrary base $b > 1$.

The *Benford mantissa distribution* with base $b \in (1, \infty)$, is a continuous distribution on $[1/b, 1)$ with distribution function F given by

$$F(y) = 1 + \log_b y, \quad y \in [1/b, 1) \quad (5.39.2)$$

The special case $b = 10$ gives the *standard Benford mantissa distribution*.

Proof

Note that F is continuous and strictly increasing on $[1/b, 1)$ with $F(1/b) = 0$ and $F(1) = 1$.

The probability density function f is given by

$$f(y) = \frac{1}{y \ln b}, \quad y \in [1/b, 1) \quad (5.39.3)$$

1. f is decreasing with mode $y = \frac{1}{b}$.
2. f is concave upward.

Proof

These results follow from the [CDF \$F\$](#) above and standard calculus. Recall that $f = F'$.

Open the [Special Distribution Simulator](#) and select the Benford mantissa distribution. Vary the base b and note the shape of the probability density function. For various values of b , run the simulation 1000 times and compare the empirical density function to the probability density function.

The quantile function F^{-1} is given by

$$F^{-1}(p) = \frac{1}{b^{1-p}}, \quad p \in [0, 1] \quad (5.39.4)$$

1. The first quartile is $F^{-1}(\frac{1}{4}) = \frac{1}{b^{3/4}}$
2. The median is $F^{-1}(\frac{1}{2}) = \frac{1}{\sqrt{b}}$
3. The third quartile is $F^{-1}(\frac{3}{4}) = \frac{1}{b^{1/4}}$

Proof

The formula for $F^{-1}(p)$ follows by solving $F(x) = p$ for x in terms of p .

Numerical values of the quartiles for the standard (base 10) distribution are given in an [exercise below](#).

Open the special distribution calculator and select the Benford mantissa distribution. Vary the base and note the shape and location of the distribution and probability density functions. For selected values of the base, compute the median and the first and third quartiles.

Moments

Assume that Y has the Benford mantissa distribution with base $b \in (1, \infty)$.

The moments of Y are

$$\mathbb{E}(Y^n) = \frac{b^n - 1}{nb^n \ln b}, \quad n \in (0, \infty) \quad (5.39.5)$$

Proof

For $n > 0$,

$$\mathbb{E}(Y^n) = \int_{1/b}^1 y^n \frac{1}{y \ln b} dy = \frac{1}{\ln b} \int_{1/b}^1 y^{n-1} dy = \frac{1 - 1/b^n}{n \ln(b)} \quad (5.39.6)$$

Note that for fixed $n > 0$, $\mathbb{E}(Y^n) \rightarrow 1$ as $b \downarrow 1$ and $\mathbb{E}(Y^n) \rightarrow 0$ as $b \rightarrow \infty$. We will learn more about the [limiting distribution](#) below. The mean and variance follow easily from the general moment result.

Mean and variance

1. The mean of Y is

$$\mathbb{E}(Y) = \frac{b-1}{b \ln b} \quad (5.39.7)$$

2. the variance of Y is

$$\text{var}(Y) = \frac{b-1}{b^2 \ln b} \left[\frac{b+1}{2} - \frac{b-1}{\ln b} \right] \quad (5.39.8)$$

Numerical values of the mean and variance for the standard (base 10) distribution are given in an [exercise below](#).

In the Special Distribution Simulator, select the Benford mantissa distribution. Vary the base b and note the size and location of the mean \pm standard deviation bar. For selected values of b , run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Related Distributions

The Benford mantissa distribution has the usual connections to the standard uniform distribution by means of the [distribution function](#) and [quantile function](#) given above.

Suppose that $b \in (1, \infty)$.

1. If U has the standard uniform distribution then $Y = b^{-U}$ has the Benford mantissa distribution with base b .
2. If Y has the Benford mantissa distribution with base b then $U = -\log_b Y$ has the standard uniform distribution.

Proof

1. If U has the standard uniform distribution then so does $1 - U$ and hence $Y = F^{-1}(1 - U) = b^{-U}$ has the Benford mantissa distribution with base b .
2. The CDF F is strictly increasing on $[b^{-1}, 1)$. Hence if Y has the Benford mantissa distribution with base b then $F(Y) = 1 + \log_b(Y)$ has the standard uniform distribution and hence so does $1 - F(Y) = -\log_b(Y)$.

Since the quantile function has a simple closed form, the Benford mantissa distribution can be simulated using the random quantile method.

Open the random quantile experiment and select the Benford mantissa distribution. Vary the base b and note again the shape and location of the distribution and probability density functions. For selected values of b , run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

Also of interest, of course, are the limiting distributions of Y with respect to the base b .

The Benford mantissa distribution with base $b \in (1, \infty)$ converges to

1. Point mass at 1 as $b \downarrow 1$.
2. Point mass at 0 as $b \uparrow \infty$.

Proof

Note that the CDF of Y above can be written as $F(y) = 1 + \ln(y)/\ln(b)$ for $1/b \leq y < 1$, and of course we also have $F(y) = 0$ for $y < 1/b$ and $F(y) = 1$ for $y \geq 1$.

1. As $b \downarrow 1$, $1/b \uparrow 1$, and $1 + \ln(y)/\ln(b) \rightarrow 1$, so in the limit we have $F(y) = 0$ for $y < 1$ and $F(y) = 1$ for $y > 1$.
2. As $b \uparrow \infty$, $1/b \downarrow 0$ and again $1 + \ln(y)/\ln(b) \rightarrow 1$, so in the limit we have $F(y) = 0$ for $y < 0$ and $F(y) = 1$ for $y > 0$.

Since the probability density function is bounded on a bounded support interval, the Benford mantissa distribution can also be simulated via the rejection method.

Open the rejection method experiment and select the Benford mantissa distribution. Vary the base b and note again the shape and location of the probability density functions. For selected values of b , run the simulation 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

Distributions of the Digits

Assume now that the base is a positive integer $b \in \{2, 3, \dots\}$, which of course is the case in standard number systems. Suppose that the sequence of digits of our mantissa Y (in base b) is (N_1, N_2, \dots) , so that

$$Y = \sum_{k=1}^{\infty} \frac{N_k}{b^k} \quad (5.39.9)$$

Thus, our *leading digit* N_1 takes values in $\{1, 2, \dots, b-1\}$, while each of the other *significant digits* takes values in $\{0, 1, \dots, b-1\}$. Note that (N_1, N_2, \dots) is a stochastic process so at least we would like to know the *finite dimensional distributions*. That is, we would like to know the joint probability density function of the first k digits for every $k \in \mathbb{N}_+$. But let's start, appropriately enough, with the *first digit law*. The leading digit is the most important one, and fortunately also the easiest to analyze mathematically.

First Digit Law

N_1 has probability density function g_1 given by $g_1(n) = \log_b(1 + \frac{1}{n}) = \log_b(n+1) - \log_b(n)$ for $n \in \{1, 2, \dots, b-1\}$. The density function g_1 is decreasing and hence the mode is $n = 1$.

Proof

Note that $N_1 = n$ if and only if $\frac{n}{b} \leq Y < \frac{n+1}{b}$ for $n \in \{1, 2, \dots, b-1\}$. Hence using the PDF of Y above,

$$\mathbb{P}(N_1 = n) = \int_{n/b}^{(n+1)/b} \frac{1}{y \ln b} dy = \log_b\left(\frac{n+1}{b}\right) - \log_b\left(\frac{n}{b}\right) = \log_b(n+1) - \log_b(n) \quad (5.39.10)$$

Note that when $b = 2$, $N_1 = 1$ deterministically, which of course has to be the case. The first significant digit of a number in base 2 must be 1. Numerical values of g_1 for the standard (base 10) distribution are given in an [exercise below](#).

In the Special Distribution Simulator, select the Benford first digit distribution. Vary the base b with the input control and note the shape of the probability density function. For various values of b , run the simulation 1000 times and compare the empirical density function to the probability density function.

N_1 has distribution function G_1 given by $G_1(x) = \log_b(\lfloor x \rfloor + 1)$ for $x \in [1, b-1]$.

Proof

Using the PDF of N_1 above note that

$$G_1(n) = \sum_{k=1}^n [\log_b(k+1) - \log_b(k)] = \log_b(n+1), \quad n \in \{1, 2, \dots, b-1\} \quad (5.39.11)$$

More generally, $G_1(x) = G_1(\lfloor x \rfloor)$ for $x \in [1, b-1]$

N_1 has quantile function G_1^{-1} given by $G_1^{-1}(p) = \lceil b^p - 1 \rceil$ for $p \in (0, 1]$.

1. The first quartile is $\lceil b^{1/4} - 1 \rceil$.

2. The median is $\lceil b^{1/2} - 1 \rceil$.
3. The third quartile is $\lceil b^{3/4} - 1 \rceil$.

Proof

As usual, the formula for $G_1^{-1}(p)$ follows from the [CDF G](#), by solving $p = G(x)$ for x in terms of p .

Numerical values of the quantiles for the standard (base 10) distribution are given in an [exercise below](#).

Open the special distribution calculator and choose the Benford first digit distribution. Vary the base and note the shape and location of the distribution and probability density functions. For selected values of the base, compute the median and the first and third quartiles.

For the most part the moments of N_1 do not have simple expressions. However, we do have the following result for the mean.

$$\mathbb{E}(N_1) = (b-1) - \log_b[(b-1)!] .$$

Proof

From the [PDF of \$N_1\$](#) above and using standard properties of the logarithm,

$$\mathbb{E}(N_1) = \sum_{n=1}^{b-1} n \log_b \left(\frac{n+1}{n} \right) = \log_b \left[\prod_{n=1}^{b-1} \left(\frac{n+1}{n} \right)^n \right] \quad (5.39.12)$$

The product in the displayed equation simplifies to $b^{b-1}/(b-1)!$, and the base b logarithm of this expression is $(b-1) - \log_b[(b-1)!]$.

Numerical values of the mean and variance for the standard (base 10) distribution are given in an [exercise below](#).

Open the Special Distribution Simulator and select the Benford first digit distribution. Vary the base b with the input control and note the size and location of the mean \pm standard deviation bar. For various values of b , run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation..

Since the quantile function has a simple, closed form, the Benford first digit distribution can be simulated via the random quantile method.

Open the random quantile experiment and select the Benford first digit distribution. Vary the base b and note again the shape and location of the probability density function. For selected values of the base, run the experiment 1000 times and compare the empirical density function, mean, and standard deviation to their distributional counterparts.

Higher Digits

Now, to compute the joint probability density function of the first k significant digits, some additional notation will help.

If $n_1 \in \{1, 2, \dots, b-1\}$ and $n_j \in \{0, 1, \dots, b-1\}$ for $j \in \{2, 3, \dots, k\}$, let

$$[n_1 n_2 \dots n_k]_b = \sum_{j=1}^k n_j b^{k-j} \quad (5.39.13)$$

Of course, this is just the base b version of what we do in our standard base 10 system: we represent integers as strings of digits between 0 and 9 (except that the first digit cannot be 0). Here is a base 5 example:

$$[324]_5 = 3 \cdot 5^2 + 2 \cdot 5^1 + 4 \cdot 5^0 = 89 \quad (5.39.14)$$

The joint probability density function h_k of (N_1, N_2, \dots, N_k) is given by

$$h_k(n_1, n_2, \dots, n_k) = \log_b \left(1 + \frac{1}{[n_1 n_2 \dots n_k]_b} \right), \quad n_1 \in \{1, 2, \dots, b-1\}, (n_2, \dots, n_k) \in \{2, \dots, b-1\}^{k-1} \quad (5.39.15)$$

Proof

Note that $\{N_1 = n_1, N_2 = n_2, \dots, N_k = n_k\} = \{l \leq Y < u\}$. where

$$l = \frac{[n_1 n_2 \dots n_k]_b}{b^k}, \quad u = \frac{[n_1 n_2 \dots n_k]_b + 1}{b^k} \quad (5.39.16)$$

Hence using the [PDF of \$Y\$](#) and properties of logarithms,

$$h_k(n_1, n_2, \dots, n_k) = \int_l^u \frac{1}{y \ln(b)} dy = \log_b(u) - \log_b(l) = \log_b([n_1 n_2 \dots n_k]_b + 1) - \log_b([n_1 n_2 \dots, n_k]_b) \quad (5.39.17)$$

The probability density function of (N_1, N_2) in the standard (base 10) case is given in an [exercise below](#). Of course, the probability density function of a given digit can be obtained by summing the joint probability density over the unwanted digits in the usual way. However, except for the first digit, these functions do not reduce to simple expressions.

The probability density function g_2 of N_2 is given by

$$g_2(n) = \sum_{k=1}^{b-1} \log_b \left(1 + \frac{1}{[k n]_b} \right) = \sum_{k=1}^{b-1} \log_b \left(1 + \frac{1}{k b + n} \right), \quad n \in \{0, 1, \dots, b-1\} \quad (5.39.18)$$

The probability density function of N_2 in the standard (base 10) case is given in an [exercise below](#).

Theoretical Explanation

Aside from the empirical evidence noted by Newcomb and Benford (and many others since), why does Benford's law work? For a theoretical explanation, see the article [A Statistical Derivation of the Significant Digit Law](#) by Ted Hill.

Computational Exercises

In the following exercises, suppose that Y has the standard Benford mantissa distribution (the base 10 decimal case), and that (N_1, N_2, \dots) are the digits of Y .

Find each of the following for the mantissa Y

1. The density function f .
2. The mean and variance
3. The quartiles

Answer

1. $f(y) = \frac{1}{0.2303y}, \quad y \in \left[\frac{1}{10}, 1\right)$
2. $\mathbb{E}(Y) = 0.3909, \text{var}(Y) = 0.0622$
3. $q_1 = 0.1778, q_2 = 0.3162, q_3 = 0.5623$

For N_1 , find each of the following numerically

1. The probability density function
2. The mean and variance
3. The quartiles

Answer

1. n	$\mathbb{P}(N_1 = n)$
1	0.3010
2	0.1761
3	0.1249
4	0.0969
5	0.0792
6	0.0669
7	0.0580
8	0.0512
9	0.0458

2. $\mathbb{E}(N_1) = 3.4402, \text{var}(N_1) = 6.0567$
3. $q_1 = 1, q_2 = 3, q_3 = 5$

Explicitly compute the values of the joint probability density function of (N_1, N_2) .

Answer

$\mathbb{P}(N_1 = n_1, N_2 = n_2)$		2	3	4	5	6	7	8	9
$n_2 = 0$	0.0414	0.0212	0.0142	0.0107	0.0086	0.0072	0.0062	0.0054	0.0048
1	0.0378	0.0202	0.0138	0.0105	0.0084	0.0071	0.0061	0.0053	0.0047
2	0.0348	0.0193	0.0134	0.0102	0.0083	0.0069	0.0060	0.0053	0.0047
3	0.0322	0.0185	0.0130	0.0100	0.0081	0.0068	0.0059	0.0052	0.0046
4	0.0300	0.0177	0.0126	0.0098	0.0080	0.0067	0.0058	0.0051	0.0046
5	0.0280	0.0170	0.0122	0.0092	0.0078	0.0066	0.0058	0.0051	0.0045
6	0.0263	0.0164	0.0119	0.0093	0.0077	0.0065	0.0057	0.0050	0.0045
7	0.0248	0.0158	0.0116	0.0091	0.0076	0.0064	0.0056	0.0050	0.0045
8	0.0235	0.0152	0.0113	0.0090	0.0074	0.0063	0.0055	0.0049	0.0044
9	0.0223	0.0147	0.0110	0.0088	0.0073	0.0062	0.0055	0.0049	0.0044

For N_2 , find each of the following numerically

1. The probability density function
2. $\mathbb{E}(N_2)$
3. $\text{var}(N_2)$

Answer

1. n	$\mathbb{P}(N_2 = n)$
0	0.1197
1	0.1139
2	0.1088
3	0.1043
4	0.1003
5	0.0967
6	0.0934
7	0.0904
8	0.0876
9	0.0850

2. $\mathbb{E}(N_2) = 4.1847$
3. $\text{var}(N_2) = 0.8254$

Comparing the [result for \$N_1\$](#) and the [result for \$N_2\$](#) , note that the distribution of N_2 is flatter than the distribution of N_1 . In general, it turns out that distribution of N_k converges to the uniform distribution on $\{0, 1, \dots, b-1\}$ as $k \rightarrow \infty$. Interestingly, the digits are dependent.

N_1 and N_2 are dependent.

Proof

This result follows from the [joint PDF](#), the [marginal PDF of \$N_1\$](#) , and the [marginal PDF of \$N_2\$](#) above.

Find each of the following.

1. $\mathbb{P}(N_1 = 5, N_2 = 3, N_3 = 1)$
2. $\mathbb{P}(N_1 = 3, N_2 = 1, N_3 = 5)$
3. $\mathbb{P}(N_1 = 1, N_2 = 3, N_3 = 5)$

This page titled [5.39: Benford's Law](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.40: The Zeta Distribution

The *zeta distribution* is used to model the size or ranks of certain types of objects randomly chosen from certain types of populations. Typical examples include the frequency of occurrence of a word randomly chosen from a text, or the population rank of a city randomly chosen from a country. The zeta distribution is also known as the *Zipf distribution*, in honor of the American linguist George Zipf.

Basic Theory

The Zeta Function

The *Riemann zeta function* ζ , named after Bernhard Riemann, is defined as follows:

$$\zeta(a) = \sum_{n=1}^{\infty} \frac{1}{n^a}, \quad a \in (1, \infty) \quad (5.40.1)$$

You might recall from calculus that the series in the zeta function converges for $a > 1$ and diverges for $a \leq 1$.




Figure 5.40.1: Graph of ζ on the interval $(1, 10]$

The zeta function satisfies the following properties:

1. ζ is decreasing.
2. ζ is concave upward.
3. $\zeta(a) \downarrow 1$ as $a \uparrow \infty$
4. $\zeta(a) \uparrow \infty$ as $a \downarrow 1$

The zeta function is transcendental, and most of its values must be approximated. However, $\zeta(a)$ can be given explicitly for even integer values of a ; in particular, $\zeta(2) = \frac{\pi^2}{6}$ and $\zeta(4) = \frac{\pi^4}{90}$.

The Probability Density Function

The *zeta distribution* with *shape parameter* $a \in (1, \infty)$ is a discrete distribution on \mathbb{N}_+ with probability density function f given by.

$$f(n) = \frac{1}{\zeta(a)n^a}, \quad n \in \mathbb{N}_+ \quad (5.40.2)$$

1. f is decreasing with mode $n = 1$.
2. When smoothed, f is concave upward.

Proof

Clearly f is a valid PDF, since by definition, $\zeta(a)$ is the normalizing constant for the function $n \mapsto \frac{1}{n^a}$ on \mathbb{N}_+ . Part (a) is clear. For part (b), note that the function $x \mapsto x^{-a}$ on $[1, \infty)$ has a positive second derivative.

Open the special distribution simulator and select the zeta distribution. Vary the shape parameter and note the shape of the probability density function. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function and quantile function do not have simple closed forms, except in terms of other special functions.

Open the special distribution calculator and select the zeta distribution. Vary the parameter and note the shape of the distribution and probability density functions. For selected values of the parameter, compute the median and the first and third quartiles.

Moments

Suppose that N has the zeta distribution with shape parameter $a \in (1, \infty)$. The moments of X can be expressed easily in terms of the zeta function.

If $k \geq a - 1$, $\mathbb{E}(X) = \infty$. If $k < a - 1$,

$$\mathbb{E}(N^k) = \frac{\zeta(a-k)}{\zeta(a)} \quad (5.40.3)$$

Proof

Note that

$$\mathbb{E}(N^k) = \sum_{n=1}^{\infty} n^k \frac{1}{\zeta(a)n^a} = \frac{1}{\zeta(a)} \sum_{n=1}^{\infty} \frac{1}{n^{a-k}} \quad (5.40.4)$$

If $a - k \leq 1$, the last sum diverges to ∞ . If $a - k > 1$, the sum converges to $\zeta(a - k)$

The mean and variance of N are as follows:

1. If $a > 2$,

$$\mathbb{E}(N) = \frac{\zeta(a-1)}{\zeta(a)} \quad (5.40.5)$$

2. If $a > 3$,

$$\text{var}(N) = \frac{\zeta(a-2)}{\zeta(a)} - \left(\frac{\zeta(a-1)}{\zeta(a)} \right)^2 \quad (5.40.6)$$

Open the special distribution simulator and select the zeta distribution. Vary the parameter and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of N are as follows:

1. If $a > 4$,

$$\text{skew}(N) = \frac{\zeta(a-3)\zeta^2(a) - 3\zeta(a-1)\zeta(a-2)\zeta(a) + 2\zeta^3(a-1)}{[\zeta(a-2)\zeta(a) - \zeta^2(a-1)]^{3/2}} \quad (5.40.7)$$

2. If $a > 5$,

$$\text{kurt}(N) = \frac{\zeta(a-4)\zeta^3(a) - 4\zeta(a-1)\zeta(a-3)\zeta^2(a) + 6\zeta^2(a-1)\zeta(a-2)\zeta(a) - 3\zeta^4(a-1)}{[\zeta(a-2)\zeta(a) - \zeta^2(a-1)]^2} \quad (5.40.8)$$

Proof

These results follow from the [general moment result](#) above and standard computational formulas for skewness and kurtosis.

The probability generating function of N can be expressed in terms of the *polylogarithm function* Li that was introduced in the section on the exponential-logarithmic distribution. Recall that the polylogarithm of order $s \in \mathbb{R}$ is defined by

$$\text{Li}_s(x) = \sum_{k=1}^{\infty} \frac{x^k}{k^s}, \quad x \in (-1, 1) \quad (5.40.9)$$

N has probability generating function P given by

$$P(t) = \mathbb{E}(t^N) = \frac{\text{Li}_a(t)}{\zeta(a)}, \quad t \in (-1, 1) \quad (5.40.10)$$

Proof

Note that

$$\mathbb{E}(t^N) = \sum_{n=1}^{\infty} t^n \frac{1}{n^a \zeta(a)} = \frac{1}{\zeta(a)} \sum_{n=1}^{\infty} \frac{t^n}{n^a} \quad (5.40.11)$$

The last sum is $\text{Li}_a(t)$.

Related Distributions

In an algebraic sense, the zeta distribution is a discrete version of the Pareto distribution. Recall that if $a > 1$, the Pareto distribution with shape parameter $a - 1$ is a continuous distribution on $[1, \infty)$ with probability density function

$$f(x) = \frac{a-1}{x^a}, \quad x \in [1, \infty) \quad (5.40.12)$$

Naturally, the limits of the zeta distribution with respect to the shape parameter a are of interest.

The zeta distribution with shape parameter $a \in (1, \infty)$ converges to point mass at 1 as $a \rightarrow \infty$.

Proof

For the PDF f above, note that $f(1) = \zeta(a) \rightarrow 1$ as $a \rightarrow \infty$ and for $n \in \{2, 3, \dots\}$, $f(n) = 1/n^a \zeta(a) \rightarrow 0$ as $a \rightarrow \infty$

Finally, the zeta distribution is a member of the family of general exponential distributions.

Suppose that N has the zeta distribution with parameter a . Then the distribution is a one-parameter exponential family with natural parameter a and natural statistic $-\ln N$.

Proof

This follows from the definition of the general exponential distribution, since the zeta PDF can be written in the form

$$f(n) = \frac{1}{\zeta(a)} \exp(-a \ln n), \quad n \in \mathbb{N}_+ \quad (5.40.13)$$

Computational Exercises

Let N denote the frequency of occurrence of a word chosen at random from a certain text, and suppose that X has the zeta distribution with parameter $a = 2$. Find $\mathbb{P}(N > 4)$.

Answer

$$\mathbb{P}(N > 4) = 1 - \frac{49}{6\pi^2} \approx 0.1725$$

Suppose that N has the zeta distribution with parameter $a = 6$. Approximate each of the following:

1. $\mathbb{E}(N)$
2. $\text{var}(N)$
3. $\text{skew}(N)$
4. $\text{kurt}(N)$

Answer

1. $\mathbb{E}(N) \approx 1.109$
2. $\text{var}(N) \approx 0.025$
3. $\text{skew}(N) \approx 11.700$
4. $\text{kurt}(N) \approx 309.19$

This page titled [5.40: The Zeta Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

5.41: The Logarithmic Series Distribution

The *logarithmic series distribution*, as the name suggests, is based on the standard power series expansion of the natural logarithm function. It is also sometimes known more simply as the *logarithmic distribution*.

Basic Theory

Distribution Functions

The *logarithmic series distribution* with shape parameter $p \in (0, 1)$ is a discrete distribution on \mathbb{N}_+ with probability density function f given by

$$f(n) = \frac{1}{-\ln(1-p)} \frac{p^n}{n}, \quad n \in \mathbb{N}_+ \quad (5.41.1)$$

1. f is decreasing with mode $n = 1$.
2. When smoothed, f is concave upward.

Proof

Recall that the standard power series for $-\ln(1-p)$, obtained by integrating the geometric series $\sum_{n=0}^{\infty} p^n = 1/(1-p)$, is

$$-\ln(1-p) = \sum_{n=1}^{\infty} \frac{p^n}{n}, \quad p \in (0, 1) \quad (5.41.2)$$

For the properties, consider the function $x \mapsto p^x/x$ on $[1, \infty)$. The first derivative is

$$\frac{p^x [x \ln(p) - 1]}{x^2} \quad (5.41.3)$$

which is negative, and the second derivative is

$$\frac{p^x [x^2 \ln^2(p) - 2x \ln(p) + 2]}{x^3} \quad (5.41.4)$$

which is positive

Open the Special Distribution Simulator and select the logarithmic series distribution. Vary the parameter and note the shape of the probability density function. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function to the probability density function.

The distribution function and the quantile function do not have simple, closed forms in terms of the standard elementary functions.

Open the special distribution calculator and select the logarithmic series distribution. Vary the parameter and note the shape of the distribution and probability density functions. For selected values of the parameters, compute the median and the first and third quartiles.

Moments

Suppose again that random variable N has the logarithmic series distribution with shape parameter $p \in (0, 1)$. Recall that the permutation formula is $n^{(k)} = n(n-1) \cdots (n-k+1)$ for $n \in \mathbb{R}$ and $k \in \mathbb{N}$. The factorial moments of N are $\mathbb{E}(N^{(k)})$ for $k \in \mathbb{N}$.

The factorial moments of N are given by

$$\mathbb{E}(N^{(k)}) = \frac{(k-1)!}{-\ln(1-p)} \left(\frac{p}{1-p} \right)^k, \quad k \in \mathbb{N}_+ \quad (5.41.5)$$

Proof

Recall that a power series can be differentiated term by term within the open interval of convergence. Hence

$$\mathbb{E}(N^{(k)}) = \sum_{n=1}^{\infty} n^{(k)} \frac{1}{-\ln(1-p)} \frac{p^n}{n} = \frac{p^k}{-\ln(1-p)} \sum_{n=k}^{\infty} n^{(k)} \frac{p^{n-k}}{n} \quad (5.41.6)$$

$$= \frac{p^k}{-\ln(1-p)} \sum_{n=k}^{\infty} \frac{d^k}{dp^k} \frac{p^n}{n} = \frac{p^k}{-\ln(1-p)} \frac{d^k}{dp^k} \sum_{n=1}^{\infty} \frac{p^n}{n} \quad (5.41.7)$$

$$= \frac{p^k}{-\ln(1-p)} \frac{d^k}{dp^k} [-\ln(1-p)] = \frac{p^k}{-\ln(1-p)} (k-1)!(1-p)^{-k} \quad (5.41.8)$$

The mean and variance of N are

$$1. \quad \mathbb{E}(N) = \frac{1}{-\ln(1-p)} \frac{p}{1-p} \quad (5.41.9)$$

$$2. \quad \text{var}(N) = \frac{1}{-\ln(1-p)} \frac{p}{(1-p)^2} \left[1 - \frac{p}{-\ln(1-p)} \right] \quad (5.41.10)$$

Proof

These results follow easily from the [factorial moments](#). For part (b), note first that

$$\mathbb{E}(N^2) = \mathbb{E}[N(N-1)] + \mathbb{E}(N) = \frac{1}{-\ln(1-p)} \frac{p}{(1-p)^2} \quad (5.41.11)$$

The result then follows from the usual computational formula $\text{var}(N) = \mathbb{E}(N^2) - [\mathbb{E}(N)]^2$.

Open the special distribution simulator and select the logarithmic series distribution. Vary the parameter and note the shape of the mean \pm standard deviation bar. For selected values of the parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The probability generating function P of N is given by

$$P(t) = \mathbb{E}(t^N) = \frac{\ln(1-pt)}{\ln(1-p)}, \quad |t| < \frac{1}{p} \quad (5.41.12)$$

Proof

$$P(t) = \sum_{n=1}^{\infty} t^n \frac{1}{-\ln(1-p)} \frac{p^n}{n} = \frac{1}{-\ln(1-p)} \sum_{n=1}^{\infty} \frac{(pt)^n}{n} = \frac{-\ln(1-pt)}{-\ln(1-p)} \quad (5.41.13)$$

The [factorial moments](#) above can also be obtained from the probability generating function, since $P^{(k)}(1) = \mathbb{E}(N^{(k)})$ for $k \in \mathbb{N}_+$.

Related Distributions

Naturally, the limits of the logarithmic series distribution with respect to the parameter p are of interest.

The logarithmic series distribution with shape parameter $p \in (0, 1)$ converges to point mass at 1 as $p \downarrow 0$.

Proof

An application of L'Hopitals rule to the [PGF \$P\$](#) above shows that $\lim_{p \downarrow 0} P(t) = t$, which is the PGF of point mass at 1.

The logarithmic series distribution is a power series distribution associated with the function $g(p) = -\ln(1-p)$ for $p \in [0, 1)$.

Proof

This follows from the definition of a power series distribution, since as noted in the [PDF proof](#),

$$\sum_{n=1}^{\infty} \frac{p^n}{n} = -\ln(1-p), \quad p \in [0, 1) \quad (5.41.14)$$

The [moment results](#) above actually follow from general results for power series distributions. The compound Poisson distribution based on the logarithmic series distribution gives a negative binomial distribution.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables each with the logarithmic series distribution with parameter $p \in (0, 1)$. Suppose also that N is independent of \mathbf{X} and has the Poisson distribution with rate parameter $r \in (0, \infty)$. Then $Y = \sum_{i=1}^N X_i$ has the negative binomial distribution on \mathbb{N} with parameters $1-p$ and $-r/\ln(1-p)$

Proof

The PGF of Y is $Q \circ P$, where P is the [PGF](#) of the logarithmic series distribution, and where Q is the PGF of the Poisson distribution so that $Q(s) = e^{r(s-1)}$ for $s \in \mathbb{R}$. Thus we have

$$(Q \circ P)(t) = \exp\left(r \left[\frac{\ln(1-pt)}{\ln(1-p)} - 1 \right]\right), \quad |t| < \frac{1}{p} \quad (5.41.15)$$

With a little algebra, this can be written in the form

$$(Q \circ P)(t) = \left(\frac{1-p}{1-pt} \right)^{-r/\ln(1-p)}, \quad |t| < \frac{1}{p} \quad (5.41.16)$$

which is the PGF of the negative binomial distribution with parameters $1-p$ and $-r/\ln(1-p)$.

This page titled [5.41: The Logarithmic Series Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

6: Random Samples

Point estimation refers to the process of estimating a parameter from a probability distribution, based on observed data from the distribution. It is one of the core topics in mathematical statistics. In this chapter, we will explore the most common methods of point estimation: the method of moments, the method of maximum likelihood, and Bayes' estimators. We also study important properties of estimators, including sufficiency and completeness, and the basic question of whether an estimator is the best possible one.

[6.1: Introduction](#)

[6.2: The Sample Mean](#)

[6.3: The Law of Large Numbers](#)

[6.4: The Central Limit Theorem](#)

[6.5: The Sample Variance](#)

[6.6: Order Statistics](#)

[6.7: Sample Correlation and Regression](#)

[6.8: Special Properties of Normal Samples](#)

This page titled [6: Random Samples](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.1: Introduction

The Basic Statistical Model

In the basic statistical model, we have a *population* of objects of interest. The objects could be persons, families, computer chips, acres of corn. In addition, we have various measurements or *variables* defined on the objects. We select a sample from the population and record the variables of interest for each object in the sample. Here are a few examples based on the data sets in this project:

- In the M&M data, the objects are bags of M&Ms of a specified size. The variables recorded for a bag of M&Ms are net weight and the counts for red, green, blue, orange, yellow, and brown candies.
- In the cicada data, the objects are cicadas from the middle Tennessee area. The variables recorded for a cicada are body weight, wing length, wing width, body length, gender, and species.
- In Fisher's iris data, the objects are irises. The variables recorded for an iris are petal length, petal width, sepal length, sepal width, and type.
- In the Polio data set, the objects are children. Although many variables were probably recorded for a child, the two crucial variables, both binary, were whether or not the child was vaccinated, and whether or not the child contracted Polio within a certain time period.
- In the Challenger data sets, the objects are Space Shuttle launches. The variables recorded are temperature at the time of launch and various measures of O-ring erosion of the solid rocket boosters.
- In Michelson's data set, the objects are beams of light and the variable recorded is speed.
- In Pearson's data set, the objects are father-son pairs. The variables are the height of the father and the height of the son.
- In Snow's data set, the objects are persons who died of cholera. The variables record the address of the person.
- In one of the SAT data sets, the objects are states and the variables are participation rate, average SAT Math score and average SAT Verbal score.

Thus, the observed outcome of a statistical experiment (the data) has the form $\mathbf{x} = (x_1, x_2, \dots, x_n)$ where x_i is the vector of measurements for the i th object chosen from the population. The set S of possible values of \mathbf{x} (before the experiment is conducted) is called the *sample space*. It is literally the space of samples. Thus, although the outcome of a statistical experiment can have quite a complicated structure (a vector of vectors), the hallmark of mathematical abstraction is the ability to *gray out* the features that are not relevant at any particular time, to treat a complex structure as a single object. This we do with the outcome \mathbf{x} of the experiment.

The techniques of statistics have been enormously successful; these techniques are widely used in just about every subject that deals with quantification—the natural sciences, the social sciences, law, and medicine. On the other hand, statistics has a legalistic quality and a great deal of terminology and jargon that can make the subject a bit intimidating at first. In the rest of this section, we begin discussing some of this terminology.

The Empirical Distribution

Suppose again that the data have the form $\mathbf{x} = (x_1, x_2, \dots, x_n)$ where x_i is the vector of measurements for the i th object chosen. The *empirical distribution* associated with \mathbf{x} is the probability distribution that places probability $1/n$ at each x_i . Thus, if the values are distinct, the empirical distribution is the discrete uniform distribution on $\{x_1, x_2, \dots, x_n\}$. More generally, if x occurs k times in the data, then the empirical distribution assigns probability k/n to x . Thus, every finite data set defines a probability distribution.

Statistics

Technically, a *statistic* $w = w(\mathbf{x})$ is an observable function of the outcome \mathbf{x} of the experiment. That is, a statistic is a computable function defined on the sample space S . The term *observable* means that the function should not contain any unknown quantities, because we need to be able to compute the value w of the statistic from the observed data \mathbf{x} . As with the data \mathbf{x} , a statistic w may have a complicated structure; typically, w is vector valued. Indeed, the outcome \mathbf{x} of the experiment is itself a statistic; all other statistics are derived from \mathbf{x} .

Statistics u and v are *equivalent* if there exists a one-to-one function r from the range of u onto the range of v such that $v = r(u)$. Equivalent statistics give equivalent information about \mathbf{x} .

Statistics u and v are equivalent if and only if the following condition holds: for any $\mathbf{x} \in S$ and $\mathbf{y} \in S$, $u(\mathbf{x}) = u(\mathbf{y})$ if and only if $v(\mathbf{x}) = v(\mathbf{y})$.

Equivalence really is an equivalence relation on the collection of statistics for a given statistical experiment. That is, if u , v , and w are arbitrary statistics then

1. u is equivalent to u (the *reflexive property*).
2. If u is equivalent to v then v is equivalent to u (the *symmetric property*).
3. If u is equivalent to v and v is equivalent to w then u is equivalent to w (the *transitive property*).

Descriptive and Inferential Statistics

There are two broad branches of statistics. The term *descriptive statistics* refers to methods for summarizing and displaying the observed data \mathbf{x} . As the name suggests, the methods of descriptive statistics usually involve computing various statistics (in the technical sense) that give useful information about the data: measures of center and spread, measures of association, and so forth. In the context of descriptive statistics, the term *parameter* refers to a characteristic of the entire population.

The deeper and more useful branch of statistics is known as *inferential statistics*. Our point of view in this branch is that the statistical experiment (before it is conducted) is a random experiment with a probability measure \mathbb{P} on an underlying sample space. Thus, the outcome \mathbf{x} of the experiment is an observed value of a random variable \mathbf{X} defined on this probability space, with the distribution of \mathbf{X} not completely known to us. Our goal is to draw inferences about the distribution of \mathbf{X} from the observed value \mathbf{x} . Thus, in a sense, inferential statistics is the dual of probability. In probability, we try to *predict* the value of \mathbf{X} *assuming* complete knowledge of the distribution. In statistics, by contrast, we *observe* the value of \mathbf{x} of the random variable \mathbf{X} and try to *infer* information about the underlying distribution of \mathbf{X} . In inferential statistics, a *statistic* (a function of \mathbf{X}) is itself a random variable with a distribution of its own. On the other hand, the term *parameter* refers to a characteristic of the distribution of \mathbf{X} . Often the inferential problem is to use various statistics to *estimate* or *test hypotheses* about a parameter. Another way to think of inferential statistics is that we are trying to infer from the *empirical* distribution associated with the observed data \mathbf{x} to the *true* distribution associated with \mathbf{X} .

There are two basic types of random experiments in the general area of inferential statistics. A *designed experiment*, as the name suggests, is carefully designed to study a particular inferential question. The experimenter has considerable control over how the objects are selected, what variables are to be recorded for these objects, and the values of certain of the variables. In an *observational study*, by contrast, the researcher has little control over these factors. Often the researcher is simply given the data set and asked to make sense out of it. For example, the Polio field trials were designed experiments to study the effectiveness of the Salk vaccine. The researchers had considerable control over how the children were selected, and how the children were assigned to the treatment and control groups. By contrast, the Challenger data sets used to explore the relationship between temperature and O-ring erosion are observational studies. Of course, just because an experiment is designed does not mean that it is *well* designed.

Difficulties

A number of difficulties can arise when trying to explore an inferential question. Often, problems arise because of *confounding variables*, which are variables that (as the name suggests) interfere with our understanding of the inferential question. In the first Polio field trial design, for example, *age* and *parental consent* are two confounding variables that interfere with the determination of the effectiveness of the vaccine. The entire point of the Berkeley admissions data, to give another example, is to illustrate how a confounding variable (*department*) can create a spurious correlation between two other variables (*gender* and *admissions status*). When we correct for the interference caused by a confounding variable, we say that we have *controlled* for the variable.

Problems also frequently arise because of *measurement errors*. Some variables are inherently difficult to measure, and systematic bias in the measurements can interfere with our understanding of the inferential question. The first Polio field trial design again provides a good example. Knowledge of the vaccination status of the children led to systematic bias by doctors attempting to diagnose polio in these children. Measurement errors are sometimes caused by hidden confounding variables.

Confounding variables and measurement errors abound in political polling, where the inferential question is who will win an election. How do confounding variables such as race, income, age, and gender (to name just a few) influence how a person will vote? How do we know that a person will vote for whom she says she will, or if she will vote at all (measurement errors)? The Literary Digest poll in the 1936 presidential election and the professional polls in the 1948 presidential election illustrate these problems.

Confounding variables, measurement errors and other causes often lead to *selection bias*, which means that the sample does not represent the population with respect to the inferential question at hand. Often randomization is used to overcome the effects of confounding variables and measurement errors.

Random Samples

The most common and important special case of the inferential statistical model occurs when the observation variable

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (6.1.1)$$

is a sequence of independent and identically distributed random variables. Again, in the standard sampling model, X_i is itself a vector of measurements for the i th object in the sample, and thus, we think of (X_1, X_2, \dots, X_n) as independent copies of an underlying measurement vector X . In this case, (X_1, X_2, \dots, X_n) is said to be a *random sample* of size n from the distribution of X .

Variables

The mathematical operations that make sense for variable in a statistical experiment depend on the *type* and *level of measurement* of the variable.

Type

Recall that a real variable x is *continuous* if the possible values form an interval of real numbers. For example, the weight variable in the M&M data set, and the length and width variables in Fisher's iris data are continuous. In contrast, a *discrete variable* is one whose set of possible values forms a discrete set. For example, the counting variables in the M&M data set, the type variable in Fisher's iris data, and the denomination and suit variables in the card experiment are discrete. Continuous variables represent quantities that can, in theory, be measured to any degree of accuracy. In practice, of course, measuring devices have limited accuracy so data collected from a continuous variable are necessarily discrete. That is, there is only a finite (but perhaps very large) set of possible values that can actually be measured. So, the distinction between a discrete and continuous variable is based on what is theoretically possible, not what is actually measured. Some additional examples may help:

- A person's *age* is usually given in years. However, one can imagine age being given in months, or weeks, or even (if the time of birth is known to a sufficient accuracy) in seconds. Age, whether of devices or persons, is usually considered to be a continuous variable.
- The *price* of an item is usually given (in the US) in dollars and cents, and of course, the smallest monetary object in circulation is the penny (\$0.01). However, taxes are sometimes given in mills (\$0.001), and one can imagine smaller divisions of a dollar, even if there are no coins to represent these divisions. Measures of wealth are usually thought of as continuous variables.
- On the other hand, the number of persons in a car at the time of an accident is a fundamentally discrete variable.

Levels of Measurement

A real variable x is also distinguished by its *level of measurement*.

Qualitative variables simply encode *types* or *names*, and thus few mathematical operations make sense, even if numbers are used for the encoding. Such variables have the *nominal level of measurement*. For example, the type variable in Fisher's iris data is qualitative. Gender, a common variable in many studies of persons and animals, is also qualitative. Qualitative variables are almost always discrete; it's hard to imagine a continuous infinity of names.

A variable for which only *order* is meaningful is said to have the *ordinal level of measurement*; differences are not meaningful even if numbers are used for the encoding. For example, in many card games, the suits are ranked, so the suit variable has the ordinal level of measurement. For another example, consider the standard 5-point scale (terrible, bad, average, good, excellent) used to rank teachers, movies, restaurants etc.

A quantitative variable for which *difference*, but not *ratios* are meaningful is said to have the *interval level of measurement*. Equivalently, a variable at this level has a relative, rather than absolute, zero value. Typical examples are temperature (in Fahrenheit or Celsius) or time (clock or calendar).

Finally, a quantitative variable for which ratios are meaningful is said to have the *ratio level of measurement*. A variable at this level has an absolute zero value. The count and weight variables in the M&M data set, and the length and width variables in Fisher's iris data are examples.

Subsamples

In the basic statistical model, subsamples corresponding to some of the variables can be constructed by *filtering* with respect to other variables. This is particularly common when the filtering variables are qualitative. Consider the cicada data for example. We might be interested in the quantitative variables body weight, body length, wing width, and wing length *by species*, that is, separately for species 0, 1, and 2. Or, we might be interested in these quantitative variables *by gender*, that is separately for males and females.

Exercises

Study Michelson's experiment to measure the velocity of light.

1. Is this a designed experiment or an observational study?
2. Classify the velocity of light variable in terms of type and level of measurement.
3. Discuss possible confounding variables and problems with measurement errors.

Answer

1. Designed experiment
2. Continuous, interval. The level of measurement is only interval because the recorded variable is the speed of light in km/hr minus 299 000(to make the numbers simpler). The actual speed in km/hr is a continuous, ratio variable.

Study Cavendish's experiment to measure the density of the earth.

1. Is this a designed experiment or an observational study?
2. Classify the density of earth variable in terms of type and level of measurement.
3. Discuss possible confounding variables and problems with measurement errors.

Answer

1. Designed experiment
2. Continuous, ratio.

Study Short's experiment to measure the parallax of the sun.

1. Is this a designed experiment or an observational study?
2. Classify the parallax of the sun variable in terms of type and level of measurement.
3. Discuss possible confounding variables and problems with measurement errors.

Answer

1. Observational study
2. Continuous, ratio.

In the M&M data, classify each variable in terms of type and level of measurement.

Answer

Each color count variable: discrete, ratio; Net weight: continuous, ratio

In the Cicada data, classify each variable in terms of type and level of measurement.

Answer

Body weight, wing length, wing width, body length: continuous, ratio. Gender, type: discrete, nominal

In Fisher's iris data, classify each variable in terms of type and level of measurement.

Answer

Petal width, petal length, sepal width, sepal length: continuous, ratio. Type: discrete, nominal

Study the Challenger experiment to explore the relationship between temperature and O-ring erosion.

1. Is this a designed experiment or an observational study?

2. Classify each variable in terms of type and level of measurement.
3. Discuss possible confounding variables and problems with measurement errors.

Answer

1. Observational study
2. Temperature: continuous, interval; Erosion: continuous, ratio; Damage index: discrete, ordinal

In the Vietnam draft data, classify each variable in terms of type and level of measurement.

Answer

Birth month: discrete, interval; Birth day: discrete, interval

In the two SAT data sets, classify each variable in terms of type and level of measurement.

Answer

SAT math and verbal scores: probably continuous, ratio; State: discrete, nominal; Year: discrete, interval

Study the Literary Digest experiment to to predict the outcome of the 1936 presidential election.

1. Is this a designed experiment or an observational study?
2. Classify each variable in terms of type and level of measurement.
3. Discuss possible confounding variables and problems with measurement errors.

Answer

1. designed experiment, although poorly designed
2. State: discrete, nominal; Electoral votes: discrete, ratio; Landon count: discrete, ratio; Roosevelt count: discrete, ratio

Study the 1948 polls to predict the outcome of the presidential election between Truman and Dewey. Are these designed experiments or an observational studies?

Answer

Designed experiments, but poorly designed

Study Pearson's experiment to explore the relationship between heights of fathers and heights of sons.

1. Is this a designed experiment or an observational study?
2. Classify each variable in terms of type and level of measurement.
3. Discuss possible confounding variables.

Answer

1. Observational study
2. height of the father: continuous ratio; height of the son: continuous ratio

Study the Polio field trials.

1. Are these designed experiments or observational studies?
2. Identify the essential variables and classify each in terms of type and level of measurement.
3. Discuss possible confounding variables and problems with measurement errors.

Answer

1. designed experiments
2. vaccination status: discrete, nominal; Polio status: discrete, nominal

Identify the parameters in each of the following:

1. Buffon's Coin Experiment
2. Buffon's Needle Experiment
3. the Bernoulli trials model

4. the Poisson model

Answer

1. radius of the coin
2. length of the needle
3. probability of success
4. rate of arrivals

Note the parameters for each of the following families of special distributions:

1. the normal distribution
2. the gamma distribution
3. the beta distribution
4. the Pareto distribution
5. the Weibull distribution

Answer

1. mean μ and standard deviation σ
2. shape parameter k and scale parameter b
3. left parameter a and right parameter b
4. shape parameter a and scale parameter b
5. shape parameter k and scale parameter b

During World War II, the Allies recorded the serial numbers of captured German tanks. Classify the underlying *serial number* variable by type and level of measurement.

Answer

discrete, ordinal.

For a discussion of how the serial numbers were used to estimate the total number of tanks, see the section on [Order Statistics](#) in the chapter on Finite Sampling Models.

This page titled [6.1: Introduction](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.2: The Sample Mean

Basic Theory

Recall the basic model of statistics: we have a population of objects of interest, and we have various measurements (variables) that we make on the objects. We select objects from the population and record the variables for the objects in the sample; these become our data. Our first discussion is from a purely descriptive point of view. That is, we do not assume that the data are generated by an underlying probability distribution. However, recall that the data themselves define a probability distribution.

Definition and Basic Properties

Suppose that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a real-valued variable. The *sample mean* is simply the arithmetic average of the sample values:

$$m = \frac{1}{n} \sum_{i=1}^n x_i \quad (6.2.1)$$

If we want to emphasize the dependence of the mean on the data, we write $m(\mathbf{x})$ instead of just m . Note that m has the same physical units as the underlying variable. For example, if we have a sample of weights of cicadas, in grams, then m is in grams also. The sample mean is frequently used as a *measure of center* of the data. Indeed, if each x_i is the location of a *point mass*, then m is the *center of mass* as defined in physics. In fact, a simple graphical display of the data is the *dotplot*: on a number line, a dot is placed at x_i for each i . If values are repeated, the dots are stacked vertically. The sample mean m is the balance point of the dotplot. The image below shows a dot plot with the mean as the balance point.




Figure 6.2.1: A dotplot

The standard notation for the sample mean corresponding to the data \mathbf{x} is \bar{x} . We break with tradition and do not use the bar notation in this text, because it's clunky and because it's inconsistent with the notation for other statistics such as the sample variance, sample standard deviation, and sample covariance. However, you should be aware of the standard notation, since you will undoubtedly see it in other sources.

The following exercises establish a few simple properties of the sample mean. Suppose that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and $\mathbf{y} = (y_1, y_2, \dots, y_n)$ are samples of size n from real-valued population variables and that c is a constant. In vector notation, recall that $\mathbf{x} + \mathbf{y} = (x_1 + y_1, x_2 + y_2, \dots, x_n + y_n)$ and $c\mathbf{x} = (cx_1, cx_2, \dots, cx_n)$.

Computing the sample mean is a *linear operation*.

1. $m(\mathbf{x} + \mathbf{y}) = m(\mathbf{x}) + m(\mathbf{y})$
2. $m(c\mathbf{x}) = c m(\mathbf{x})$

Proof

1.
$$m(\mathbf{x} + \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n (x_i + y_i) = \frac{1}{n} \sum_{i=1}^n x_i + \frac{1}{n} \sum_{i=1}^n y_i = m(\mathbf{x}) + m(\mathbf{y}) \quad (6.2.2)$$

2.
$$m(c\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n cx_i = c \frac{1}{n} \sum_{i=1}^n x_i = c m(\mathbf{x}) \quad (6.2.3)$$

The sample mean preserves order.

1. If $x_i \geq 0$ for each i then $m(\mathbf{x}) \geq 0$.
2. If $x_i \geq 0$ for each i and $x_j > 0$ for some j then $m(\mathbf{x}) > 0$

3. If $x_i \leq y_i$ for each i then $m(\mathbf{x}) \leq m(\mathbf{y})$
4. If $x_i \leq y_i$ for each i and $x_j < y_j$ for some j then $m(\mathbf{x}) < m(\mathbf{y})$

Proof

Parts (a) and (b) are obvious from the definition. Part (c) follows from part (a) and the linearity of expected value. Specifically, if $\mathbf{x} \leq \mathbf{y}$ (in the product ordering), then $\mathbf{y} - \mathbf{x} \geq \mathbf{0}$. Hence by (a), $m(\mathbf{y} - \mathbf{x}) \geq 0$. But $m(\mathbf{y} - \mathbf{x}) = m(\mathbf{y}) - m(\mathbf{x})$. Hence $m(\mathbf{y}) \geq m(\mathbf{x})$. Similarly, (d) follows from (b) and the linearity of expected value.

Trivially, the mean of a constant sample is simply the constant. .

If $\mathbf{c} = (c, c, \dots, c)$ is a constant sample then $m(\mathbf{c}) = c$.

Proof

Note that

$$m(\mathbf{c}) = \frac{1}{n} \sum_{i=1}^n c_i = \frac{nc}{n} = c \quad (6.2.4)$$

As a special case of these results, suppose that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n corresponding to a real variable x , and that a and b are constants. Then the sample corresponding to the variable $y = a + bx$, in our vector notation, is $\mathbf{a} + b\mathbf{x}$. The sample means are related in precisely the same way, that is, $m(\mathbf{a} + b\mathbf{x}) = a + bm(\mathbf{x})$. Linear transformations of this type, when $b > 0$, arise frequently when physical units are changed. In this case, the transformation is often called a *location-scale* transformation; a is the location parameter and b is the scale parameter. For example, if x is the length of an object in inches, then $y = 2.54x$ is the length of the object in centimeters. If x is the temperature of an object in degrees Fahrenheit, then $y = \frac{5}{9}(x - 32)$ is the temperature of the object in degree Celsius.

Sample means are ubiquitous in statistics. In the next few paragraphs we will consider a number of special statistics that are based on sample means.

The Empirical Distribution

Suppose now that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a general variable taking values in a set S . For $A \subseteq S$, the *frequency* of A corresponding to \mathbf{x} is the number of data values that are in A :

$$n(A) = \#\{i \in \{1, 2, \dots, n\} : x_i \in A\} = \sum_{i=1}^n \mathbf{1}(x_i \in A) \quad (6.2.5)$$

The *relative frequency* of A corresponding to \mathbf{x} is the proportion of data values that are in A :

$$p(A) = \frac{n(A)}{n} = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i \in A) \quad (6.2.6)$$

Note that for fixed A , $p(A)$ is itself a sample mean, corresponding to the data $\{\mathbf{1}(x_i \in A) : i \in \{1, 2, \dots, n\}\}$. This fact bears repeating: *every sample proportion is a sample mean*, corresponding to an indicator variable. In the picture below, the red dots represent the data, so $p(A) = 4/15$.




Figure 6.2.2: The empirical probability of A

p is a probability measure on S .

1. $p(A) \geq 0$ for every $A \subseteq S$
2. $p(S) = 1$

3. If $\{A_j : j \in J\}$ is a countable collection of pairwise disjoint subsets of S then $p\left(\bigcup_{j \in J} A_j\right) = \sum_{j \in J} p(A_j)$

Proof

Parts (a) and (b) are obvious. For part (c) note that since the sets are disjoint,

$$p\left(\bigcup_{i \in I} A_i\right) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\left(x_i \in \bigcup_{j \in J} A_j\right) = \frac{1}{n} \sum_{i=1}^n \sum_{j \in J} \mathbf{1}(x_i \in A_j) \quad (6.2.7)$$

$$= \sum_{j \in J} \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i \in A_j) = \sum_{j \in J} p(A_j) \quad (6.2.8)$$

This probability measure is known as the *empirical probability distribution* associated with the data set \mathbf{x} . It is a discrete distribution that places probability $\frac{1}{n}$ at each point x_i . In fact this observation supplies a simpler proof of previous theorem. Thus, if the data values are distinct, the empirical distribution is the discrete uniform distribution on $\{x_1, x_2, \dots, x_n\}$. More generally, if $x \in S$ occurs k times in the data then the empirical distribution assigns probability k/n to x .

If the underlying variable is real-valued, then clearly the sample mean is simply the mean of the empirical distribution. It follows that the sample mean satisfies *all* properties of expected value, not just the [linear properties](#) and [increasing properties](#) given above. These properties are just the most important ones, and so were repeated for emphasis.

Empirical Density

Suppose now that the population variable x takes values in a set $S \subseteq \mathbb{R}^d$ for some $d \in \mathbb{N}_+$. Recall that the *standard measure* on \mathbb{R}^d is given by

$$\lambda_d(A) = \int_A 1 \, dx, \quad A \subseteq \mathbb{R}^d \quad (6.2.9)$$

In particular $\lambda_1(A)$ is the length of A , for $A \subseteq \mathbb{R}$; $\lambda_2(A)$ is the area of A , for $A \subseteq \mathbb{R}^2$; and $\lambda_3(A)$ is the volume of A , for $A \subseteq \mathbb{R}^3$. Suppose that x is a continuous variable in the sense that $\lambda_d(S) > 0$. Typically, S is an interval if $d = 1$ and a Cartesian product of intervals if $d > 1$. Now for $A \subseteq S$ with $\lambda_d(A) > 0$, the *empirical density* of A corresponding to \mathbf{x} is

$$D(A) = \frac{p(A)}{\lambda_d(A)} = \frac{1}{n \lambda_d(A)} \sum_{i=1}^n \mathbf{1}(x_i \in A) \quad (6.2.10)$$

Thus, the empirical density of A is the proportion of data values in A , divided by the size of A . In the picture below (corresponding to $d = 2$), if A has area 5, say, then $D(A) = 4/75$.




Figure 6.2.3: The empirical density of A

The Empirical Distribution Function

Suppose again that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a real-valued variable. For $x \in \mathbb{R}$, let $F(x)$ denote the [relative frequency](#) (empirical probability) of $(-\infty, x]$ corresponding to the data set \mathbf{x} . Thus, for each $x \in \mathbb{R}$, $F(x)$ is the sample mean of the data $\{\mathbf{1}(x_i \leq x) : i \in \{1, 2, \dots, n\}\}$:

$$F(x) = p((-\infty, x]) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i \leq x) \quad (6.2.11)$$

F is a distribution function.

1. F increases from 0 to 1.
2. F is a step function with jumps at the distinct sample values $\{x_1, x_2, \dots, x_n\}$.

Proof

Suppose that (y_1, y_2, \dots, y_k) are the distinct values of the data, ordered from smallest to largest, and that y_j occurs n_j times in the data. Then $F(x) = 0$ for $x < y_1$, $F(x) = n_1/n$ for $y_1 \leq x < y_2$, $F(x) = (n_1 + n_2)/n$ for $y_2 \leq x < y_3$, and so forth.

Appropriately enough, F is called the *empirical distribution function* associated with \mathbf{x} and is simply the distribution function of the *empirical distribution* corresponding to \mathbf{x} . If we know the sample size n and the empirical distribution function F , we can recover the data, except for the order of the observations. The distinct values of the data are the places where F jumps, and the number of data values at such a point is the size of the jump, times the sample size n .

The Empirical Discrete Density Function

Suppose now that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a discrete variable that takes values in a countable set S . For $x \in S$, let $f(x)$ be the *relative frequency* (empirical probability) of x corresponding to the data set \mathbf{x} . Thus, for each $x \in S$, $f(x)$ is the sample mean of the data $\{\mathbf{1}(x_i = x) : i \in \{1, 2, \dots, n\}\}$:

$$f(x) = p(\{x\}) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i = x) \quad (6.2.12)$$

In the picture below, the dots are the possible values of the underlying variable. The red dots represent the data, and the numbers indicate repeated values. The blue dots are possible values of the variable that did not happen to occur in the data. So, the sample size is 12, and for the value x that occurs 3 times, we have $f(x) = 3/12$.




Figure 6.2.4: The discrete probability density function

f is a discrete probability density function:

1. $f(x) \geq 0$ for $x \in S$
2. $\sum_{x \in S} f(x) = 1$

Proof

Part (a) is obvious. For part (b), note that

$$\sum_{x \in S} f(x) = \sum_{x \in S} p(\{x\}) = 1 \quad (6.2.13)$$

Appropriately enough, f is called the *empirical probability density function* or the *relative frequency function* associated with \mathbf{x} , and is simply the probability density function of the *empirical distribution* corresponding to \mathbf{x} . If we know the empirical PDF f and the sample size n , then we can recover the data set, except for the order of the observations.

If the underlying population variable is real-valued, then the sample mean is the expected value computed relative to the empirical density function. That is,

$$\frac{1}{n} \sum_{i=1}^n x_i = \sum_{x \in S} x f(x) \quad (6.2.14)$$

Proof

Note that

$$\sum_{x \in S} x f(x) = \sum_{x \in S} x \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i = x) = \frac{1}{n} \sum_{i=1}^n \sum_{x \in S} x \mathbf{1}(x_i = x) = \frac{1}{n} \sum_{i=1}^n x_i \quad (6.2.15)$$

As we noted earlier, if the population variable is real-valued then the sample mean is the mean of the empirical distribution.

The Empirical Continuous Density Function

Suppose now that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a continuous variable that takes values in a set $S \subseteq \mathbb{R}^d$. Let $\mathcal{A} = \{A_j : j \in J\}$ be a partition of S into a countable number of subsets, each of positive, finite measure. Recall that the word *partition* means that the subsets are pairwise disjoint and their union is S . Let f be the function on S defined by the rule that $f(x)$ is the **empirical density** of A_j , corresponding to the data set \mathbf{x} , for each $x \in A_j$. Thus, f is constant on each of the partition sets:

$$f(x) = D(A_j) = \frac{p(A_j)}{\lambda_d(A_j)} = \frac{1}{n\lambda_d(A_j)} \sum_{i=1}^n \mathbf{1}(x_i \in A_j), \quad x \in A_j \quad (6.2.16)$$

f is a continuous probability density function.

1. $f(x) \geq 0$ for $x \in S$
2. $\int_S f(x) dx = 1$

Proof

Part (a) is obvious. For part (b) note that since f is constant on A_j for each $j \in J$ we have

$$\int_S f(x) dx = \sum_{j \in J} \int_{A_j} f(x) dx = \sum_{j \in J} \lambda_d(A_j) \frac{p(A_j)}{\lambda_d(A_j)} = \sum_{j \in J} p(A_j) = 1 \quad (6.2.17)$$

The function f is called the *empirical probability density function* associated with the data \mathbf{x} and the partition \mathcal{A} . For the probability distribution defined by f , the empirical probability $p(A_j)$ is uniformly distributed over A_j for each $j \in J$. In the picture below, the red dots represent the data and the black lines define a partition of S into 9 rectangles. For the partition set A in the upper right, the empirical distribution would distribute probability $3/15 = 1/5$ uniformly over A . If the area of A is, say, 4, then $f(x) = 1/20$ for $x \in A$.




Figure 6.2.5: Empirical probability density function

Unlike the discrete case, we cannot recover the data from the empirical PDF. If we know the sample size, then of course we can determine the number of data points in A_j for each j , but not the precise location of these points in A_j . For this reason, the mean of the empirical PDF is not in general the same as the sample mean when the underlying variable is real-valued.

Histograms

Our next discussion is closely related to the previous one. Suppose again that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a variable that takes values in a set S and that $\mathcal{A} = (A_1, A_2, \dots, A_k)$ is a partition of S into k subsets. The sets in the partition are sometimes known as *classes*. The underlying variable may be discrete or continuous.

- The mapping that assigns frequencies to classes is known as a *frequency distribution* for the data set and the given partition.
- The mapping that assigns relative frequencies to classes is known as a *relative frequency distribution* for the data set and the given partition.
- In the case of a continuous variable, the mapping that assigns densities to classes is known as a *density distribution* for the data set and the given partition.

In dimensions 1 or 2, the bar graph any of these distributions, is known as a *histogram*. The histogram of a frequency distribution and the histogram of the corresponding relative frequency distribution look the same, except for a change of scale on the vertical axis. If the classes all have the same size, the histogram of the corresponding density histogram also looks the same, again except for a change of scale on the vertical axis. If the underlying variable is real-valued, the classes are usually intervals (discrete or continuous) and the midpoints of these intervals are sometimes referred to as *class marks*.




Figure 6.2.6: A density histogram

The whole purpose of constructing a partition and graphing one of these empirical distributions corresponding to the partition is to summarize and display the data in a meaningful way. Thus, there are some general guidelines in choosing the classes:

1. The number of classes should be moderate.
2. If possible, the classes should have the same size.

For highly skewed distributions, classes of different sizes are appropriate, to avoid numerous classes with very small frequencies. For a continuous variable with classes of different sizes, it is essential to use a density histogram, rather than a frequency or relative frequency histogram, otherwise the graphic is visually misleading, and in fact mathematically wrong.

It is important to realize that frequency data is inevitable for a continuous variable. For example, suppose that our variable represents the weight of a bag of M&Ms (in grams) and that our measuring device (a scale) is accurate to 0.01 grams. If we measure the weight of a bag as 50.32, then we are really saying that the weight is in the interval $[50.315, 50.324]$ (or perhaps some other interval, depending on how the measuring device works). Similarly, when two bags have the same *measured* weight, the apparent equality of the weights is really just an artifact of the imprecision of the measuring device; actually the two bags almost certainly do *not* have the exact same weight. Thus, two bags with the same measured weight really give us a frequency count of 2 for a certain interval.

Again, there is a trade-off between the *number* of classes and the *size* of the classes; these determine the *resolution* of the empirical distribution corresponding to the partition. At one extreme, when the class size is smaller than the accuracy of the recorded data, each class contains a single datum or no datum. In this case, there is no loss of information and we can recover the original data set from the frequency distribution (except for the *order* in which the data values were obtained). On the other hand, it can be hard to discern the shape of the data when we have many classes with small frequency. At the other extreme is a frequency distribution with one class that contains all of the possible values of the data set. In this case, all information is lost, except the number of the values in the data set. Between these two extreme cases, an empirical distribution gives us partial information, but not complete information. These intermediate cases can organize the data in a useful way.

Ogives

Suppose now the underlying variable is real-valued and that the set of possible values is partitioned into intervals (A_1, A_2, \dots, A_k) , with the endpoints of the intervals ordered from smallest to largest. Let n_j denote the frequency of class A_j , so that $p_j = n_j/n$ is the relative frequency of class A_j . Let t_j denote the class mark (midpoint) of class A_j . The *cumulative frequency* of class A_j is $N_j = \sum_{i=1}^j n_i$ and the *cumulative relative frequency* of class A_j is $P_j = \sum_{i=1}^j p_i = N_j/n$. Note that the cumulative frequencies increase from n_1 to n and the cumulative relative frequencies increase from p_1 to 1.

- The mapping that assigns cumulative frequencies to classes is known as a *cumulative frequency distribution* for the data set and the given partition. The polygonal graph that connects the points (t_j, N_j) for $j \in \{1, 2, \dots, k\}$ is the *cumulative frequency ogive*.
- The mapping that assigns cumulative relative frequencies to classes is known as a *cumulative relative frequency distribution* for the data set and the given partition. The polygonal graph that connects the points (t_j, P_j) for $j \in \{1, 2, \dots, k\}$ is the *cumulative relative frequency ogive*.

Note that the relative frequency ogive is simply the graph of the distribution function corresponding to the probability distribution that places probability p_j at t_j for each j .

Approximating the Mean

In the setting of the last subsection, suppose that we do not have the actual data \mathbf{x} , but just the frequency distribution. An approximate value of the sample mean is

$$\frac{1}{n} \sum_{j=1}^k n_j t_j = \sum_{j=1}^k p_j t_j \quad (6.2.18)$$

This approximation is based on the hope that the mean of the data values in each class is close to the midpoint of that class. In fact, the expression on the right is the expected value of the distribution that places probability p_j on class mark t_j for each j .

Exercises

Basic Properties

Suppose that x is the temperature (in degrees Fahrenheit) for a certain type of electronic component after 10 hours of operation.

1. Classify x by type and level of measurement.
2. A sample of 30 components has mean 113° . Find the sample mean if the temperature is converted to degrees Celsius. The transformation is $y = \frac{5}{9}(x - 32)$.

Answer

1. continuous, interval
2. 45°

Suppose that x is the length (in inches) of a machined part in a manufacturing process.

1. Classify x by type and level of measurement.
2. A sample of 50 parts has mean 10.0. Find the sample mean if length is measured in centimeters. The transformation is $y = 2.54x$.

Answer

1. continuous, ratio
2. 25.4

Suppose that x is the number of brothers and y the number of sisters for a person in a certain population. Thus, $z = x + y$ is the number of siblings.

1. Classify the variables by type and level of measurement.
2. For a sample of 100 persons, $m(\mathbf{x}) = 0.8$ and $m(\mathbf{y}) = 1.2$. Find $m(\mathbf{z})$.

Answer

1. discrete, ratio
2. 2.0

Professor Moriarity has a class of 25 students in her section of Stat 101 at Enormous State University (ESU). The mean grade on the first midterm exam was 64 (out of a possible 100 points). Professor Moriarity thinks the grades are a bit low and is considering various transformations for increasing the grades. In each case below give the mean of the transformed grades, or state that there is not enough information.

1. Add 10 points to each grade, so the transformation is $y = x + 10$.
2. Multiply each grade by 1.2, so the transformation is $z = 1.2x$
3. Use the transformation $w = 10\sqrt{x}$. Note that this is a non-linear transformation that curves the grades greatly at the low end and very little at the high end. For example, a grade of 100 is still 100, but a grade of 36 is transformed to 60.

One of the students did not study at all, and received a 10 on the midterm. Professor Moriarity considers this score to be an outlier.

4. What would the mean be if this score is omitted?

Answer

1. 74
2. 76.8
3. Not enough information
4. 66.25

Computational Exercises

All statistical software packages will compute means and proportions, draw dotplots and histograms, and in general perform the numerical and graphical procedures discussed in this section. For real statistical experiments, particularly those with large data sets, the use of statistical software is essential. On the other hand, there is some value in performing the computations by hand, with small, artificial data sets, in order to master the concepts and definitions. In this subsection, do the computations and draw the graphs with minimal technological aids.

Suppose that x is the number of math courses completed by an ESU student. A sample of 10 ESU students gives the data $x = (3, 1, 2, 0, 2, 4, 3, 2, 1, 2)$

1. Classify x by type and level of measurement.
2. Sketch the dotplot.
3. Compute the sample mean m from the definition and indicate its location on the dotplot.
4. Find the empirical density function f and sketch the graph.
5. Compute the sample mean m using f .
6. Find the empirical distribution function F and sketch the graph.

Answer

1. discrete, ratio
3. 2
4. $f(0) = 1/10, f(1) = 2/10, f(2) = 4/10, f(3) = 2/10, f(4) = 1/10$
5. 2
6. $F(x) = 0$ for $x < 0$, $F(x) = 1/10$ for $0 \leq x < 1$, $F(x) = 3/10$ for $1 \leq x < 2$, $F(x) = 7/10$ for $2 \leq x < 3$, $F(x) = 9/10$ for $3 \leq x < 4$, $F(x) = 1$ for $x \geq 4$

Suppose that a sample of size 12 from a discrete variable x has empirical density function given by $f(-2) = 1/12$, $f(-1) = 1/4$, $f(0) = 1/3$, $f(1) = 1/6$, $f(2) = 1/6$.

1. Sketch the graph of f .
2. Compute the sample mean m using f .
3. Find the empirical distribution function F
4. Give the sample values, ordered from smallest to largest.

Answer

2. $1/12$
3. $F(x) = 0$ for $x < -2$, $F(x) = 1/12$ for $-2 \leq x < -1$, $F(x) = 1/3$ for $-1 \leq x < 0$, $F(x) = 2/3$ for $0 \leq x < 1$, $F(x) = 5/6$ for $1 \leq x < 2$, $F(x) = 1$ for $x \geq 2$
4. $(-2, -1, -1, -1, 0, 0, 0, 0, 1, 1, 2, 2)$

The following table gives a frequency distribution for the commuting distance to the math/stat building (in miles) for a sample of ESU students.

Class	Freq	Rel Freq	Density	Cum Freq	Cum Rel Freq	Midpoint
(0, 2]	6					
(2, 6]	16					
(6, 10]	18					
Total		Density	Cum Freq	Cum Rel Freq	Midpoint	

Class	Freq	Rel Freq	Density	Cum Freq	Cum Rel Freq	Midpoint
(10, 20]	10					
Total		Density	Cum Freq	Cum Rel Freq	Midpoint	

1. Complete the table
2. Sketch the density histogram
3. Sketch the cumulative relative frequency ogive.
4. Compute an approximation to the mean

Answer

1.	Class	Freq	Rel Freq	Density	Cum Freq	Cum Rel Freq	Midpoint
	(0, 2]	6	0.12	0.06	6	0.12	1
	(2, 6]	16	0.32	0.08	22	0.44	4
	(6, 10]	18	0.36	0.09	40	0.80	8
	(10, 20])	10	0.20	0.02	50	1	15
	Total	50	1				

4. 7.28

App Exercises

In the interactive histogram, click on the x -axis at various points to generate a data set with at least 20 values. Vary the number of classes and switch between the frequency histogram and the relative frequency histogram. Note how the shape of the histogram changes as you perform these operations. Note in particular how the histogram loses resolution as you decrease the number of classes.

In the interactive histogram, click on the axis to generate a distribution of the given type with at least 30 points. Now vary the number of classes and note how the shape of the distribution changes.

1. A uniform distribution
2. A symmetric unimodal distribution
3. A unimodal distribution that is skewed right.
4. A unimodal distribution that is skewed left.
5. A symmetric bimodal distribution
6. A u -shaped distribution.

Data Analysis Exercises

Statistical software should be used for the problems in this subsection.

Consider the petal length and species variables in Fisher's iris data.

1. Classify the variables by type and level of measurement.
2. Compute the sample mean and plot a density histogram for petal length.
3. Compute the sample mean and plot a density histogram for petal length by species.

Answers

1. petal length: continuous, ratio. species: discrete, nominal
2. $m = 37.8$
3. $m(0) = 14.6$, $m(1) = 55.5$, $m(2) = 43.2$

Consider the erosion variable in the Challenger data set.

1. Classify the variable by type and level of measurement.
2. Compute the mean
3. Plot a density histogram with the classes $[0, 5)$, $[5, 40)$, $[40, 50)$, $[50, 60)$.

Answer

1. continuous, ratio
2. $m = 7.7$

Consider Michelson's velocity of light data.

1. Classify the variable by type and level of measurement.
2. Plot a density histogram.
3. Compute the sample mean.
4. Find the sample mean if the variable is converted to km/hr. The transformation is $y = x + 299\,000$

Answer

1. continuous, interval
3. $m = 852.4$
4. $m = 299\,852.4$

Consider Short's parallax of the sun data.

1. Classify the variable by type and level of measurement.
2. Plot a density histogram.
3. Compute the sample mean.
4. Find the sample mean if the variable is converted to degrees. There are 3600 seconds in a degree.
5. Find the sample mean if the variable is converted to radians. There are $\pi/180$ radians in a degree.

Answer

1. continuous, ratio
3. 8.616
4. 0.00239
5. 0.0000418

Consider Cavendish's density of the earth data.

1. Classify the variable by type and level of measurement.
2. Compute the sample mean.
3. Plot a density histogram.

Answer

1. continuous, ratio
2. $m = 5.448$

Consider the M&M data.

1. Classify the variables by type and level of measurement.
2. Compute the sample mean for each color count variable.
3. Compute the sample mean for the total number of candies, using the results from (b).
4. Plot a relative frequency histogram for the total number of candies.
5. Compute the sample mean and plot a density histogram for the net weight.

Answer

1. color counts: discrete ratio. net weight: continuous ratio.
2. $m(r) = 9.60$, $m(g) = 7.40$, $m(bl) = 7.23$, $m(o) = 6.63$, $m(y) = 13.77$, $m(br) = 12.47$
3. $m(n) = 57.10$
5. $m(w) = 49.215$

Consider the body weight, species, and gender variables in the Cicada data.

1. Classify the variables by type and level of measurement.
2. Compute the relative frequency function for species and plot the graph.
3. Compute the relative frequency function for gender and plot the graph.
4. Compute the sample mean and plot a density histogram for body weight.
5. Compute the sample mean and plot a density histogram for body weight by species.
6. Compute the sample mean and plot a density histogram for body weight by gender.

Answer

1. body weight: continuous, ratio. species: discrete, nominal. gender: discrete, nominal.
2. $f(0) = 0.423$, $f(1) = 0.519$, $f(2) = 0.058$
3. $f(0) = 0.567$, $f(1) = 0.433$
4. $m = 0.180$
5. $m(0) = 0.168$, $m(1) = 0.185$, $m(2) = 0.225$
6. $m(0) = 0.206$, $m(1) = 0.145$

Consider Pearson's height data.

1. Classify the variables by type and level of measurement.
2. Compute the sample mean and plot a density histogram for the height of the father.
3. Compute the sample mean and plot a density histogram for the height of the son.

Answer

1. continuous ratio
2. $m(f) = 67.69$
3. $m(s) = 68.68$

This page titled [6.2: The Sample Mean](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.3: The Law of Large Numbers

Basic Theory

This section continues the discussion of the sample mean from the last section, but we now consider the more interesting setting where the variables are random. Specifically, suppose that we have a basic random experiment with an underlying probability measure \mathbb{P} , and that X is random variable for the experiment. Suppose now that we perform n independent replications of the basic experiment. This defines a new, compound experiment with a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots, X_n)$, each with the same distribution as X . Recall that in statistical terms, \mathbf{X} is a random sample of size n from the distribution of X . All of the relevant statistics discussed in the the previous section, are defined for \mathbf{X} , but of course now these statistics are random variables with distributions of their own. For the most part, we use the notation established previously, except that for the usual convention of denoting random variables with capital letters. Of course, the deterministic properties and relations established previously apply as well. When we acutally *run* the experiment and observe the values $\mathbf{x} = (x_1, x_2, \dots, x_n)$ of the random variables, then we are precisely in the setting of the previous section.

Suppose now that the basic variable X is real valued, and let $\mu = \mathbb{E}(X)$ denote the expected value of X and $\sigma^2 = \text{var}(X)$ the variance of X (assumed finite). The *sample mean* is

$$M = \frac{1}{n} \sum_{i=1}^n X_i \quad (6.3.1)$$

Often the distribution mean μ is unknown and the sample mean M is used as an estimator of this unknown parameter.

Moments

The mean and variance of M are

1. $\mathbb{E}(M) = \mu$
2. $\text{var}(M) = \sigma^2/n$

Proof

1. This follows from the linear property of expected value:

$$\mathbb{E}(M) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}(X_i) = \frac{1}{n} \sum_{i=1}^n \mu = \frac{1}{n} n\mu = \mu \quad (6.3.2)$$

2. This follows from basic properties of variance. Recall in particular that the variance of the sum of independent variables is the sum of the variances.

$$\text{var}(M) = \frac{1}{n^2} \sum_{i=1}^n \text{var}(X_i) = \frac{1}{n^2} \sum_{i=1}^n \sigma^2 = \frac{1}{n^2} n\sigma^2 = \frac{\sigma^2}{n} \quad (6.3.3)$$

Part (a) means that the sample mean M is an *unbiased* estimator of the distribution mean μ . Therefore, the variance of M is the *mean square error*, when M is used as an estimator of μ . Note that the variance of M is an increasing function of the distribution variance and a decreasing function of the sample size. Both of these make intuitive sense if we think of the sample mean M as an estimator of the distribution mean μ . The fact that the mean square error (variance in this case) decreases to 0 as the sample size n increases to ∞ means that the sample mean M is a *consistent* estimator of the distribution mean μ .

Recall that $X_i - M$ is the *deviation* of X_i from M , that is, the directed distance from M to X_i . The following theorem states that the sample mean is uncorrelated with each deviation, a result that will be crucial for showing the independence of the sample mean and the sample variance when the sampling distribution is normal.

M and $X_i - M$ are uncorrelated.

Proof

This result follows from simple properties of covariance. Note that $\text{cov}(M, X_i - M) = \text{cov}(M, X_i) - \text{cov}(M, M)$. By independence,

$$\text{cov}(M, X_i) = \text{cov}\left(\frac{1}{n} \sum_{j=1}^n X_j, X_i\right) = \frac{1}{n} \sum_{j=1}^n \text{cov}(X_j, X_i) = \frac{1}{n} \text{cov}(X_i, X_i) = \frac{1}{n} \text{var}(X_i) = \frac{\sigma^2}{n} \quad (6.3.4)$$

But by previous theorem, $\text{cov}(M, M) = \text{var}(M) = \sigma^2/n$.

The Weak and Strong Laws of Large Numbers

The *law of large numbers* states that the sample mean converges to the distribution mean as the sample size increases, and is one of the fundamental theorems of probability. There are different versions of the law, depending on the *mode* of convergence.

Suppose again that X is a real-valued random variable for our basic experiment, with mean μ and standard deviation σ (assumed finite). We repeat the basic experiment indefinitely to create a new, compound experiment with an infinite sequence of independent random variables (X_1, X_2, \dots) , each with the same distribution as X . In statistical terms, we are *sampling* from the distribution of X . In probabilistic terms, we have an independent, identically distributed (IID) sequence. For each n , let M_n denote the sample mean of the first n sample variables:

$$M_n = \frac{1}{n} \sum_{i=1}^n X_i \quad (6.3.5)$$

From the result above on [variance](#), note that $\text{var}(M_n) = \mathbb{E}[(M_n - \mu)^2] \rightarrow 0$ as $n \rightarrow \infty$. This means that $M_n \rightarrow \mu$ as $n \rightarrow \infty$ in mean square. As stated in the next theorem, $M_n \rightarrow \mu$ as $n \rightarrow \infty$ in probability as well.

$\mathbb{P}(|M_n - \mu| > \epsilon) \rightarrow 0$ as $n \rightarrow \infty$ for every $\epsilon > 0$.

Proof

This follows from Chebyshev's inequality:

$$\mathbb{P}(|M_n - \mu| > \epsilon) \leq \frac{\text{var}(M_n)}{\epsilon^2} = \frac{\sigma^2}{n\epsilon^2} \rightarrow 0 \text{ as } n \rightarrow \infty \quad (6.3.6)$$

Recall that in general, convergence in mean square implies convergence in probability. The convergence of the sample mean to the distribution mean in mean square and in probability are known as *weak laws of large numbers*.

Finally, the *strong law of large numbers* states that the sample mean M_n converges to the distribution mean μ with probability 1. As the name suggests, this is a much stronger result than the weak laws. We will need some additional notation for the proof. First let $Y_n = \sum_{i=1}^n X_i$ so that $M_n = Y_n/n$. Next, recall the definitions of the *positive and negative parts* a real number x : $x^+ = \max\{x, 0\}$, $x^- = \max\{-x, 0\}$. Note that $x^+ \geq 0$, $x^- \geq 0$, $x = x^+ - x^-$, and $|x| = x^+ + x^-$.

$M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1.

Proof

The proof is in three major steps. The first step is to show that with probability 1, $M_{n^2} \rightarrow \mu$ as $n \rightarrow \infty$. From Chebyshev's inequality, $\mathbb{P}(|M_{n^2} - \mu| > \epsilon) \leq \sigma^2/n^2\epsilon^2$ for every $n \in \mathbb{N}_+$ and every $\epsilon > 0$. Since $\sum_{n=1}^{\infty} \sigma^2/n^2\epsilon^2 < \infty$, it follows from the first Borel-Cantelli lemma that for every $\epsilon > 0$,

$$\mathbb{P}(|M_{n^2} - \mu| > \epsilon \text{ for infinitely many } n \in \mathbb{N}_+) = 0 \quad (6.3.7)$$

Next, from Boole's inequality it follows that

$$\mathbb{P}(\text{For some rational } \epsilon > 0, |M_{n^2} - \mu| > \epsilon \text{ for infinitely many } n \in \mathbb{N}_+) = 0 \quad (6.3.8)$$

This is equivalent to the statement that $M_{n^2} \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1.

For our next step, we will show that if the underlying sampling variable is nonnegative, so that $\mathbb{P}(X \geq 0) = 1$, then $M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1. Note first that with probability 1, Y_n is increasing in n . For $n \in \mathbb{N}_+$, let k_n be the unique positive integer such that $k_n^2 \leq n < (k_n + 1)^2$. From the increasing property and simple algebra, it follows that with probability 1,

$$\frac{Y_{k_n^2}}{(k_n + 1)^2} \leq \frac{Y_n}{n} \leq \frac{Y_{(k_n+1)^2}}{k_n^2} \quad (6.3.9)$$

From our first step, with probability 1,

$$\frac{Y_{k_n^2}}{(k_n + 1)^2} = \frac{Y_{k_n^2}}{k_n^2} \frac{k_n^2}{(k_n + 1)^2} \rightarrow \mu \text{ as } n \rightarrow \infty \quad (6.3.10)$$

Similarly with probability 1

$$\frac{Y_{(k_n+1)^2}}{k_n^2} = \frac{Y_{(k_n+1)^2}}{(k_n + 1)^2} \frac{(k_n + 1)^2}{k_n^2} \rightarrow \mu \text{ as } n \rightarrow \infty \quad (6.3.11)$$

Finally by the squeeze theorem for limits it follows that with probability 1, $M_n = Y_n/n \rightarrow \mu$ as $n \rightarrow \infty$.

Finally we relax the condition that the underlying sampling variable X is nonnegative. From step two, it follows that $\frac{1}{n} \sum_{i=1}^n X_i^+ \rightarrow \mathbb{E}(X^+)$ as $n \rightarrow \infty$ with probability 1, and $\frac{1}{n} \sum_{i=1}^n X_i^- \rightarrow \mathbb{E}(X^-)$ as $n \rightarrow \infty$ with probability 1. Now from algebra and

the linearity of expected value, with probability 1,

$$\frac{1}{n} \sum_{i=1}^n X_i = \frac{1}{n} \sum_{i=1}^n (X_i^+ - X_i^-) = \frac{1}{n} \sum_{i=1}^n X_i^+ - \sum_{i=1}^n X_i^- \rightarrow \mathbb{E}(X^+) - \mathbb{E}(X^-) = \mathbb{E}(X^+ - X^-) = \mathbb{E}(X) \text{ as } n \rightarrow \infty \quad (6.3.12)$$

The proof of the strong law of large numbers given above requires that the variance of the sampling distribution be finite (note that this is critical in the first step). However, there are better proofs that only require that $\mathbb{E}(|X|) < \infty$. An elegant proof showing that $M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1 and in mean, using backwards martingales, is given in the chapter on martingales. In the next few paragraphs, we apply the law of large numbers to some of the special statistics studied in the previous section.

Empirical Probability

Suppose that X is the outcome random variable for a basic experiment, with sample space S and probability measure \mathbb{P} . Now suppose that we repeat the basic experiment indefinitely to form a sequence of independent random variables (X_1, X_2, \dots) each with the same distribution as X . That is, we sample from the distribution of X . For $A \subseteq S$, let $P_n(A)$ denote the empirical probability of A corresponding to the sample (X_1, X_2, \dots, X_n) :

$$P_n(A) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i \in A) \quad (6.3.13)$$

Now of course, $P_n(A)$ is a random variable for each event A . In fact, the sum $\sum_{i=1}^n \mathbf{1}(X_i \in A)$ has the binomial distribution with parameters n and $\mathbb{P}(A)$.

For each event A ,

1. $\mathbb{E}[P_n(A)] = \mathbb{P}(A)$
2. $\text{var}[P_n(A)] = \frac{1}{n} \mathbb{P}(A)[1 - \mathbb{P}(A)]$
3. $P_n(A) \rightarrow \mathbb{P}(A)$ as $n \rightarrow \infty$ with probability 1.

Proof

These results follow from the results of this section, since $P_n(A)$ is the sample mean for the random sample $\{\mathbf{1}(X_i \in A) : i \in \{1, 2, \dots, n\}\}$ from the distribution of $\mathbf{1}(X \in A)$.

This special case of the law of large numbers is central to the very concept of probability: the relative frequency of an event converges to the probability of the event as the experiment is repeated.

The Empirical Distribution Function

Suppose now that X is a real-valued random variable for a basic experiment. Recall that the distribution function of X is the function F given by

$$F(x) = \mathbb{P}(X \leq x), \quad x \in \mathbb{R} \quad (6.3.14)$$

Now suppose that we repeat the basic experiment indefinitely to form a sequence of independent random variables (X_1, X_2, \dots) , each with the same distribution as X . That is, we sample from the distribution of X . Let F_n denote the empirical distribution function corresponding to the sample (X_1, X_2, \dots, X_n) :

$$F_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i \leq x), \quad x \in \mathbb{R} \quad (6.3.15)$$

Now, of course, $F_n(x)$ is a random variable for each $x \in \mathbb{R}$. In fact, the sum $\sum_{i=1}^n \mathbf{1}(X_i \leq x)$ has the binomial distribution with parameters n and $F(x)$.

For each $x \in \mathbb{R}$,

1. $\mathbb{E}[F_n(x)] = F(x)$
2. $\text{var}[F_n(x)] = \frac{1}{n} F(x)[1 - F(x)]$
3. $F_n(x) \rightarrow F(x)$ as $n \rightarrow \infty$ with probability 1.

Proof

These results follow immediately from the results in this section, since $F_n(x)$ is the sample mean for the random sample $\{\mathbf{1}(X_i \leq x) : i \in \{1, 2, \dots, n\}\}$ from the distribution of $\mathbf{1}(X \leq x)$.

Empirical Density for a Discrete Variable

Suppose now that X is a random variable for a basic experiment with a discrete distribution on a countable set S . Recall that the probability density function of X is the function f given by

$$f(x) = \mathbb{P}(X = x), \quad x \in S \quad (6.3.16)$$

Now suppose that we repeat the basic experiment to form a sequence of independent random variables (X_1, X_2, \dots) each with the same distribution as X . That is, we sample from the distribution of X . Let f_n denote the empirical probability density function corresponding to the sample (X_1, X_2, \dots, X_n) :

$$f_n(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i = x), \quad x \in S \quad (6.3.17)$$

Now, of course, $f_n(x)$ is a random variable for each $x \in S$. In fact, the sum $\sum_{i=1}^n \mathbf{1}(X_i = x)$ has the binomial distribution with parameters n and $f(x)$.

For each $x \in S$,

1. $\mathbb{E}[f_n(x)] = f(x)$
2. $\text{var}[f_n(x)] = \frac{1}{n} f(x) [1 - f(x)]$
3. $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$ with probability 1.

Proof

These results follow immediately from the results in this section, since $f_n(x)$ is the sample mean for the random sample $\{\mathbf{1}(X_i = x) : i \in \{1, 2, \dots, n\}\}$ from the distribution of $\mathbf{1}(X = x)$.

Recall that a countable intersection of events with probability 1 still has probability 1. Thus, in the context of the previous theorem, we actually have

$$\mathbb{P}[f_n(x) \rightarrow f(x) \text{ as } n \rightarrow \infty \text{ for every } x \in S] = 1 \quad (6.3.18)$$

Empirical Density for a Continuous Variable

Suppose now that X is a random variable for a basic experiment, with a continuous distribution on $S \subseteq \mathbb{R}^d$, and that X has probability density function f . Technically, f is the probability density function with respect to the standard (Lebesgue) measure λ_d . Thus, by definition,

$$\mathbb{P}(X \in A) = \int_A f(x) dx, \quad A \subseteq S \quad (6.3.19)$$

Again we repeat the basic experiment to generate a sequence of independent random variables (X_1, X_2, \dots) each with the same distribution as X . That is, we sample from the distribution of X . Suppose now that $\mathcal{A} = \{A_j : j \in J\}$ is a partition of S into a countable number of subsets, each with positive, finite size. Let f_n denote the empirical probability density function corresponding to the sample (X_1, X_2, \dots, X_n) and the partition \mathcal{A} :

$$f_n(x) = \frac{P_n(A_j)}{\lambda_d(A_j)} = \frac{1}{n \lambda_d(A_j)} \sum_{i=1}^n \mathbf{1}(X_i \in A_j); \quad j \in J, x \in A_j \quad (6.3.20)$$

Of course now, $f_n(x)$ is a random variable for each $x \in S$. If the partition is sufficiently fine (so that $\lambda_d(A_j)$ is small for each j), and if the sample size n is sufficiently large, then by the law of large numbers,

$$f_n(x) \approx f(x), \quad x \in S \quad (6.3.21)$$

Exercises

Simulation Exercises

In the dice experiment, recall that the dice scores form a random sample from the specified die distribution. Select the average random variable, which is the sample mean of the sample of dice scores. For each die distribution, start with 1 die and increase the sample size n . Note how the distribution of the sample mean begins to resemble a point mass distribution. Note also that the mean of the sample mean stays the same, but the standard deviation of the sample mean decreases. For selected values of n and selected die distributions, run the simulation 1000 times and compare the relative frequency function of the sample mean to the true probability density function, and compare the empirical moments of the sample mean to the true moments.

Several apps in this project are simulations of random experiments with events of interest. When you run the experiment, you are performing independent replications of the experiment. In most cases, the app displays the relative frequency of the event and its complement, both graphically in blue, and numerically in a table. When you run the experiment, the relative frequencies are shown graphically in red and also numerically.

In the simulation of Buffon's coin experiment, the event of interest is that the coin crosses a crack. For various values of the parameter (the radius of the coin), run the experiment 1000 times and compare the relative frequency of the event to the true probability.

In the simulation of Bertrand's experiment, the event of interest is that a "random chord" on a circle will be longer than the length of a side of the inscribed equilateral triangle. For each of the various models, run the experiment 1000 times and compare the relative frequency of the event to the true probability.

Many of the apps in this project are simulations of experiments which result in discrete variables. When you run the simulation, you are performing independent replications of the experiment. In most cases, the app displays the true probability density function numerically in a table and visually as a blue bar graph. When you run the simulation, the relative frequency function is also shown numerically in the table and visually as a red bar graph.

In the simulation of the binomial coin experiment, select the number of heads. For selected values of the parameters, run the simulation 1000 times and compare the sample mean to the distribution mean, and compare the empirical density function to the probability density function.

In the simulation of the matching experiment, the random variable is the number of matches. For selected values of the parameter, run the simulation 1000 times and compare the sample mean and the distribution mean, and compare the empirical density function to the probability density function.

In the poker experiment, the random variable is the type of hand. Run the simulation 1000 times and compare the empirical density function to the true probability density function.

Many of the apps in this project are simulations of experiments which result in variables with continuous distributions. When you run the simulation, you are performing independent replications of the experiment. In most cases, the app displays the true probability density function visually as a blue graph. When you run the simulation, an empirical density function, based on a partition, is also shown visually as a red bar graph.

In the simulation of the gamma experiment, the random variable represents a random arrival time. For selected values of the parameters, run the experiment 1000 times and compare the sample mean to the distribution mean, and compare the empirical density function to the probability density function.

In the special distribution simulator, select the normal distribution. For various values of the parameters (the mean and standard deviation), run the experiment 1000 times and compare the sample mean to the distribution mean, and compare the empirical density function to the probability density function.

Probability Exercises

Suppose that X has probability density function $f(x) = 12x^2(1-x)$ for $0 \leq x \leq 1$. The distribution of X is a member of the beta family. Compute each of the following

1. $\mathbb{E}(X)$
2. $\text{var}(X)$
3. $\mathbb{P}(X \leq \frac{1}{2})$

Answer

1. $\frac{3}{5}$
2. $\frac{1}{25}$
3. $\frac{5}{16}$

Suppose now that (X_1, X_2, \dots, X_9) is a random sample of size 9 from the distribution in the previous problem. Find the expected value and variance of each of the following random variables:

1. The sample mean M
2. The empirical probability $P([0, \frac{1}{2}])$

Answer

1. $\frac{3}{5}, \frac{1}{225}$
2. $\frac{5}{16}, \frac{55}{2304}$

Suppose that X has probability density function $f(x) = \frac{3}{x^4}$ for $1 \leq x < \infty$. The distribution of X is a member of the Pareto family. Compute each of the following

1. $\mathbb{E}(X)$
2. $\text{var}(X)$

3. $\mathbb{P}(2 \leq X \leq 3)$

Answer

1. $\frac{3}{2}$
2. $\frac{3}{4}$
3. $\frac{19}{216}$

Suppose now that $(X_1, X_2, \dots, X_{16})$ is a random sample of size 16 from the distribution in the previous problem. Find the expected value and variance of each of the following random variables:

1. The sample mean M
2. The empirical probability $P([2, 3])$

Answer

1. $\frac{3}{2}, \frac{3}{64}$
2. $\frac{19}{216}, \frac{3743}{746496}$

Recall that for an *ace-six flat die*, faces 1 and 6 have probability $\frac{1}{4}$ each, while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each. Let X denote the score when an ace-six flat die is thrown. Compute each of the following:

1. The probability density function $f(x)$ for $x \in \{1, 2, 3, 4, 5, 6\}$
2. The distribution function $F(x)$ for $x \in \{1, 2, 3, 4, 5, 6\}$
3. $\mathbb{E}(X)$
4. $\text{var}(X)$

Answer

1. $f(x) = \frac{1}{4}, x \in \{1, 6\}; f(x) = \frac{1}{8}, x \in \{2, 3, 4, 5\}$
2. $F(1) = \frac{1}{4}, F(2) = \frac{3}{8}, F(3) = \frac{1}{2}, F(4) = \frac{5}{8}, F(5) = \frac{3}{4}, F(6) = 1$
3. $\mathbb{E}(X) = \frac{7}{2}$
4. $\text{var}(X) = \frac{15}{4}$

Suppose now that an ace-six flat die is thrown n times. Find the expected value and variance of each of the following random variables:

1. The empirical probability density function $f_n(x)$ for $x \in \{1, 2, 3, 4, 5, 6\}$
2. The empirical distribution function $F_n(x)$ for $x \in \{1, 2, 3, 4, 5, 6\}$
3. The average score M

Answer

1. $\mathbb{E}[f_n(x)] = \frac{1}{4}, x \in \{1, 6\}; \mathbb{E}[f_n(x)] = \frac{1}{8}, x \in \{2, 3, 4, 5\}$
2. $\text{var}[f_n(x)] = \frac{3}{16n}, x \in \{1, 6\}; \text{var}[f_n(x)] = \frac{7}{64n}, x \in \{2, 3, 4, 5\}$
3. $\mathbb{E}[F_n(1)] = \frac{1}{4}, \mathbb{E}[F_n(2)] = \frac{3}{8}, \mathbb{E}[F_n(3)] = \frac{1}{2}, \mathbb{E}[F_n(4)] = \frac{5}{8}, \mathbb{E}[F_n(5)] = \frac{3}{4}, \mathbb{E}[F_n(6)] = 1$
4. $\text{var}[F_n(1)] = \frac{3}{16n}, \text{var}[F_n(2)] = \frac{15}{64n}, \text{var}[F_n(3)] = \frac{1}{4n}, \text{var}[F_n(4)] = \frac{15}{64n}, \text{var}[F_n(5)] = \frac{3}{16n}, \text{var}[F_n(6)] = 0$
5. $\mathbb{E}(M) = \frac{7}{2}, \text{var}(M) = \frac{15}{4n}$

This page titled [6.3: The Law of Large Numbers](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.4: The Central Limit Theorem

The central limit theorem and the law of large numbers are the two *fundamental theorems* of probability. Roughly, the central limit theorem states that the distribution of the sum (or average) of a large number of independent, identically distributed variables will be approximately normal, regardless of the underlying distribution. The importance of the central limit theorem is hard to overstate; indeed it is the reason that many statistical procedures work.

Partial Sum Processes

Definitions

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent, identically distributed, real-valued random variables with common probability density function f , mean μ , and variance σ^2 . We assume that $0 < \sigma < \infty$, so that in particular, the random variables really are *random* and not constants. Let

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (6.4.1)$$

Note that by convention, $Y_0 = 0$, since the sum is over an empty index set. The random process $\mathbf{Y} = (Y_0, Y_1, Y_2, \dots)$ is called the *partial sum process* associated with \mathbf{X} . Special types of partial sum processes have been studied in many places in this text; in particular see

- the binomial distribution in the setting of Bernoulli trials
- the negative binomial distribution in the setting of Bernoulli trials
- the gamma distribution in the Poisson process
- the arrival times in a general renewal process

Recall that in statistical terms, the sequence \mathbf{X} corresponds to sampling from the underlying distribution. In particular, (X_1, X_2, \dots, X_n) is a random sample of size n from the distribution, and the corresponding sample mean is

$$M_n = \frac{Y_n}{n} = \frac{1}{n} \sum_{i=1}^n X_i \quad (6.4.2)$$

By the law of large numbers, $M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1.

Stationary, Independent Increments

The partial sum process corresponding to a sequence of independent, identically distributed variables has two important properties, and these properties essentially characterize such processes.

If $m \leq n$ then $Y_n - Y_m$ has the same distribution as Y_{n-m} . Thus the process \mathbf{Y} has *stationary increments*.

Proof

Note that $Y_n - Y_m = \sum_{i=m+1}^n X_i$ and is the sum of $n - m$ independent variables, each with the common distribution. Of course, Y_{n-m} is also the sum of $n - m$ independent variables, each with the common distribution.

Note however that $Y_n - Y_m$ and Y_{n-m} are very different random variables; the theorem simply states that they have the same *distribution*.

If $n_1 \leq n_2 \leq n_3 \leq \dots$ then $(Y_{n_1}, Y_{n_2} - Y_{n_1}, Y_{n_3} - Y_{n_2}, \dots)$ is a sequence of independent random variables. Thus the process \mathbf{Y} has *independent increments*.

Proof

The terms in the sequence of increments $(Y_{n_1}, Y_{n_2} - Y_{n_1}, Y_{n_3} - Y_{n_2}, \dots)$ are sums over disjoint collections of terms in the sequence \mathbf{X} . Since the sequence \mathbf{X} is independent, so is the sequence of increments.

Conversely, suppose that $\mathbf{V} = (V_0, V_1, V_2, \dots)$ is a random process with stationary, independent increments. Define $U_i = V_i - V_{i-1}$ for $i \in \mathbb{N}_+$. Then $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent, identically distributed variables and \mathbf{V} is

the partial sum process associated with U .

Thus, partial sum processes are the only discrete-time random processes that have stationary, independent increments. An interesting, and much harder problem, is to characterize the continuous-time processes that have stationary independent increments. The Poisson counting process has stationary independent increments, as does the Brownian motion process.

Moments

If $n \in \mathbb{N}$ then

1. $\mathbb{E}(Y_n) = n\mu$
2. $\text{var}(Y_n) = n\sigma^2$

Proof

The results follow from basic properties of expected value and variance. Expected value is a linear operation so $\mathbb{E}(Y_n) = \sum_{i=1}^n \mathbb{E}(X_i) = n\mu$. By independence, $\text{var}(Y_n) = \sum_{i=1}^n \text{var}(X_i) = n\sigma^2$.

If $n \in \mathbb{N}_+$ and $m \in \mathbb{N}$ with $m \leq n$ then

1. $\text{cov}(Y_m, Y_n) = m\sigma^2$
2. $\text{cor}(Y_m, Y_n) = \sqrt{\frac{m}{n}}$
3. $\mathbb{E}(Y_m Y_n) = m\sigma^2 + mn\mu^2$

Proof

1. Note that $Y_n = Y_m + (Y_n - Y_m)$. This follows from basic properties of covariance, and [Theorem 1](#) and [Theorem 2](#):

$$\text{cov}(Y_m, Y_n) = \text{cov}(Y_m, Y_m) + \text{cov}(Y_m, Y_n - Y_m) = \text{var}(Y_m) + 0 = m\sigma^2 \quad (6.4.3)$$

2. This result follows from part (a) and [Theorem 4](#)

$$\text{cor}(Y_m, Y_n) = \frac{\text{cov}(Y_m, Y_n)}{\text{sd}(Y_m)\text{sd}(Y_n)} = \frac{m\sigma^2}{\sqrt{m\sigma^2}\sqrt{n\sigma^2}} = \sqrt{\frac{m}{n}} \quad (6.4.4)$$

3. This result also follows from part (a) and [Theorem 4](#): $\mathbb{E}(Y_m Y_n) = \text{cov}(Y_m, Y_n) + \mathbb{E}(Y_m)\mathbb{E}(Y_n) = m\sigma^2 + m\mu n\mu$

If X has moment generating function G then Y_n has moment generating function G^n .

Proof

This follows from a basic property of generating functions: the generating function of a sum of independent variables is the product of the generating functions of the terms.

Distributions

Suppose that X has either a discrete distribution or a continuous distribution with probability density function f . Then the probability density function of Y_n is $f^{*n} = f * f * \dots * f$, the convolution power of f of order n .

Proof

This follows from a basic property of PDFs: the pdf of a sum of independent variables is the convolution of the PDFs of the terms.

More generally, we can use the stationary and independence properties to find the joint distributions of the partial sum process:

If $n_1 < n_2 < \dots < n_k$ then $(Y_{n_1}, Y_{n_2}, \dots, Y_{n_k})$ has joint probability density function

$$f_{n_1, n_2, \dots, n_k}(y_1, y_2, \dots, y_k) = f^{*n_1}(y_1) f^{*(n_2 - n_1)}(y_2 - y_1) \dots f^{*(n_k - n_{k-1})}(y_k - y_{k-1}), \quad (y_1, y_2, \dots, y_k) \in \mathbb{R}^k \quad (6.4.5)$$

Proof

This follows from the multivariate change of variables theorem.

The Central Limit Theorem

First, let's make the central limit theorem more precise. From [Theorem 4](#), we cannot expect Y_n itself to have a limiting distribution. Note that $\text{var}(Y_n) \rightarrow \infty$ as $n \rightarrow \infty$ since $\sigma > 0$, and $\mathbb{E}(Y_n) \rightarrow \infty$ as $n \rightarrow \infty$ if $\mu > 0$ while $\mathbb{E}(Y_n) \rightarrow -\infty$ as $n \rightarrow \infty$ if $\mu < 0$. Similarly, we know that $M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1, so the limiting distribution of the sample mean is degenerate. Thus, to obtain a limiting distribution of Y_n or M_n that is not degenerate, we need to consider, not these variables themselves, but rather the common *standard score*. Thus, let

$$Z_n = \frac{Y_n - n\mu}{\sqrt{n}\sigma} = \frac{M_n - \mu}{\sigma/\sqrt{n}} \quad (6.4.6)$$

Z_n has mean 0 and variance 1.

1. $\mathbb{E}(Z_n) = 0$
2. $\text{var}(Z_n) = 1$

Proof

These results follow from basic properties of expected value and variance, and are true for the standard score associated with any random variable. Recall also that the standard score of a variable is invariant under linear transformations with positive slope. The fact that the standard score of Y_n and the standard score of M_n are the same is a special case of this.

The precise statement of the central limit theorem is that the distribution of the standard score Z_n converges to the standard normal distribution as $n \rightarrow \infty$. Recall that the standard normal distribution has probability density function

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}, \quad z \in \mathbb{R} \quad (6.4.7)$$

and is studied in more detail in the chapter on special distributions. A special case of the central limit theorem (to Bernoulli trials), dates to Abraham De Moivre. The term *central limit theorem* was coined by George Pólya in 1920. By definition of convergence in distribution, the central limit theorem states that $F_n(z) \rightarrow \Phi(z)$ as $n \rightarrow \infty$ for each $z \in \mathbb{R}$, where F_n is the distribution function of Z_n and Φ is the standard normal distribution function:

$$\Phi(z) = \int_{-\infty}^z \phi(x) dx = \int_{-\infty}^z \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} dx, \quad z \in \mathbb{R} \quad (6.4.8)$$

An equivalent statment of the central limit theorem involves convergence of the corresponding characteristic functions. This is the version that we will give and prove, but first we need a generalization of a famous limit from calculus.

Suppose that (a_1, a_2, \dots) is a sequence of real numbers and that $a_n \rightarrow a \in \mathbb{R}$ as $n \rightarrow \infty$. Then

$$\left(1 + \frac{a_n}{n}\right)^n \rightarrow e^a \text{ as } n \rightarrow \infty \quad (6.4.9)$$

Now let χ denote the characteristic function of the standard score of the sample variable X , and let χ_n denote the characteristic function of the standard score Z_n :

$$\chi(t) = \mathbb{E} \left[\exp \left(it \frac{X - \mu}{\sigma} \right) \right], \quad \chi_n(t) = \mathbb{E}[\exp(itZ_n)]; \quad t \in \mathbb{R} \quad (6.4.10)$$

Recall that $t \mapsto e^{-\frac{1}{2}t^2}$ is the characteristic function of the standard normal distribution. We can now give a proof.

The central limit theorem. The distribution of Z_n converges to the standard normal distribution as $n \rightarrow \infty$. That is, $\chi_n(t) \rightarrow e^{-\frac{1}{2}t^2}$ as $n \rightarrow \infty$ for each $t \in \mathbb{R}$.

Proof

Note that $\chi(0) = 1$, $\chi'(0) = 0$, $\chi''(0) = -1$. Next

$$Z_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{X_i - \mu}{\sigma} \quad (6.4.11)$$

From properties of characteristic functions, $\chi_n(t) = \chi^n(t/\sqrt{n})$ for $t \in \mathbb{R}$. By Taylor's theorem (named after Brook Taylor),

$$\chi\left(\frac{t}{\sqrt{n}}\right) = 1 + \frac{1}{2}\chi''(s_n)\frac{t^2}{n} \text{ where } |s_n| \leq \frac{|t|}{n} \quad (6.4.12)$$

But $s_n \rightarrow 0$ and hence $\chi''(s_n) \rightarrow -1$ as $n \rightarrow \infty$. Finally,

$$\chi_n(t) = \left[1 + \frac{1}{2}\chi''(s_n)\frac{t^2}{n}\right]^n \rightarrow e^{-\frac{1}{2}t^2} \text{ as } n \rightarrow \infty \quad (6.4.13)$$

Normal Approximations

The central limit theorem implies that if the sample size n is “large” then the distribution of the partial sum Y_n is approximately normal with mean $n\mu$ and variance $n\sigma^2$. Equivalently the sample mean M_n is approximately normal with mean μ and variance σ^2/n . The central limit theorem is of fundamental importance, because it means that we can approximate the distribution of certain statistics, even if we know very little about the underlying sampling distribution.

Of course, the term “large” is relative. Roughly, the more “abnormal” the basic distribution, the larger n must be for normal approximations to work well. The rule of thumb is that a sample size n of at least 30 will usually suffice if the basic distribution is not too weird; although for many distributions smaller n will do.

Let Y denote the sum of the variables in a random sample of size 30 from the uniform distribution on $[0, 1]$. Find normal approximations to each of the following:

1. $\mathbb{P}(13 < Y < 18)$
2. The 90th percentile of Y

Answer

1. 0.8682
2. 17.03

Random variable Y in the previous exercise has the *Irwin-Hall distribution* of order 30. The Irwin-Hall distributions are studied in more detail in the chapter on Special Distributions and are named for Joseph Irwin and Phillip Hall.

In the special distribution simulator, select the Irwin-Hall distribution. Vary n from 1 to 10 and note the shape of the probability density function. With $n = 10$ run the experiment 1000 times and compare the empirical density function to the true probability density function.

Let M denote the sample mean of a random sample of size 50 from the distribution with probability density function $f(x) = \frac{3}{x^4}$ for $1 \leq x < \infty$. This is a *Pareto distribution*, named for Vilfredo Pareto. Find normal approximations to each of the following:

1. $\mathbb{P}(M > 1.6)$
2. The 60th percentile of M

Answer

1. 0.2071
2. 1.531

The Continuity Correction

A slight technical problem arises when the sampling distribution is discrete. In this case, the partial sum also has a discrete distribution, and hence we are approximating a discrete distribution with a continuous one. Suppose that X takes integer values (the most common case) and hence so does the partial sum Y_n . For any $k \in \mathbb{Z}$ and $h \in [0, 1]$, note that the event $\{k - h \leq Y_n \leq k + h\}$ is equivalent to the event $\{Y = k\}$. Different values of h lead to different normal approximations, even though the events are equivalent. The smallest approximation would be 0 when $h = 0$, and the approximations increase as h increases. It is customary to split the difference by using $h = \frac{1}{2}$ for the normal approximation. This is sometimes called the *half-*

unit continuity correction or the histogram correction. The continuity correction is extended to other events in the natural way, using the additivity of probability.

Suppose that $j, k \in \mathbb{Z}$ with $j \leq k$.

1. For the event $\{j \leq Y_n \leq k\} = \{j-1 < Y_n < k+1\}$, use $\{j - \frac{1}{2} \leq Y_n \leq k + \frac{1}{2}\}$ in the normal approximation.
2. For the event $\{j \leq Y_n\} = \{j-1 < Y_n\}$, use $\{j - \frac{1}{2} \leq Y_n\}$ in the normal approximation.
3. For the event $\{Y_n \leq k\} = \{Y_n < k+1\}$, use $\{Y_n \leq k + \frac{1}{2}\}$ in the normal approximation.

Let Y denote the sum of the scores of 20 fair dice. Compute the normal approximation to $\mathbb{P}(60 \leq Y \leq 75)$.

Answer

0.6741

In the dice experiment, set the die distribution to fair, select the sum random variable Y , and set $n = 20$. Run the simulation 1000 times and find each of the following. Compare with the result in the previous exercise:

1. $\mathbb{P}(60 \leq Y \leq 75)$
2. The relative frequency of the event $\{60 \leq Y \leq 75\}$ (from the simulation)

Normal Approximation to the Gamma Distribution

Recall that the *gamma distribution* with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ is a continuous distribution on $(0, \infty)$ with probability density function f given by

$$f(x) = \frac{1}{\Gamma(k)b^k} x^{k-1} e^{-x/b}, \quad x \in (0, \infty) \quad (6.4.14)$$

The mean is kb and the variance is kb^2 . The gamma distribution is widely used to model random times (particularly in the context of the Poisson model) and other positive random variables. The general gamma distribution is studied in more detail in the chapter on Special Distributions. In the context of the Poisson model (where $k \in \mathbb{N}_+$), the gamma distribution is also known as the *Erlang distribution*, named for Agner Erlang; it is studied in more detail in the chapter on the Poisson Process. Suppose now that Y_k has the gamma (Erlang) distribution with shape parameter $k \in \mathbb{N}_+$ and scale parameter $b > 0$ then

$$Y_k = \sum_{i=1}^k X_i \quad (6.4.15)$$

where (X_1, X_2, \dots) is a sequence of independent variables, each having the exponential distribution with scale parameter b . (The exponential distribution is a special case of the gamma distribution with shape parameter 1.) It follows that if k is large, the gamma distribution can be approximated by the normal distribution with mean kb and variance kb^2 . The same statement actually holds when k is not an integer. Here is the precise statement:

Suppose that Y_k has the gamma distribution with scale parameter $b \in (0, \infty)$ and shape parameter $k \in (0, \infty)$. Then the distribution of the standardized variable Z_k below converges to the standard normal distribution as $k \rightarrow \infty$:

$$Z_k = \frac{Y_k - kb}{\sqrt{kb}} \quad (6.4.16)$$

In the special distribution simulator, select the gamma distribution. Vary k and b and note the shape of the probability density function. With $k = 10$ and various values of b , run the experiment 1000 times and compare the empirical density function to the true probability density function.

Suppose that Y has the gamma distribution with shape parameter $k = 10$ and scale parameter $b = 2$. Find normal approximations to each of the following:

1. $\mathbb{P}(18 \leq Y \leq 23)$
2. The 80th percentile of Y

Answer

1. 0.3063
2. 25.32

Normal Approximation to the Chi-Square Distribution

Recall that the *chi-square distribution* with $n \in (0, \infty)$ degrees of freedom is a special case of the gamma distribution, with shape parameter $k = n/2$ and scale parameter $b = 2$. Thus, the chi-square distribution with n degrees of freedom has probability density function

$$f(x) = \frac{1}{\Gamma(n/2)2^{n/2}} x^{n/2-1} e^{-x/2}, \quad 0 < x < \infty \quad (6.4.17)$$

When n is a positive, integer, the chi-square distribution governs the sum of n independent, standard normal variables. For this reason, it is one of the most important distributions in statistics. The chi-square distribution is studied in more detail in the chapter on Special Distributions. From the [previous discussion](#), it follows that if n is large, the chi-square distribution can be approximated by the normal distribution with mean n and variance $2n$. Here is the precise statement:

Suppose that Y_n has the chi-square distribution with $n \in (0, \infty)$ degrees of freedom. Then the distribution of the standardized variable Z_n below converges to the standard normal distribution as $n \rightarrow \infty$:

$$Z_n = \frac{Y_n - n}{\sqrt{2n}} \quad (6.4.18)$$

In the special distribution simulator, select the chi-square distribution. Vary n and note the shape of the probability density function. With $n = 20$, run the experiment 1000 times and compare the empirical density function to the probability density function.

Suppose that Y has the chi-square distribution with $n = 20$ degrees of freedom. Find normal approximations to each of the following:

1. $\mathbb{P}(18 < Y < 25)$
2. The 75th percentile of Y

Answer

1. 0.4107
2. 24.3

Normal Approximation to the Binomial Distribution

Recall that a *Bernoulli trials sequence*, named for Jacob Bernoulli, is a sequence (X_1, X_2, \dots) of independent, identically distributed indicator variables with $\mathbb{P}(X_i = 1) = p$ for each i , where $p \in (0, 1)$ is the parameter. In the usual language of reliability, X_i is the outcome of trial i , where 1 means success and 0 means failure. The common mean is p and the common variance is $p(1 - p)$.

Let $Y_n = \sum_{i=1}^n X_i$, so that Y_n is the number of successes in the first n trials. Recall that Y_n has the *binomial distribution* with parameters n and p , and has probability density function

$$f(k) = \binom{n}{k} p^k (1 - p)^{n-k}, \quad k \in \{0, 1, \dots, n\} \quad (6.4.19)$$

The binomial distribution is studied in more detail in the chapter on Bernoulli trials.

It follows from the central limit theorem that if n is large, the binomial distribution with parameters n and p can be approximated by the normal distribution with mean np and variance $np(1 - p)$. The rule of thumb is that n should be large enough for $np \geq 5$ and $n(1 - p) \geq 5$. (The first condition is the important one when $p < \frac{1}{2}$ and the second condition is the important one when $p > \frac{1}{2}$.) Here is the precise statement:

Suppose that Y_n has the binomial distribution with trial parameter $n \in \mathbb{N}_+$ and success parameter $p \in (0, 1)$. Then the distribution of the standardized variable Z_n given below converges to the standard normal distribution as $n \rightarrow \infty$:

$$Z_n = \frac{Y_n - np}{\sqrt{np(1-p)}} \quad (6.4.20)$$

In the binomial timeline experiment, vary n and p and note the shape of the probability density function. With $n = 50$ and $p = 0.3$, run the simulation 1000 times and compute the following:

1. $\mathbb{P}(12 \leq Y \leq 16)$
2. The relative frequency of the event $\{12 \leq Y \leq 16\}$ (from the simulation)

Answer

1. 0.5448

Suppose that Y has the binomial distribution with parameters $n = 50$ and $p = 0.3$. Compute the normal approximation to $\mathbb{P}(12 \leq Y \leq 16)$ (don't forget the continuity correction) and compare with the results of the previous exercise.

Answer

0.5383

Normal Approximation to the Poisson Distribution

Recall that the *Poisson distribution*, named for Simeon Poisson, is a discrete distribution on \mathbb{N} with probability density function f given by

$$f(x) = e^{-\theta} \frac{\theta^x}{x!}, \quad x \in \mathbb{N} \quad (6.4.21)$$

where $\theta > 0$ is a parameter. The parameter is both the mean and the variance of the distribution. The Poisson distribution is widely used to model the number of “random points” in a region of time or space, and is studied in more detail in the chapter on the Poisson Process. In this context, the parameter is proportional to the size of the region.

Suppose now that Y_n has the Poisson distribution with parameter $n \in \mathbb{N}_+$. Then

$$Y_n = \sum_{i=1}^n X_i \quad (6.4.22)$$

where (X_1, X_2, \dots, X_n) is a sequence of independent variables, each with the Poisson distribution with parameter 1. It follows from the central limit theorem that if n is large, the Poisson distribution with parameter n can be approximated by the normal distribution with mean n and variance n . The same statement holds when the parameter n is not an integer. Here is the precise statement:

. Suppose that Y_θ has the Poisson distribution with parameter $\theta \in (0, \infty)$. Then the distribution of the standardized variable Z_θ below converges to the standard normal distribution as $\theta \rightarrow \infty$:

$$Z_\theta = \frac{Y_\theta - \theta}{\sqrt{\theta}} \quad (6.4.23)$$

Suppose that Y has the Poisson distribution with mean 20.

1. Compute the true value of $\mathbb{P}(16 \leq Y \leq 23)$.
2. Compute the normal approximation to $\mathbb{P}(16 \leq Y \leq 23)$.

Answer

1. 0.6310
2. 0.6259

In the Poisson experiment, vary the time and rate parameters t and r (the parameter of the Poisson distribution in the experiment is the product rt). Note the shape of the probability density function. With $r = 5$ and $t = 4$, run the experiment 1000 times and compare the empirical density function to the true probability density function.

Normal Approximation to the Negative Binomial Distribution

The general version of the *negative binomial distribution* is a discrete distribution on \mathbb{N} , with shape parameter $k \in (0, \infty)$ and success parameter $p \in (0, 1)$. The probability density function f is given by

$$f(n) = \binom{n+k-1}{n} p^k (1-p)^n, \quad n \in \mathbb{N}_+ \quad (6.4.24)$$

The mean is $k(1-p)/p$ and the variance is $k(1-p)/p^2$. The negative binomial distribution is studied in more detail in the chapter on Bernoulli trials. If $k \in \mathbb{N}_+$, the distribution governs the number of failures Y_k before success number k in a sequence of Bernoulli trials with success parameter p . Thus in this case,

$$Y_k = \sum_{i=1}^k X_i \quad (6.4.25)$$

where (X_1, X_2, \dots, X_k) is a sequence of independent variables, each having the geometric distribution on \mathbb{N} with parameter p . (The geometric distribution is a special case of the negative binomial, with parameters 1 and p .) In the context of the Bernoulli trials, X_1 is the number of failures before the first success, and for $i \in \{2, 3, \dots\}$, X_i is the number of failures between success number $i-1$ and success number i . It follows that if k is large, the negative binomial distribution can be approximated by the normal distribution. The same statement holds if k is not an integer. Here is the precise statement:

Suppose that Y_k has the negative binomial distribution with shape parameter $k \in (0, 1)$ and scale parameter $p \in (0, 1)$. Then the distribution of the standardized variable Z_k below converges to the standard normal distribution as $k \rightarrow \infty$:

$$Z_k = \frac{pY_k - k(1-p)}{\sqrt{k(1-p)}} \quad (6.4.26)$$

Another version of the negative binomial distribution is the distribution of the trial number V_k of success number $k \in \mathbb{N}_+$. So $V_k = k + Y_k$ and V_k has mean k/p and variance $k(1-p)/p^2$. The normal approximation applies to the distribution of V_k as well, if k is large, and since the distributions are related by a location transformation, the standard scores are the same. That is

$$\frac{pV_k - k}{\sqrt{k(1-p)}} = \frac{pY_k - k(1-p)}{\sqrt{k(1-p)}} \quad (6.4.27)$$

In the negative binomial experiment, vary k and p and note the shape of the probability density function. With $k = 5$ and $p = 0.4$, run the experiment 1000 times and compare the empirical density function to the true probability density function.

Suppose that Y has the negative binomial distribution with trial parameter $k = 10$ and success parameter $p = 0.4$. Find normal approximations to each of the following:

1. $\mathbb{P}(20 < Y < 30)$
2. The 80th percentile of Y

Answer

1. 0.6318
2. 30.1

Partial Sums with a Random Number of Terms

Our last topic is a bit more esoteric, but still fits with the general setting of this section. Recall that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent, identically distributed real-valued random variables with common mean μ and variance σ^2 . Suppose now that N is a random variable (on the same probability space) taking values in \mathbb{N} , also with finite mean and variance. Then

$$Y_N = \sum_{i=1}^N X_i \quad (6.4.28)$$

is a random sum of the independent, identically distributed variables. That is, the terms are random of course, but so also is the number of terms N . We are primarily interested in the moments of Y_N .

Independent Number of Terms

Suppose first that N , the number of terms, is independent of \mathbf{X} , the sequence of terms. Computing the moments of Y_N is a good exercise in conditional expectation.

The conditional expected value of Y_N given N , and the expected value of Y_N are

1. $\mathbb{E}(Y_N | N) = N\mu$
2. $\mathbb{E}(Y_N) = \mathbb{E}(N)\mu$

The conditional variance of Y_N given N and the variance of Y_N are

1. $\text{var}(Y_N | N) = N\sigma^2$
2. $\text{var}(Y_N) = \mathbb{E}(N)\sigma^2 + \text{var}(N)\mu^2$

Let H denote the probability generating function of N . Show that the moment generating function of Y_N is $H \circ G$.

1. $\mathbb{E}(e^{tY_N} | N) = [G(t)]^N$
2. $\mathbb{E}(e^{tY_N}) = H(G(t))$

Wald's Equation

The result in [Exercise 29 \(b\)](#) generalizes to the case where the random number of terms N is a *stopping time* for the sequence \mathbf{X} . This means that the event $\{N = n\}$ depends only on (technically, is measurable with respect to) (X_1, X_2, \dots, X_n) for each $n \in \mathbb{N}$. The generalization is known as *Wald's equation*, and is named for Abraham Wald. Stopping times are studied in much more technical detail in the section on Filtrations and Stopping Times.

If N is a stopping time for \mathbf{X} then $\mathbb{E}(Y_N) = \mathbb{E}(N)\mu$.

Proof

First note that $Y_N = \sum_{i=1}^{\infty} X_i \mathbf{1}(i \leq N)$. But $\{i \leq N\} = \{N \geq i\}^c$ depends only on $\{X_1, \dots, X_{i-1}\}$ and hence is independent of X_i . Thus $\mathbb{E}[X_i \mathbf{1}(i \leq N)] = \mu \mathbb{P}(N \geq i)$. Suppose that $X_i \geq 0$ for each i . Taking expected values term by term gives Wald's equation in this special case. The interchange of sum and expected value is justified by the monotone convergence theorem. Now Wald's equation can be established in general by using the dominated convergence theorem.

An elegant proof of Wald's equation is given in the chapter on Martingales.

Suppose that the number of customers arriving at a store during a given day has the Poisson distribution with parameter 50. Each customer, independently of the others (and independently of the number of customers), spends an amount of money that is uniformly distributed on the interval $[0, 20]$. Find the mean and standard deviation of the amount of money that the store takes in during a day.

Answer

500, 81.65

When a certain critical component in a system fails, it is immediately replaced by a new, statistically identical component. The components are independent, and the lifetime of each (in hours) is exponentially distributed with scale parameter b . During the life of the system, the number of critical components used has a geometric distribution on \mathbb{N}_+ with parameter p . For the total life of the critical component,

1. Find the mean.
2. Find the standard deviation.
3. Find the moment generating function.

4. Identify the distribution by name.

Answer

1. b/p

2. b/p

3. $t \mapsto \frac{1}{1-(b/p)t}$

4. Exponential distribution with scale parameter b/p

This page titled [6.4: The Central Limit Theorem](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.5: The Sample Variance

Descriptive Theory

Recall the basic model of statistics: we have a population of objects of interest, and we have various measurements (variables) that we make on these objects. We select objects from the population and record the variables for the objects in the sample; these become our data. Once again, our first discussion is from a descriptive point of view. That is, we do not assume that the data are generated by an underlying probability distribution. Remember however, that the data themselves form a probability distribution.

Variance and Standard Deviation

Suppose that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a real-valued variable x . Recall that the sample mean is

$$m = \frac{1}{n} \sum_{i=1}^n x_i \quad (6.5.1)$$

and is the most important measure of the center of the data set. The *sample variance* is defined to be

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - m)^2 \quad (6.5.2)$$

If we need to indicate the dependence on the data vector \mathbf{x} , we write $s^2(\mathbf{x})$. The difference $x_i - m$ is the *deviation* of x_i from the mean m of the data set. Thus, the variance is the *mean square deviation* and is a measure of the spread of the data set with respect to the mean. The reason for dividing by $n-1$ rather than n is best understood in terms of the inferential point of view that we discuss in the next section; this definition makes the sample variance an unbiased estimator of the distribution variance. However, the reason for the averaging can also be understood in terms of a related concept.

$$\sum_{i=1}^n (x_i - m) = 0.$$

Proof

$$\sum_{i=1}^n (x_i - m) = \sum_{i=1}^n x_i - \sum_{i=1}^n m = nm - nm = 0.$$

Thus, if we know $n-1$ of the deviations, we can compute the last one. This means that there are only $n-1$ freely varying deviations, that is to say, $n-1$ *degrees of freedom* in the set of deviations. In the definition of sample variance, we average the squared deviations, not by dividing by the number of terms, but rather by dividing by the number of degrees of freedom in those terms. However, this argument notwithstanding, it would be reasonable, *from a purely descriptive point of view*, to divide by n in the definition of the sample variance. Moreover, when n is sufficiently large, it hardly matters whether we divide by n or by $n-1$.

In any event, the square root s of the sample variance s^2 is the sample standard deviation. It is the *root mean square deviation* and is also a measure of the spread of the data with respect to the mean. Both measures of spread are important. Variance has nicer mathematical properties, but its physical unit is the square of the unit of x . For example, if the underlying variable x is the height of a person in inches, the variance is in square inches. On the other hand, the standard deviation has the same physical unit as the original variable, but its mathematical properties are not as nice.

Recall that the data set \mathbf{x} naturally gives rise to a probability distribution, namely the *empirical distribution* that places probability $\frac{1}{n}$ at x_i for each i . Thus, if the data are distinct, this is the uniform distribution on $\{x_1, x_2, \dots, x_n\}$. The sample mean m is simply the expected value of the empirical distribution. Similarly, if we were to divide by n rather than $n-1$, the sample variance would be the variance of the empirical distribution. Most of the properties and results this section follow from much more general properties and results for the variance of a probability distribution (although for the most part, we give independent proofs).

Measures of Center and Spread

Measures of center and measures of spread are best thought of together, in the context of an *error function*. The error function measures how well a single number a represents the entire data set \mathbf{x} . The values of a (if they exist) that minimize the error functions are our measures of center; the minimum value of the error function is the corresponding measure of spread. Of course, we hope for a *single* value of a that minimizes the error function, so that we have a unique measure of center.

Let's apply this procedure to the *mean square error function* defined by

$$\text{mse}(a) = \frac{1}{n-1} \sum_{i=1}^n (x_i - a)^2, \quad a \in \mathbb{R} \quad (6.5.3)$$

Minimizing mse is a standard problem in calculus.

The graph of mse is a parabola opening upward.

1. mse is minimized when $a = m$, the sample mean.
2. The minimum value of mse is s^2 , the sample variance.

Proof

We can tell from the form of mse that the graph is a parabola opening upward. Taking the derivative gives

$$\frac{d}{da} \text{mse}(a) = -\frac{2}{n-1} \sum_{i=1}^n (x_i - a) = -\frac{2}{n-1} (nm - na) \quad (6.5.4)$$

Hence $a = m$ is the unique value that minimizes mse. Of course, $\text{mse}(m) = s^2$.

Trivially, if we defined the mean square error function by dividing by n rather than $n-1$, then the minimum value would still occur at m , the sample mean, but the minimum value would be the alternate version of the sample variance in which we divide by n . On the other hand, if we were to use the *root mean square deviation* function $\text{rmse}(a) = \sqrt{\text{mse}(a)}$, then because the square root function is strictly increasing on $[0, \infty)$, the minimum value would again occur at m , the sample mean, but the minimum value would be s , the sample standard deviation. The important point is that with all of these error functions, the unique measure of center is the sample mean, and the corresponding measures of spread are the various ones that we are studying.

Next, let's apply our procedure to the *mean absolute error function* defined by

$$\text{mae}(a) = \frac{1}{n-1} \sum_{i=1}^n |x_i - a|, \quad a \in \mathbb{R} \quad (6.5.5)$$

The mean absolute error function satisfies the following properties:

1. mae is a continuous function.
2. The graph of mae consists of lines.
3. The slope of the line at a depends on where a is in the data set \mathbf{x} .

Proof

For parts (a) and (b), note that for each i , $|x_i - a|$ is a continuous function of a with the graph consisting of two lines (of slopes ± 1) meeting at x_i .

Mathematically, mae has some problems as an error function. First, the function will not be smooth (differentiable) at points where two lines of different slopes meet. More importantly, the values that minimize mae may occupy an entire interval, thus leaving us without a unique measure of center. The [error function exercises](#) below will show you that these pathologies can really happen. It turns out that mae is minimized at any point in the median interval of the data set \mathbf{x} . The proof of this result follows from a much more general result for probability distributions. Thus, the medians are the natural measures of center associated with mae as a measure of error, in the same way that the sample mean is the measure of center associated with the mse as a measure of error.

Properties

In this section, we establish some essential properties of the sample variance and standard deviation. First, the following alternate formula for the sample variance is better for computational purposes, and for certain theoretical purposes as well.

The sample variance can be computed as

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n x_i^2 - \frac{n}{n-1} m^2 \quad (6.5.6)$$

Proof

Note that

$$\sum_{i=1}^n (x_i - m)^2 = \sum_{i=1}^n (x_i^2 - 2mx_i + m^2) = \sum_{i=1}^n x_i^2 - 2m \sum_{i=1}^n x_i + \sum_{i=1}^n m^2 \quad (6.5.7)$$

$$= \sum_{i=1}^n x_i^2 - 2nm^2 + nm^2 = \sum_{i=1}^n x_i^2 - nm^2 \quad (6.5.8)$$

Dividing by $n - 1$ gives the result.

If we let $\mathbf{x}^2 = (x_1^2, x_2^2, \dots, x_n^2)$ denote the sample from the variable x^2 , then the computational formula in the last exercise can be written succinctly as

$$s^2(\mathbf{x}) = \frac{n}{n-1} [m(\mathbf{x}^2) - m^2(\mathbf{x})] \quad (6.5.9)$$

The following theorem gives another computational formula for the sample variance, directly in terms of the variables and thus without the computation of an intermediate statistic.

The sample variance can be computed as

$$s^2 = \frac{1}{2n(n-1)} \sum_{i=1}^n \sum_{j=1}^n (x_i - x_j)^2 \quad (6.5.10)$$

Proof

Note that

$$\frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n (x_i - x_j)^2 = \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n (x_i - m + m - x_j)^2 \quad (6.5.11)$$

$$= \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n [(x_i - m)^2 + 2(x_i - m)(m - x_j) + (m - x_j)^2] \quad (6.5.12)$$

$$= \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n (x_i - m)^2 + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (x_i - m)(m - x_j) + \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n (m - x_j)^2 \quad (6.5.13)$$

$$= \frac{1}{2} \sum_{i=1}^n (x_i - m)^2 + 0 + \frac{1}{2} \sum_{j=1}^n (m - x_j)^2 \quad (6.5.14)$$

$$= \sum_{i=1}^n (x_i - m)^2 \quad (6.5.15)$$

Dividing by $n - 1$ gives the result.

The sample variance is nonnegative:

1. $s^2 \geq 0$
2. $s^2 = 0$ if and only if $x_i = x_j$ for each $i, j \in \{1, 2, \dots, n\}$.

Proof

Part (a) is obvious. For part (b) note that if $s^2 = 0$ then $x_i = m$ for each i . Conversely, if \mathbf{x} is a constant vector, then m is that same constant.

Thus, $s^2 = 0$ if and only if the data set is constant (and then, of course, the mean is the common value).

If c is a constant then

1. $s^2(c\mathbf{x}) = c^2 s^2(\mathbf{x})$
2. $s(c\mathbf{x}) = |c| s(\mathbf{x})$

Proof

For part (a), recall that $m(c\mathbf{x}) = cm(\mathbf{x})$. Hence

$$s^2(c\mathbf{x}) = \frac{1}{n-1} \sum_{i=1}^n [cx_i - cm(\mathbf{x})]^2 = \frac{1}{n-1} \sum_{i=1}^n c^2 [x_i - m(\mathbf{x})]^2 = c^2 s^2(\mathbf{x}) \quad (6.5.16)$$

If \mathbf{c} is a sample of size n from a constant c then

1. $s^2(\mathbf{x} + \mathbf{c}) = s^2(\mathbf{x})$.
2. $s(\mathbf{x} + \mathbf{c}) = s(\mathbf{x})$

Proof

Recall that $m(\mathbf{x} + \mathbf{c}) = m(\mathbf{x}) + c$. Hence

$$s^2(\mathbf{x} + \mathbf{c}) = \frac{1}{n-1} \sum_{i=1}^n \{(x_i + c) - [m(\mathbf{x}) + c]\}^2 = \frac{1}{n-1} \sum_{i=1}^n [x_i - m(\mathbf{x})]^2 = s^2(\mathbf{x}) \quad (6.5.17)$$

As a special case of these results, suppose that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n corresponding to a real variable x , and that a and b are constants. The sample corresponding to the variable $y = a + bx$, in our vector notation, is $\mathbf{a} + b\mathbf{x}$. Then $m(\mathbf{a} + b\mathbf{x}) = a + bm(\mathbf{x})$ and $s(\mathbf{a} + b\mathbf{x}) = |b|s(\mathbf{x})$. Linear transformations of this type, when $b > 0$, arise frequently when physical units are changed. In this case, the transformation is often called a *location-scale* transformation; a is the location parameter and b is the scale parameter. For example, if x is the length of an object in inches, then $y = 2.54x$ is the length of the object in centimeters. If x is the temperature of an object in degrees Fahrenheit, then $y = \frac{5}{9}(x - 32)$ is the temperature of the object in degree Celsius.

Now, for $i \in \{1, 2, \dots, n\}$, let $z_i = (x_i - m)/s$. The number z_i is the *standard score* associated with x_i . Note that since x_i , m , and s have the same physical units, the standard score z_i is *dimensionless* (that is, has no physical units); it measures the directed distance from the mean m to the data value x_i in standard deviations.

The sample of standard scores $\mathbf{z} = (z_1, z_2, \dots, z_n)$ has mean 0 and variance 1. That is,

1. $m(\mathbf{z}) = 0$
2. $s^2(\mathbf{z}) = 1$

Proof

These results follow from Theorems 7 and 8. In vector notation, note that $\mathbf{z} = (\mathbf{x} - \mathbf{m})/s$. Hence $m(\mathbf{z}) = (m - m)/s = 0$ and $s(\mathbf{z}) = s/s = 1$.

Approximating the Variance

Suppose that instead of the actual data \mathbf{x} , we have a frequency distribution corresponding to a partition with classes (intervals) (A_1, A_2, \dots, A_k) , class marks (midpoints of the intervals) (t_1, t_2, \dots, t_k) , and frequencies (n_1, n_2, \dots, n_k) . Recall that the relative frequency of class A_j is $p_j = n_j/n$. In this case, approximate values of the sample mean and variance are, respectively,

$$m = \frac{1}{n} \sum_{j=1}^k n_j t_j = \sum_{j=1}^k p_j t_j \quad (6.5.18)$$

$$s^2 = \frac{1}{n-1} \sum_{j=1}^k n_j (t_j - m)^2 = \frac{n}{n-1} \sum_{j=1}^k p_j (t_j - m)^2 \quad (6.5.19)$$

These approximations are based on the hope that the data values in each class are well represented by the class mark. In fact, these are the standard definitions of sample mean and variance for the data set in which t_j occurs n_j times for each j .

Inferential Statistics

We continue our discussion of the sample variance, but now we assume that the variables are random. Thus, suppose that we have a basic random experiment, and that X is a real-valued random variable for the experiment with mean μ and standard deviation σ . We will need some higher order moments as well. Let $\sigma_3 = \mathbb{E}[(X - \mu)^3]$ and $\sigma_4 = \mathbb{E}[(X - \mu)^4]$ denote the 3rd and 4th moments about the mean. Recall that $\sigma_3/\sigma^3 = \text{skew}(X)$, the skewness of X , and $\sigma_4/\sigma^4 = \text{kurt}(X)$, the kurtosis of X . We assume that $\sigma_4 < \infty$.

We repeat the basic experiment n times to form a new, compound experiment, with a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots, X_n)$, each with the same distribution as X . In statistical terms, \mathbf{X} is a random sample of size n from the distribution of X . All of the statistics above make sense for \mathbf{X} , of course, but now these statistics are random variables. We will use the same notation, except for the usual convention of denoting random variables by capital letters. Finally, note that the deterministic properties and relations established above still hold.

In addition to being a measure of the center of the data \mathbf{X} , the sample mean

$$M = \frac{1}{n} \sum_{i=1}^n X_i \quad (6.5.20)$$

is a natural estimator of the distribution mean μ . In this section, we will derive statistics that are natural estimators of the distribution variance σ^2 . The statistics that we will derive are different, depending on whether μ is known or unknown; for this reason, μ is referred to as a *nuisance parameter* for the problem of estimating σ^2 .

A Special Sample Variance

First we will assume that μ is known. Although this is almost always an artificial assumption, it is a nice place to start because the analysis is relatively easy and will give us insight for the standard case. A natural estimator of σ^2 is the following statistic, which we will refer to as the *special sample variance*.

$$W^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 \quad (6.5.21)$$

W^2 is the sample mean for a random sample of size n from the distribution of $(X - \mu)^2$, and satisfies the following properties:

1. $\mathbb{E}(W^2) = \sigma^2$
2. $\text{var}(W^2) = \frac{1}{n}(\sigma_4 - \sigma^4)$
3. $W^2 \rightarrow \sigma^2$ as $n \rightarrow \infty$ with probability 1
4. The distribution of $\sqrt{n}(W^2 - \sigma^2) / \sqrt{\sigma_4 - \sigma^4}$ converges to the standard normal distribution as $n \rightarrow \infty$.

Proof

These result follow immediately from standard results in the section on the Law of Large Numbers and the section on the Central Limit Theorem. For part (b), note that

$$\text{var}[(X - \mu)^2] = \mathbb{E}[(X - \mu)^4] - (\mathbb{E}[(X - \mu)^2])^2 = \sigma_4 - \sigma^4 \quad (6.5.22)$$

In particular part (a) means that W^2 is an *unbiased* estimator of σ^2 . From part (b), note that $\text{var}(W^2) \rightarrow 0$ as $n \rightarrow \infty$; this means that W^2 is a *consistent* estimator of σ^2 . The square root of the special sample variance is a special version of the *sample standard deviation*, denoted W .

$\mathbb{E}(W) \leq \sigma$. Thus, W is a *negatively biased* estimator that tends to underestimate σ .

Proof

This follows from the unbiased property and Jensen's inequality. Since $w \mapsto \sqrt{w}$ is concave downward on $[0, \infty)$, we have $\mathbb{E}(W) = \mathbb{E}(\sqrt{W^2}) \leq \sqrt{\mathbb{E}(W^2)} = \sqrt{\sigma^2} = \sigma$.

Next we compute the covariance and correlation between the sample mean and the special sample variance.

The covariance and correlation of M and W^2 are

1. $\text{cov}(M, W^2) = \sigma_3/n$.
2. $\text{cor}(M, W^2) = \sigma^3 / \sqrt{\sigma^2(\sigma_4 - \sigma^4)}$

Proof

1. From the bilinearity of the covariance operator and by independence,

$$\text{cov}(M, W^2) = \text{cov}\left[\frac{1}{n} \sum_{i=1}^n X_i, \frac{1}{n} \sum_{j=1}^n (X_j - \mu)^2\right] = \frac{1}{n^2} \sum_{i=1}^n \text{cov}[X_i, (X_i - \mu)^2] \quad (6.5.23)$$

But $\text{cov}[X_i, (X_i - \mu)^2] = \text{cov}[X_i - \mu, (X_i - \mu)^2] = \mathbb{E}[(X_i - \mu)^3] - \mathbb{E}(X_i - \mu)\mathbb{E}[(X_i - \mu)^2] = \sigma_3$. Substituting gives the result.

2. This follows from part (a), the unbiased property, and our previous result that $\text{var}(M) = \sigma^2/n$.

Note that the correlation does not depend on the sample size, and that the sample mean and the special sample variance are uncorrelated if $\sigma_3 = 0$ (equivalently $\text{skew}(X) = 0$).

The Standard Sample Variance

Consider now the more realistic case in which μ is unknown. In this case, a natural approach is to average, in some sense, the squared deviations $(X_i - M)^2$ over $i \in \{1, 2, \dots, n\}$. It might seem that we should average by dividing by n . However, another approach is to divide by whatever constant would give us an unbiased estimator of σ^2 . This constant turns out to be $n - 1$, leading to the *standard sample variance*:

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M)^2 \quad (6.5.24)$$

$\mathbb{E}(S^2) = \sigma^2$.

Proof

By expanding (as was shown in the last section),

$$\sum_{i=1}^n (X_i - M)^2 = \sum_{i=1}^n X_i^2 - nM^2 \quad (6.5.25)$$

Recall that $\mathbb{E}(M) = \mu$ and $\text{var}(M) = \sigma^2/n$. Taking expected values in the displayed equation gives

$$\mathbb{E}\left(\sum_{i=1}^n (X_i - M)^2\right) = \sum_{i=1}^n (\sigma^2 + \mu^2) - n\left(\frac{\sigma^2}{n} + \mu^2\right) = n(\sigma^2 + \mu^2) - n\left(\frac{\sigma^2}{n} + \mu^2\right) = (n-1)\sigma^2 \quad (6.5.26)$$

Of course, the square root of the sample variance is the *sample standard deviation*, denoted S .

$\mathbb{E}(S) \leq \sigma$. Thus, S is a negatively biased estimator than tends to underestimate σ .

Proof

The proof is exactly the same as for the [special standard variance](#).

$S^2 \rightarrow \sigma^2$ as $n \rightarrow \infty$ with probability 1.

Proof

This follows from the strong law of large numbers. Recall again that

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n X_i^2 - \frac{n}{n-1} M^2 = \frac{n}{n-1} [M(\mathbf{X}^2) - M^2(\mathbf{X})] \quad (6.5.27)$$

But with probability 1, $M(\mathbf{X}^2) \rightarrow \sigma^2 + \mu^2$ as $n \rightarrow \infty$ and $M^2(\mathbf{X}) \rightarrow \mu^2$ as $n \rightarrow \infty$.

Since S^2 is an unbiased estimator of σ^2 , the variance of S^2 is the *mean square error*, a measure of the quality of the estimator.

$$\text{var}(S^2) = \frac{1}{n} \left(\sigma_4 - \frac{n-3}{n-1} \sigma^4 \right).$$

Proof

Recall from the [result above](#) that

$$S^2 = \frac{1}{2n(n-1)} \sum_{i=1}^n \sum_{j=1}^n (X_i - X_j)^2 \quad (6.5.28)$$

Hence, using the bilinear property of covariance we have

$$\text{var}(S^2) = \text{cov}(S^2, S^2) = \frac{1}{4n^2(n-1)^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{l=1}^n \text{cov}[(X_i - X_j)^2, (X_k - X_l)^2] \quad (6.5.29)$$

We compute the covariances in this sum by considering disjoint cases:

- $\text{cov}[(X_i - X_j)^2, (X_k - X_l)^2] = 0$ if $i = j$ or $k = l$, and there are $2n^3 - n^2$ such terms.
- $\text{cov}[(X_i - X_j)^2, (X_k - X_l)^2] = 0$ if i, j, k, l are distinct, and there are $n(n-1)(n-2)(n-3)$ such terms.
- $\text{cov}[(X_i - X_j)^2, (X_k - X_l)^2] = 2\sigma_4 + 2\sigma^4$ if $i \neq j$ and $\{k, l\} = \{i, j\}$, and there are $2n(n-1)$ such terms.
- $\text{cov}[(X_i - X_j)^2, (X_k - X_l)^2] = \sigma_4 - \sigma^4$ if $i \neq j, k \neq l$ and $\#\{i, j\} \cap \{k, l\} = 1$, and there are $4n(n-1)(n-2)$ such terms.

Substituting gives the result.

Note that $\text{var}(S^2) \rightarrow 0$ as $n \rightarrow \infty$, and hence S^2 is a consistent estimator of σ^2 . On the other hand, it's not surprising that the variance of the standard sample variance (where we assume that μ is unknown) is greater than the variance of the special standard variance (in which we assume μ is known).

$$\text{var}(S^2) > \text{var}(W^2).$$

Proof

From the [formula above](#) for the variance of W^2 , the previous result for the variance of S^2 , and simple algebra,

$$\text{var}(S^2) - \text{var}(W^2) = \frac{2}{n(n-1)} \sigma^4 \quad (6.5.30)$$

Note however that the difference goes to 0 as $n \rightarrow \infty$.

Next we compute the covariance between the sample mean and the sample variance.

The covariance and correlation between the sample mean and sample variance are

1. $\text{cov}(M, S^2) = \sigma_3/n$
2. $\text{cor}(M, S^2) = \frac{\sigma_3}{\sigma \sqrt{\sigma_4 - \sigma^4(n-3)/(n-1)}}$

Proof

1. Recall again that

$$M = \frac{1}{n} \sum_{i=1}^n X_i, \quad S^2 = \frac{1}{2n(n-1)} \sum_{j=1}^n \sum_{k=1}^n (X_j - X_k)^2 \quad (6.5.31)$$

Hence, using the bilinear property of covariance we have

$$\text{cov}(M, S^2) = \frac{1}{2n^2(n-1)} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \text{cov}[X_i, (X_j - X_k)^2] \quad (6.5.32)$$

We compute the covariances in this sum by considering disjoint cases:

- $\text{cov}[X_i, (X_j - X_k)^2] = 0$ if $j = k$, and there are n^2 such terms.
- $\text{cov}[X_i, (X_j - X_k)^2] = 0$ if i, j, k are distinct, and there are $n(n-1)(n-2)$ such terms.
- $\text{cov}[X_i, (X_j - X_k)^2] = \sigma_3$ if $j \neq k$ and $i \in \{j, k\}$, and there are $2n(n-1)$ such terms.

Substituting gives the result.

2. This follows from part(a), the result above on the variance of S^2 , and $\text{var}(M) = \sigma^2/n$.

In particular, note that $\text{cov}(M, S^2) = \text{cov}(M, W^2)$. Again, the sample mean and variance are uncorrelated if $\sigma_3 = 0$ so that $\text{skew}(X) = 0$. Our last result gives the covariance and correlation between the special sample variance and the standard one. Curiously, the covariance the same as the variance of the special sample variance.

The covariance and correlation between W^2 and S^2 are

$$\begin{aligned} 1. \text{cov}(W^2, S^2) &= (\sigma_4 - \sigma^4)/n \\ 2. \text{cor}(W^2, S^2) &= \sqrt{\frac{\sigma_4 - \sigma^4}{\sigma_4 - \sigma^4(n-3)/(n-1)}} \end{aligned}$$

Proof

1. Recall again that

$$W^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2, \quad S^2 = \frac{1}{2n(n-1)} \sum_{j=1}^n \sum_{k=1}^n (X_j - X_k)^2 \quad (6.5.33)$$

so by the bilinear property of covariance we have

$$\text{cov}(W^2, S^2) = \frac{1}{2n^2(n-1)} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \text{cov}[(X_i - \mu)^2, (X_j - X_k)^2] \quad (6.5.34)$$

Once again, we compute the covariances in this sum by considering disjoint cases:

- $\text{cov}[(X_i - \mu)^2, (X_j - X_k)^2] = 0$ if $j = k$, and there are n^2 such terms.
- $\text{cov}[(X_i - \mu)^2, (X_j - X_k)^2] = 0$ if i, j, k are distinct, and there are $n(n-1)(n-2)$ such terms.
- $\text{cov}[(X_i - \mu)^2, (X_j - X_k)^2] = \sigma_4 - \sigma^4$ if $j \neq k$ and $i \in \{j, k\}$, and there are $2n(n-1)$ such terms.

Substituting gives the results.

2. This follows from part (a) and the formulas above for the variance of W^2 and the variance of V^2

Note that $\text{cor}(W^2, S^2) \rightarrow 1$ as $n \rightarrow \infty$, not surprising since with probability 1, $S^2 \rightarrow \sigma^2$ and $W^2 \rightarrow \sigma^2$ as $n \rightarrow \infty$.

A particularly important special case occurs when the sampling distribution is normal. This case is explored in the section on Special Properties of Normal Samples.

Exercises

Basic Properties

Suppose that x is the temperature (in degrees Fahrenheit) for a certain type of electronic component after 10 hours of operation. A sample of 30 components has mean 113° and standard deviation 18° .

1. Classify x by type and level of measurement.
2. Find the sample mean and standard deviation if the temperature is converted to degrees Celsius. The transformation is $y = \frac{5}{9}(x - 32)$.

Answer

1. continuous, interval
2. $m = 45^\circ$, $s = 10^\circ$

Suppose that x is the length (in inches) of a machined part in a manufacturing process. A sample of 50 parts has mean 10.0 and standard deviation 2.0.

1. Classify x by type and level of measurement.
2. Find the sample mean if length is measured in centimeters. The transformation is $y = 2.54x$.

Answer

1. continuous, ratio
2. $m = 25.4$, $s = 5.08$

Professor Moriarity has a class of 25 students in her section of Stat 101 at Enormous State University (ESU). The mean grade on the first midterm exam was 64 (out of a possible 100 points) and the standard deviation was 16. Professor Moriarity thinks the grades are a bit low and is considering various transformations for increasing the grades. In each case below give the mean and standard deviation of the transformed grades, or state that there is not enough information.

1. Add 10 points to each grade, so the transformation is $y = x + 10$.
2. Multiply each grade by 1.2, so the transformation is $z = 1.2x$
3. Use the transformation $w = 10\sqrt{x}$. Note that this is a non-linear transformation that curves the grades greatly at the low end and very little at the high end. For example, a grade of 100 is still 100, but a grade of 36 is transformed to 60.

One of the students did not study at all, and received a 10 on the midterm. Professor Moriarity considers this score to be an outlier.

4. Find the mean and standard deviation if this score is omitted.

Answer

1. $m = 74, s = 16$
2. $m = 76.8, s = 19.2$
3. Not enough information
4. $m = 66.25, s = 11.62$

Computational Exercises

All statistical software packages will compute means, variances and standard deviations, draw dotplots and histograms, and in general perform the numerical and graphical procedures discussed in this section. For real statistical experiments, particularly those with large data sets, the use of statistical software is essential. On the other hand, there is some value in performing the computations by hand, with small, artificial data sets, in order to master the concepts and definitions. In this subsection, do the computations and draw the graphs with minimal technological aids.

Suppose that x is the number of math courses completed by an ESU student. A sample of 10 ESU students gives the data $\mathbf{x} = (3, 1, 2, 0, 2, 4, 3, 2, 1, 2)$

1. Classify x by type and level of measurement.
2. Sketch the dotplot.
3. Construct a table with rows corresponding to cases and columns corresponding to i , x_i , $x_i - m$, and $(x_i - m)^2$. Add rows at the bottom in the i column for totals and means.

Answer

1. discrete, ratio

3. i	x_i	$x_i - m$	$(x_i - m)^2$
1	3	1	1
2	1	-1	1
3	2	0	0
4	0	-2	4
5	2	0	0
6	4	2	4
7	3	1	1
8	2	0	0
9	1	-1	1
Total	20	0	14
Mean	2	0	14/9

i	x_i	$x_i - m$	$(x_i - m)^2$
10	2	0	0
Total	20	0	14
Mean	2	0	14/9

Suppose that a sample of size 12 from a discrete variable x has empirical density function given by $f(-2) = 1/12$, $f(-1) = 1/4$, $f(0) = 1/3$, $f(1) = 1/6$, $f(2) = 1/6$.

1. Sketch the graph of f .
2. Compute the sample mean and variance.
3. Give the sample values, ordered from smallest to largest.

Answer

2. $m = 1/12$, $s^2 = 203/121$
3. $(-2, -1, -1, -1, 0, 0, 0, 0, 1, 1, 2, 2)$

The following table gives a frequency distribution for the commuting distance to the math/stat building (in miles) for a sample of ESU students.

Class	Freq	Rel Freq	Density	Cum Freq	Cum Rel Freq	Midpoint
(0, 2]	6					
(2, 6]	16					
(6, 10]	18					
(10, 20]	10					
Total						

1. Complete the table
2. Sketch the density histogram
3. Sketch the cumulative relative frequency ogive.
4. Compute an approximation to the mean and standard deviation.

Answer

1.	Class	Freq	Rel Freq	Density	Cum Freq	Cum Rel Freq	Midpoint
	(0, 2]	6	0.12	0.06	6	0.12	1
	(2, 6]	16	0.32	0.08	22	0.44	4
	(6, 10]	18	0.36	0.09	40	0.80	8
	(10, 20]	10	0.20	0.02	50	1	15
	Total	50	1				

4. $m = 7.28$, $s = 4.549$

Error Function Exercises

In the error function app, select root mean square error. As you add points, note the shape of the graph of the error function, the value that minimizes the function, and the minimum value of the function.

In the error function app, select mean absolute error. As you add points, note the shape of the graph of the error function, the values that minimize the function, and the minimum value of the function.

Suppose that our data vector is $(2, 1, 5, 7)$. Explicitly give mae as a piecewise function and sketch its graph. Note that

1. All values of $a \in [2, 5]$ minimize mae.
2. mae is not differentiable at $a \in \{1, 2, 5, 7\}$.

Suppose that our data vector is $(3, 5, 1)$. Explicitly give mae as a piecewise function and sketch its graph. Note that

1. mae is minimized at $a = 3$.
2. mae is not differentiable at $a \in \{1, 3, 5\}$.

Simulation Exercises

Many of the apps in this project are simulations of experiments with a basic random variable of interest. When you run the simulation, you are performing independent replications of the experiment. In most cases, the app displays the standard deviation of the distribution, both numerically in a table and graphically as the radius of the blue, horizontal bar in the graph box. When you run the simulation, the sample standard deviation is also displayed numerically in the table and graphically as the radius of the red horizontal bar in the graph box.

In the binomial coin experiment, the random variable is the number of heads. For various values of the parameters n (the number of coins) and p (the probability of heads), run the simulation 1000 times and compare the sample standard deviation to the distribution standard deviation.

In the simulation of the matching experiment, the random variable is the number of matches. For selected values of n (the number of balls), run the simulation 1000 times and compare the sample standard deviation to the distribution standard deviation.

Run the simulation of the gamma experiment 1000 times for various values of the rate parameter r and the shape parameter k . Compare the sample standard deviation to the distribution standard deviation.

Probability Exercises

Suppose that X has probability density function $f(x) = 12x^2(1-x)$ for $0 \leq x \leq 1$. The distribution of X is a member of the beta family. Compute each of the following

1. $\mu = \mathbb{E}(X)$
2. $\sigma^2 = \text{var}(X)$
3. $d_3 = \mathbb{E}[(X - \mu)^3]$
4. $d_4 = \mathbb{E}[(X - \mu)^4]$

Answer

1. $3/5$
2. $1/25$
3. $-2/875$
4. $33/8750$

Suppose now that $(X_1, X_2, \dots, X_{10})$ is a random sample of size 10 from the beta distribution in the previous problem. Find each of the following:

1. $\mathbb{E}(M)$
2. $\text{var}(M)$
3. $\mathbb{E}(W^2)$
4. $\text{var}(W^2)$
5. $\mathbb{E}(S^2)$
6. $\text{var}(S^2)$

7. $\text{cov}(M, W^2)$
8. $\text{cov}(M, S^2)$
9. $\text{cov}(W^2, S^2)$

Answer

1. $3/5$
2. $1/250$
3. $1/25$
4. $19/87500$
5. $1/25$
6. $199/787500$
7. $-2/8750$
8. $-2/8750$
9. $19/87500$

Suppose that X has probability density function $f(x) = \lambda e^{-\lambda x}$ for $0 \leq x < \infty$, where $\lambda > 0$ is a parameter. Thus X has the exponential distribution with rate parameter λ . Compute each of the following

1. $\mu = \mathbb{E}(X)$
2. $\sigma^2 = \text{var}(X)$
3. $d_3 = \mathbb{E}[(X - \mu)^3]$
4. $d_4 = \mathbb{E}[(X - \mu)^4]$

Answer

1. $1/\lambda$
2. $1/\lambda^2$
3. $2/\lambda^3$
4. $9/\lambda^4$

Suppose now that (X_1, X_2, \dots, X_5) is a random sample of size 5 from the exponential distribution in the previous problem. Find each of the following:

1. $\mathbb{E}(M)$
2. $\text{var}(M)$
3. $\mathbb{E}(W^2)$
4. $\text{var}(W^2)$
5. $\mathbb{E}(S^2)$
6. $\text{var}(S^2)$
7. $\text{cov}(M, W^2)$
8. $\text{cov}(M, S^2)$
9. $\text{cov}(W^2, S^2)$

Answer

1. $1/\lambda$
2. $1/5\lambda^2$
3. $1/\lambda^2$
4. $8/5\lambda^4$
5. $1/\lambda^2$
6. $17/10\lambda^4$
7. $2/5\lambda^3$
8. $2/5\lambda^3$
9. $8/5\lambda^4$

Recall that for an *ace-six flat die*, faces 1 and 6 have probability $\frac{1}{4}$ each, while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each. Let X denote the score when an ace-six flat die is thrown. Compute each of the following:

1. $\mu = \mathbb{E}(X)$
2. $\sigma^2 = \text{var}(X)$
3. $d_3 = \mathbb{E}[(X - \mu)^3]$
4. $d_4 = \mathbb{E}[(X - \mu)^4]$

Answer

1. $7/2$
2. $15/4$
3. 0
4. $333/16$

Suppose now that an ace-six flat die is tossed 8 times. Find each of the following:

1. $\mathbb{E}(M)$
2. $\text{var}(M)$
3. $\mathbb{E}(W^2)$
4. $\text{var}(W^2)$
5. $\mathbb{E}(S^2)$
6. $\text{var}(S^2)$
7. $\text{cov}(M, W^2)$
8. $\text{cov}(M, S^2)$
9. $\text{cov}(W^2, S^2)$

Answer

1. $7/2$
2. $15/32$
3. $15/4$
4. $27/32$
5. $15/4$
6. $207/512$
7. 0
8. 0
9. $27/32$

Data Analysis Exercises

Statistical software should be used for the problems in this subsection.

Consider the petal length and species variables in Fisher's iris data.

1. Classify the variables by type and level of measurement.
2. Compute the sample mean and standard deviation, and plot a density histogram for petal length.
3. Compute the sample mean and standard deviation, and plot a density histogram for petal length by species.

Answers

1. petal length: continuous, ratio. species: discrete, nominal
2. $m = 37.8$, $s = 17.8$
3. $m(0) = 14.6$, $s(0) = 1.7$; $m(1) = 55.5$, $s(1) = 30.5$; $m(2) = 43.2$, $s(2) = 28.7$

Consider the erosion variable in the Challenger data set.

1. Classify the variable by type and level of measurement.
2. Compute the mean and standard deviation
3. Plot a density histogram with the classes $[0, 5)$, $[5, 40)$, $[40, 50)$, $[50, 60)$.

Answer

1. continuous, ratio

2. $m = 7.7$, $s = 17.2$

Consider Michelson's velocity of light data.

1. Classify the variable by type and level of measurement.
2. Plot a density histogram.
3. Compute the sample mean and standard deviation.
4. Find the sample mean and standard deviation if the variable is converted to km/hr. The transformation is $y = x + 299\,000$

Answer

1. continuous, interval
3. $m = 852.4$, $s = 79.0$
4. $m = 299\,852.4$, $s = 79.0$

Consider Short's parallax of the sun data.

1. Classify the variable by type and level of measurement.
2. Plot a density histogram.
3. Compute the sample mean and standard deviation.
4. Find the sample mean and standard deviation if the variable is converted to degrees. There are 3600 seconds in a degree.
5. Find the sample mean and standard deviation if the variable is converted to radians. There are $\pi/180$ radians in a degree.

Answer

1. continuous, ratio
3. $m = 8.616$, $s = 0.749$
4. $m = 0.00239$, $s = 0.000208$
5. $m = 0.0000418$, $s = 0.00000363$

Consider Cavendish's density of the earth data.

1. Classify the variable by type and level of measurement.
2. Compute the sample mean and standard deviation.
3. Plot a density histogram.

Answer

1. continuous, ratio
2. $m = 5.448$, $s = 0.221$

Consider the M&M data.

1. Classify the variables by type and level of measurement.
2. Compute the sample mean and standard deviation for each color count variable.
3. Compute the sample mean and standard deviation for the total number of candies.
4. Plot a relative frequency histogram for the total number of candies.
5. Compute the sample mean and standard deviation, and plot a density histogram for the net weight.

Answer

1. color counts: discrete ratio. net weight: continuous ratio.
2. $m(r) = 9.60$, $s(r) = 4.12$; $m(g) = 7.40$, $s(g) = 0.57$; $m(bl) = 7.23$, $s(bl) = 4.35$; $m(o) = 6.63$, $s(o) = 3.69$;
 $m(y) = 13.77$, $s(y) = 6.06$; $m(br) = 12.47$, $s(br) = 5.13$
3. $m(n) = 57.10$, $s(n) = 2.4$
5. $m(w) = 49.215$, $s(w) = 1.522$

Consider the body weight, species, and gender variables in the Cicada data.

1. Classify the variables by type and level of measurement.
2. Compute the relative frequency function for species and plot the graph.

3. Compute the relative frequency function for gender and plot the graph.
4. Compute the sample mean and standard deviation, and plot a density histogram for body weight.
5. Compute the sample mean and standard deviation, and plot a density histogram for body weight by species.
6. Compute the sample mean and standard deviation, and plot a density histogram for body weight by gender.

Answer

1. body weight: continuous, ratio. species: discrete, nominal. gender: discrete, nominal.
2. $f(0) = 0.423$, $f(1) = 0.519$, $f(2) = 0.058$
3. $f(0) = 0.567$, $f(1) = 0.433$
4. $m = 0.180$, $s = 0.059$
5. $m(0) = 0.168$, $s(0) = 0.054$; $m(1) = 0.185$, $s(1) = 0.185$; $m(2) = 0.225$, $s(2) = 0.107$
6. $m(0) = 0.206$, $s(0) = 0.052$; $m(1) = 0.145$, $s(1) = 0.051$

Consider Pearson's height data.

1. Classify the variables by type and level of measurement.
2. Compute the sample mean and standard deviation, and plot a density histogram for the height of the father.
3. Compute the sample mean and standard deviation, and plot a density histogram for the height of the son.

Answer

1. continuous ratio
2. $m(x) = 67.69$, $s(x) = 2.75$
3. $m(y) = 68.68$, $s(y) = 2.82$

This page titled [6.5: The Sample Variance](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.6: Order Statistics

Descriptive Theory

Recall again the basic model of statistics: we have a population of objects of interest, and we have various measurements (variables) that we make on these objects. We select objects from the population and record the variables for the objects in the sample; these become our data. Our first discussion is from a purely descriptive point of view. That is, we do not assume that the data are generated by an underlying probability distribution. But as always, remember that the data themselves define a probability distribution, namely the *empirical distribution*.

Order Statistics

Suppose that x is a real-valued variable for a population and that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ are the observed values of a sample of size n corresponding to this variable. The *order statistic* of rank k is the k th smallest value in the data set, and is usually denoted $x_{(k)}$. To emphasize the dependence on the sample size, another common notation is $x_{n:k}$. Thus,

$$x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(n-1)} \leq x_{(n)} \quad (6.6.1)$$

Naturally, the underlying variable x should be at least at the ordinal level of measurement. The order statistics have the same physical units as x . One of the first steps in *exploratory data analysis* is to order the data, so order statistics occur naturally. In particular, note that the *extreme order statistics* are

$$x_{(1)} = \min\{x_1, x_2, \dots, x_n\}, \quad x_{(n)} = \max\{x_1, x_2, \dots, x_n\} \quad (6.6.2)$$

The *sample range* is $r = x_{(n)} - x_{(1)}$ and the *sample midrange* is $\frac{r}{2} = \frac{1}{2}[x_{(n)} - x_{(1)}]$. These statistics have the same physical units as x and are measures of the dispersion of the data set.

The Sample Median

If n is odd, the *sample median* is the middle of the ordered observations, namely $x_{(k)}$ where $k = \frac{n+1}{2}$. If n is even, there is not a single middle observation, but rather two middle observations. Thus, the *median interval* is $[x_{(k)}, x_{(k+1)}]$ where $k = \frac{n}{2}$. In this case, the *sample median* is defined to be the midpoint of the median interval, namely $\frac{1}{2}[x_{(k)} + x_{(k+1)}]$ where $k = \frac{n}{2}$. In a sense, this definition is a bit arbitrary because there is no compelling reason to prefer one point in the median interval over another. For more on this issue, see the discussion of error functions in the section on Sample Variance. In any event, sample median is a natural statistic that gives a measure of the center of the data set.

Sample Quantiles

We can generalize the sample median discussed above to other sample quantiles. Thus, suppose that $p \in [0, 1]$. Our goal is to find the value that is the fraction p of the way through the (ordered) data set. We define the *rank* of the value that we are looking for as $(n-1)p + 1$. Note that the rank is a linear function of p , and that the rank is 1 when $p = 0$ and n when $p = 1$. But of course, the rank will not be an integer in general, so we let $k = \lfloor (n-1)p + 1 \rfloor$, the integer part of the desired rank, and we let $t = [(n-1)p + 1] - k$, the fractional part of the desired rank. Thus, $(n-1)p + 1 = k + t$ where $k \in \{1, 2, \dots, n\}$ and $t \in [0, 1]$. So, using *linear interpolation*, we define the *sample quantile* of order p to be

$$x_{[p]} = x_{(k)} + t[x_{(k+1)} - x_{(k)}] = (1-t)x_{(k)} + tx_{(k+1)} \quad (6.6.3)$$

Sample quantiles have the same physical units as the underlying variable x . The algorithm really does generalize the results for [sample medians](#).

The sample quantile of order $p = \frac{1}{2}$ is the median as defined earlier, in both cases where n is odd and where n is even.

The sample quantile of order $\frac{1}{4}$ is known as the *first quartile* and is frequently denoted q_1 . The sample quantile of order $\frac{3}{4}$ is known as the *third quartile* and is frequently denoted q_3 . The sample median which is the quartile of order $\frac{1}{2}$ is sometimes denoted q_2 . The *interquartile range* is defined to be $iqr = q_3 - q_1$. Note that iqr is a statistic that measures the spread of the distribution about the median, but of course this number gives less information than the *interval* $[q_1, q_3]$.

The statistic $q_1 - \frac{3}{2}iqr$ is called the *lower fence* and the statistic $q_3 + \frac{3}{2}iqr$ is called the *upper fence*. Sometimes *lower limit* and *upper limit* are used instead of lower fence and upper fence. Values in the data set that are below the lower fence or above the upper fence are potential *outliers*, that is, values that don't seem to fit the overall pattern of the data. An outlier can be due to a measurement error, or may be a valid but rather extreme value. In any event, outliers usually deserve additional study.

The five statistics $(x_{(1)}, q_1, q_2, q_3, x_{(n)})$ are often referred to as the *five-number summary*. Together, these statistics give a great deal of information about the data set in terms of the center, spread, and skewness. The five numbers roughly separate the data set into four intervals

each of which contains approximately 25% of the data. Graphically, the five numbers, and the outliers, are often displayed as a *boxplot*, sometimes called a *box and whisker plot*. A boxplot consists of an axis that extends across the range of the data. A line is drawn from smallest value that is not an outlier (of course this may be the minimum $x_{(1)}$) to the largest value that is not an outlier (of course, this may be the maximum $x_{(n)}$). Vertical marks (“whiskers”) are drawn at the ends of this line. A rectangular box extends from the first quartile q_1 to the third quartile q_3 and with an additional whisker at the median q_2 . Finally, the outliers are denoted as points (beyond the extreme whiskers). All statistical packages will compute the quartiles and most will draw boxplots. The picture below shows a boxplot with 3 outliers.




Figure 6.6.1: Boxplot

Alternate Definitions

The algorithm given above is not the only reasonable way to define sample quantiles, and indeed there are lots of alternatives. One natural method would be to first compute the empirical distribution function

$$F(x) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(x_i \leq x), \quad x \in \mathbb{R} \quad (6.6.4)$$

Recall that F has the mathematical properties of a distribution function, and in fact F is the distribution function of the empirical distribution of the data. Recall that this is the distribution that places probability $\frac{1}{n}$ at each data value x_i (so this is the discrete uniform distribution on $\{x_1, x_2, \dots, x_n\}$ if the data values are distinct). Thus, $F(x) = \frac{k}{n}$ for $x \in [x_{(k)}, x_{(k+1)})$. Then, we could define the quantile function to be the inverse of the distribution function, as we usually do for probability distributions:

$$F^{-1}(p) = \min\{x \in \mathbb{R} : F(x) \geq p\}, \quad p \in (0, 1) \quad (6.6.5)$$

It's easy to see that with this definition, the quantile of order $p \in (0, 1)$ is simply $x_{(k)}$ where $k = \lceil np \rceil$.

Another method is to compute the rank of the quantile of order $p \in (0, 1)$ as $(n+1)p$, rather than $(n-1)p+1$, and then use linear interpolation just as we have done. To understand the reasoning behind this method, suppose that the underlying variable x takes value in an interval (a, b) . Then the n points in the data set \mathbf{x} separate this interval into $n+1$ subintervals, so it's reasonable to think of $x_{(k)}$ as the quantile of order $\frac{k}{n+1}$. This method also reduces to the standard calculation for the median when $p = \frac{1}{2}$. However, the method will fail if p is so small that $(n+1)p < 1$ or so large that $(n+1)p > n$.

The primary definition that we give above is the one that is most commonly used in statistical software and spreadsheets. Moreover, when the sample size n is large, it doesn't matter very much which of these competing quantile definitions is used. All will give similar results.

Transformations

Suppose again that $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is a sample of size n from a population variable x , but now suppose also that $y = a + bx$ is a new variable, where $a \in \mathbb{R}$ and $b \in (0, \infty)$. Recall that transformations of this type are *location-scale transformations* and often correspond to changes in units. For example, if x is the length of an object in inches, then $y = 2.54x$ is the length of the object in centimeters. If x is the temperature of an object in degrees Fahrenheit, then $y = \frac{5}{9}(x - 32)$ is the temperature of the object in degrees Celsius. Let $\mathbf{y} = \mathbf{a} + b\mathbf{x}$ denote the sample from the variable y .

Order statistics and quantiles are preserved under location-scale transformations:

1. $y_{(i)} = a + bx_{(i)}$ for $i \in \{1, 2, \dots, n\}$
2. $y_{[p]} = a + bx_{[p]}$ for $p \in [0, 1]$

Proof

Part (a) follows easily from the fact that the location-scale transformation is strictly increasing and hence preserves order: $x_i < x_j$ if and only if $a + bx_i < a + bx_j$. For part (b), let $p \in [0, 1]$ and let $k \in \{1, 2, \dots, n\}$ and $t \in [0, 1]$ be as above in the definition of the sample quantile or order p . Then

$$y_{[p]} = y_{(k)} + t[y_{(k+1)} - y_{(k)}] = a + bx_{(k)} + t[a + bx_{(k+1)} - (a + bx_{(k)})] = a + b(x_{(k)} + t[x_{(k+1)} - x_{(k)}]) = a + bx_{[p]} \quad (6.6.6)$$

Like standard deviation (our most important measure of spread), range and interquartile range are not affected by the location parameter, but are scaled by the scale parameter.

The range and interquartile range of \mathbf{y} are

1. $r(\mathbf{y}) = b r(\mathbf{x})$
2. $\text{iqr}(\mathbf{y}) = b \text{iqr}(\mathbf{x})$

Proof

These results follow immediately from the previous result.

More generally, suppose $y = g(x)$ where g is a strictly increasing real-valued function on the set of possible values of x . Let $\mathbf{y} = (g(x_1), g(x_2), \dots, g(x_n))$ denote the sample corresponding to the variable y . Then (as in the proof of Theorem 2), the order statistics are preserved so $y_{(i)} = g(x_{(i)})$. However, if g is nonlinear, the quantiles are not preserved (because the quantiles involve *linear* interpolation). That is, $y_{[p]}$ and $g(x_{[p]})$ are not usually the same. When g is convex or concave we can at least give an inequality for the sample quantiles.

Suppose that $y = g(x)$ where g is strictly increasing. Then

1. $y_{(i)} = g(x_{(i)})$ for $i \in \{1, 2, \dots, n\}$
2. If g is convex then $y_{[p]} \geq g(x_{[p]})$ for $p \in [0, 1]$
3. If g is concave then $y_{[p]} \leq g(x_{[p]})$ for $p \in [0, 1]$

Proof

As noted, part (a) follows since g is strictly increasing and hence preserves order. Part (b) follows from the definition of convexity. For $p \in [0, 1]$, and $k \in \{1, 2, \dots, n\}$ and $t \in [0, 1]$ as in the definition of the sample quantile of order p , we have

$$y_{[p]} = (1-t)y_{(k)} + ty_{(k+1)} = (1-t)g(x_{(k)}) + tg(x_{(k+1)}) \geq g[(1-t)x_{(k)} + tx_{(k+1)}] = g(x_{[p]}) \quad (6.6.7)$$

Part (c) follows by the same argument.

Stem and Leaf Plots

A *stem and leaf plot* is a graphical display of the order statistics $(x_{(1)}, x_{(2)}, \dots, x_{(n)})$. It has the benefit of showing the data in a graphical way, like a histogram, and at the same time, preserving the ordered data. First we assume that the data have a fixed number format: a fixed number of digits, then perhaps a decimal point and another fixed number of digits. A stem and leaf plot is constructed by using an initial part of this string as the *stem*, and the remaining parts as the *leaves*. There are lots of variations in how to do this, so rather than give an exhaustive, complicated definition, we will just look at a couple of examples in the exercise below.

Probability Theory

We continue our discussion of order statistics except that now we assume that the variables are random variables. Specifically, suppose that we have a basic random experiment, and that X is a real-valued random variable for the experiment with distribution function F . We perform n independent replications of the basic experiment to generate a random sample $\mathbf{X} = (X_1, X_2, \dots, X_n)$ of size n from the distribution of X . Recall that this is a sequence of independent random variables, each with the distribution of X . All of the statistics defined in the previous section make sense, but now of course, they are random variables. We use the notation established previously, except that we follow our usual convention of denoting random variables with capital letters. Thus, for $k \in \{1, 2, \dots, n\}$, $X_{(k)}$ is the *kth order statistic*, that is, the k smallest of (X_1, X_2, \dots, X_n) . Our interest now is on the distribution of the order statistics and statistics derived from them.

Distribution of the k th order statistic

Finding the distribution function of an order statistic is a nice application of Bernoulli trials and the binomial distribution.

The distribution function F_k of $X_{(k)}$ is given by

$$F_k(x) = \sum_{j=k}^n \binom{n}{j} [F(x)]^j [1 - F(x)]^{n-j}, \quad x \in \mathbb{R} \quad (6.6.8)$$

Proof

For $x \in \mathbb{R}$, let

$$N_x = \sum_{i=1}^n \mathbf{1}(X_i \leq x) \quad (6.6.9)$$

so that N_x is the number of sample variables that fall in the interval $(-\infty, x]$. The indicator variables in the sum are independent, and each takes the value 1 with probability $F(x)$. Thus, N_x has the binomial distribution with parameters n and $F(x)$. Next note that $X_{(k)} \leq x$ if and only if $N_x \geq k$ for $x \in \mathbb{R}$ and $k \in \{1, 2, \dots, n\}$, since both events mean that there are at least k sample variables in the interval $(-\infty, x]$. Hence

$$\mathbb{P}(X_{(k)} \leq x) = \mathbb{P}(N_x \geq k) = \sum_{j=k}^n \binom{n}{j} [F(x)]^j [1 - F(x)]^{n-j} \quad (6.6.10)$$

As always, the extreme order statistics are particularly interesting.

The distribution functions F_1 of $X_{(1)}$ and F_n of $X_{(n)}$ are given by

1. $F_1(x) = 1 - [1 - F(x)]^n$ for $x \in \mathbb{R}$
2. $F_n(x) = [F(x)]^n$ for $x \in \mathbb{R}$

The quantile functions F_1^{-1} and F_n^{-1} of $X_{(1)}$ and $X_{(n)}$ are given by

1. $F_1^{-1}(p) = F^{-1}[1 - (1 - p)^{1/n}]$ for $p \in (0, 1)$
2. $F_n^{-1}(p) = F^{-1}(p^{1/n})$ for $p \in (0, 1)$

Proof

The formulas follow from the [previous theorem](#) and simple algebra. Recall that if G is a distribution function, then the corresponding quantile function is given by $G^{-1}(p) = \min\{x \in \mathbb{R} : G(x) \geq p\}$ for $p \in (0, 1)$.

When the underlying distribution is continuous, we can give a simple formula for the probability density function of an order statistic.

Suppose now that X has a continuous distribution with probability density function f . Then $X_{(k)}$ has a continuous distribution with probability density function f_k given by

$$f_k(x) = \frac{n!}{(k-1)!(n-k)!} [F(x)]^{k-1} [1 - F(x)]^{n-k} f(x), \quad x \in \mathbb{R} \quad (6.6.11)$$

Proof

Of course, $f_k(x) = F'_k(x)$. We take the derivatives term by term and use the product rule on

$$\frac{d}{dx} [F(x)]^j [1 - F(x)]^{n-j} = j[F(x)]^{j-1} f(x) [1 - F(x)]^{n-j} - (n-j)[F(x)]^j [1 - F(x)]^{n-j-1} f(x) \quad (6.6.12)$$

We use the binomial identities $j \binom{n}{j} = n \binom{n-1}{j-1}$ and $(n-j) \binom{n}{j} = n \binom{n-1}{j}$. The net effect is

$$f_k(x) = n f(x) \left[\sum_{j=k}^n \binom{n-1}{j-1} [F(x)]^{j-1} [1 - F(x)]^{(n-1)-(j-1)} - \sum_{j=k}^{n-1} \binom{n-1}{j} [F(x)]^j [1 - F(x)]^{(n-1)-j} \right] \quad (6.6.13)$$

The sums cancel, leaving only the $j = k$ term in the first sum. Hence

$$f_k(x) = n f(x) \binom{n-1}{k-1} [F(x)]^{k-1} [1 - F(x)]^{n-k} \quad (6.6.14)$$

$$\text{But } n \binom{n-1}{k-1} = \frac{n!}{(k-1)!(n-k)!}.$$

Heuristic Proof

There is a simple heuristic argument for this result. First, $f_k(x) dx$ is the probability that $X_{(k)}$ is in an infinitesimal interval of size dx about x . On the other hand, this event means that one of sample variables is in the infinitesimal interval, $k-1$ sample variables are less than x , and $n-k$ sample variables are greater than x . The number of ways of choosing these variables is the multinomial coefficient

$$\binom{n}{k-1, 1, n-k} = \frac{n!}{(k-1)!(n-k)!} \quad (6.6.15)$$

By independence, the probability that the chosen variables are in the specified intervals is

$$[F(x)]^{k-1} [1 - F(x)]^{n-k} f(x) dx \quad (6.6.16)$$

Here are the special cases for the extreme order statistics.

The probability density function f_1 of $X_{(1)}$ and f_n of $X_{(n)}$ are given by

1. $f_1(x) = n[1 - F(x)]^{n-1} f(x)$ for $x \in \mathbb{R}$
2. $f_n(x) = n[F(x)]^{n-1} f(x)$ for $x \in \mathbb{R}$

Joint Distributions

We assume again that X has a continuous distribution with distribution function F and probability density function f .

Suppose that $j, k \in \{1, 2, \dots, n\}$ with $j < k$. The joint probability density function $f_{j,k}$ of $(X_{(j)}, X_{(k)})$ is given by

$$f_{j,k}(x, y) = \frac{n!}{(j-1)!(k-j-1)!(n-k)!} [F(x)]^{j-1} [F(y) - F(x)]^{k-j-1} [1 - F(y)]^{n-k} f(x)f(y); \quad x, y \in \mathbb{R}, x < y \quad (6.6.17)$$

Heuristic Proof

We want to compute the probability that $X_{(j)}$ is in an infinitesimal interval dx about x and $X_{(k)}$ is in an infinitesimal interval dy about y . Note that there must be $j-1$ sample variables that are less than x , one variable in the infinitesimal interval about x , $k-j-1$ sample variables that are between x and y , one variable in the infinitesimal interval about y , and $n-k$ sample variables that are greater than y . The number of ways to select the variables is the multinomial coefficient

$$\binom{n}{j-1, 1, k-j-1, 1, n-k} = \frac{n!}{(j-1)!(k-j-1)!(n-k)!} \quad (6.6.18)$$

By independence, the probability that the chosen variables are in the specified intervals is

$$[F(x)]^{j-1} f(x) dx [F(y) - F(x)]^{k-j-1} f(y) dy [1 - F(y)]^{n-k} \quad (6.6.19)$$

From the joint distribution of two order statistics we can, in principle, find the distribution of various other statistics: the sample range R ; sample quantiles $X_{[p]}$ for $p \in [0, 1]$, and in particular the sample quartiles Q_1, Q_2, Q_3 ; and the inter-quartile range IQR. The joint distribution of the extreme order statistics $(X_{(1)}, X_{(n)})$ is a particularly important case.

The joint probability density function $f_{1,n}$ of $(X_{(1)}, X_{(n)})$ is given by

$$f_{1,n}(x, y) = n(n-1)[F(y) - F(x)]^{n-2} f(x)f(y); \quad x, y \in \mathbb{R}, x < y \quad (6.6.20)$$

Proof

This is a corollary of Theorem 7 with $j = 1$ and $k = n$.

Arguments similar to the one above can be used to obtain the joint probability density function of any number of the order statistics. Of course, we are particularly interested in the joint probability density function of *all* of the order statistics. It turns out that this density function has a remarkably simple form.

$(X_{(1)}, X_{(2)}, \dots, X_{(n)})$ has joint probability density function g given by

$$g(x_1, x_2, \dots, x_n) = n! f(x_1) f(x_2) \cdots f(x_n), \quad x_1 < x_2 < \cdots < x_n \quad (6.6.21)$$

Proof

For each permutation $\mathbf{i} = (i_1, i_2, \dots, i_n)$ of $(1, 2, \dots, n)$, let $S_{\mathbf{i}} = \{\mathbf{x} \in \mathbb{R}^n : x_{i_1} < x_{i_2} < \cdots < x_{i_n}\}$. On $S_{\mathbf{i}}$, the mapping $(x_1, x_2, \dots, x_n) \mapsto (x_{i_1}, x_{i_2}, \dots, x_{i_n})$ is one-to-one, has continuous first partial derivatives, and has Jacobian 1. The sets $S_{\mathbf{i}}$ where \mathbf{i} ranges over the $n!$ permutations of $(1, 2, \dots, n)$ are disjoint. The probability that (X_1, X_2, \dots, X_n) is not in one of these sets is 0. The result now follows from the multivariate change of variables formula.

Heuristic Proof

Again, there is a simple heuristic argument for this result. For each $\mathbf{x} \in \mathbb{R}^n$ with $x_1 < x_2 < \cdots < x_n$, there are $n!$ permutations of the coordinates of \mathbf{x} . The probability density of (X_1, X_2, \dots, X_n) at each of these points is $f(x_1) f(x_2) \cdots f(x_n)$. Hence the probability density of $(X_{(1)}, X_{(2)}, \dots, X_{(n)})$ at \mathbf{x} is $n!$ times this product.

Probability Plots

A *probability plot*, also called a *quantile-quantile plot* or a *Q-Q plot* for short, is an informal, graphical test to determine if observed data come from a specified distribution. Thus, suppose that we observe real-valued data (x_1, x_2, \dots, x_n) from a random sample of size n . We are interested in the question of whether the data could reasonably have come from a continuous distribution with distribution function F . First, we order that data from smallest to largest; this gives us the sequence of observed values of the order statistics: $(x_{(1)}, x_{(2)}, \dots, x_{(n)})$.

Note that we can view $x_{(i)}$ has the *sample* quantile of order $\frac{i}{n+1}$. Of course, by definition, the *distribution* quantile of order $\frac{i}{n+1}$ is $y_i = F^{-1}\left(\frac{i}{n+1}\right)$. If the data really do come from the distribution, then we would expect the points $((x_{(1)}, y_1), (x_{(2)}, y_2), \dots, (x_{(n)}, y_n))$ to

be close to the diagonal line $y = x$; conversely, strong deviation from this line is evidence that the distribution did not produce the data. The plot of these points is referred to as a *probability plot*.

Usually however, we are not trying to see if the data come from a *particular* distribution, but rather from a parametric *family* of distributions (such as the normal, uniform, or exponential families). We are usually forced into this situation because we don't know the parameters; indeed the next step, after the probability plot, may be to estimate the parameters. Fortunately, the probability plot method has a simple extension for any location-scale family of distributions. Thus, suppose that G is a given distribution function. Recall that the location-scale family associated with G has distribution function $F(x) = G\left(\frac{x-a}{b}\right)$ for, $x \in \mathbb{R}$, where $a \in \mathbb{R}$ is the location parameter and $b \in (0, \infty)$ is the scale parameter. Recall also that for $p \in (0, 1)$, if $z_p = G^{-1}(p)$ denote the quantile of order p for G and $y_p = F^{-1}(p)$ the quantile of order p for F . Then $y_p = a + bz_p$. It follows that if the probability plot constructed with distribution function F is nearly linear (and in particular, if it is close to the diagonal line), then the probability plot constructed with distribution function G will be nearly linear. Thus, we can use the distribution function G without having to know the location and scale parameters.

In the exercises below, you will explore probability plots for the normal, exponential, and uniform distributions. We will study a formal, quantitative procedure, known as the chi-square goodness of fit test in the chapter on Hypothesis Testing.

Exercises and Applications

Basic Properties

Suppose that x is the temperature (in degrees Fahrenheit) for a certain type of electronic component after 10 hours of operation. A sample of 30 components has five number summary (84, 102, 113, 120, 135.)

1. Classify x by type and level of measurement.
2. Find the range and interquartile range.
3. Find the five number summary, range, and interquartile range if the temperature is converted to degrees Celsius. The transformation is $y = \frac{5}{9}(x - 32)$.

Answer

1. continuous, interval
2. 51, 18
3. (28.89, 38.89, 45.00, 48.89, 57.22), 33, 10

Suppose that x is the length (in inches) of a machined part in a manufacturing process. A sample of 50 parts has five number summary (9.6, 9.8, 10.0, 10.1, 10.3).

1. Classify x by type and level of measurement.
2. Find the range and interquartile range.
3. Find the five number summary, range, and interquartile if length is measured in centimeters. The transformation is $y = 2.54x$.

Answer

1. continuous, ratio
2. 0.7, 0.3
3. (24.38, 24.89, 25.40, 25.65, 26.16), 78, 0.76
- 4.

Professor Moriarity has a class of 25 students in her section of Stat 101 at Enormous State University (ESU). For the first midterm exam, the five number summary was (16, 52, 64, 72, 81) (out of a possible 100 points). Professor Moriarity thinks the grades are a bit low and is considering various transformations for increasing the grades.

1. Find the range and interquartile range.
2. Suppose she adds 10 points to each grade. Find the five number summary, range, and interquartile range for the transformed grades.
3. Suppose she multiplies each grade by 1.2. Find the five number summary, range, and interquartile range for the transformed grades.
4. Suppose she uses the transformation $w = 10\sqrt{x}$, which curves the grades greatly at the low end and very little at the high end. Give whatever information you can about the five number summary of the transformed grades.
5. Determine whether the low score of 16 is an outlier.

Answer

1. 65, 20
2. (26, 62, 74, 82, 91), 65, 20
3. (19.2, 62.4, 76.8, 86.4, 97.2), 78, 24
4. $y_{(1)} = 40$, $q_1 \leq 72.11$, $q_2 \leq 80$, $q_3 \leq 84.85$, $y_{(25)} = 90$

5. The lower fence is 27, so yes 16 is an outlier.

Computational Exercises

All statistical software packages will compute order statistics and quantiles, draw stem-and-leaf plots and boxplots, and in general perform the numerical and graphical procedures discussed in this section. For real statistical experiments, particularly those with large data sets, the use of statistical software is essential. On the other hand, there is some value in performing the computations by hand, with small, artificial data sets, in order to master the concepts and definitions. In this subsection, do the computations and draw the graphs with minimal technological aids.

Suppose that x is the number of math courses completed by an ESU student. A sample of 10 ESU students gives the data $x = (3, 1, 2, 0, 2, 4, 3, 2, 1, 2)$

1. Classify x by type and level of measurement.
2. Give the order statistics
3. Compute the five number summary and draw the boxplot.
4. Compute the range and the interquartile range.

Answer

1. discrete, ratio
2. (0, 1, 1, 2, 2, 2, 2, 3, 3, 4)
3. (0, 1.25, 2, 2.75, 4)
4. 4, 1.5

Suppose that a sample of size 12 from a discrete variable x has empirical density function given by $f(-2) = 1/12$, $f(-1) = 1/4$, $f(0) = 1/3$, $f(1) = 1/6$, $f(2) = 1/6$.

1. Give the order statistics.
2. Compute the five number summary and draw the boxplot.
3. Compute the range and the interquartile range.

Answer

1. (-2, -1, -1, -1, 0, 0, 0, 0, 1, 1, 2, 2)
2. (-2, -1, 0, 1, 2)
3. 4, 2

The stem and leaf plot below gives the grades for a 100-point test in a probability course with 38 students. The first digit is the stem and the second digit is the leaf. Thus, the low score was 47 and the high score was 98. The scores in the 6 row are 60, 60, 62, 63, 65, 65, 67, 68.

4	7
5	0346
6	00235578
7	0112346678899
8	0367889
9	1368

Compute the five number summary and draw the boxplot.

Answer

(47, 65, 75, 83, 98)

App Exercises

In the histogram app, construct a distribution with at least 30 values of each of the types indicated below. Note the five number summary.

1. A uniform distribution.
2. A symmetric, unimodal distribution.
3. A unimodal distribution that is skewed right.
4. A unimodal distribution that is skewed left.
5. A symmetric bimodal distribution.
6. A u -shaped distribution.

In the error function app, Start with a distribution and add additional points as follows. Note the effect on the five number summary:

1. Add a point below $x_{(1)}$.
2. Add a point between $x_{(1)}$ and q_1 .
3. Add a point between q_1 and q_2 .
4. Add a point between q_2 and q_3 .
5. Add a point between q_3 and $x_{(n)}$.
6. Add a point above $x_{(n)}$.

In the last problem, you may have noticed that when you add an additional point to the distribution, one or more of the five statistics does not change. In general, quantiles can be relatively insensitive to changes in the data.

The Uniform Distribution

Recall that the standard uniform distribution is the uniform distribution on the interval $[0, 1]$.

Suppose that \mathbf{X} is a random sample of size n from the standard uniform distribution. For $k \in \{1, 2, \dots, n\}$, $X_{(k)}$ has the beta distribution, with left parameter k and right parameter $n - k + 1$. The probability density function f_k is given by

$$f_k(x) = \frac{n!}{(k-1)!(n-k)!} x^{k-1} (1-x)^{n-k}, \quad 0 \leq x \leq 1 \quad (6.6.22)$$

Proof

This follows immediately from the [basic theorem above](#) since $f(x) = 1$ and $F(x) = x$ for $0 \leq x \leq 1$. From the form of f_k we can identify the distribution as beta with left parameter k and right parameter $n - k + 1$.

In the order statistic experiment, select the standard uniform distribution and $n = 5$. Vary k from 1 to 5 and note the shape of the probability density function of $X_{(k)}$. For each value of k , run the simulation 1000 times and compare the empirical density function to the true probability density function.

It's easy to extend the results for the standard uniform distribution to the general uniform distribution on an interval.

Suppose that \mathbf{X} is a random sample of size n from the uniform distribution on the interval $[a, a + h]$ where $a \in \mathbb{R}$ and $h \in (0, \infty)$. For $k \in \{1, 2, \dots, n\}$, $X_{(k)}$ has the beta distribution with left parameter k , right parameter $n - k + 1$, location parameter a , and scale parameter h . In particular,

1. $\mathbb{E}(X_{(k)}) = a + h \frac{k}{n+1}$
2. $\text{var}(X_{(k)}) = h^2 \frac{k(n-k+1)}{(n+1)^2(n+2)}$

Proof

Suppose that $\mathbf{U} = (U_1, U_2, \dots, U_n)$ is a random sample of size n from the standard uniform distribution, and let $X_i = a + hU_i$ for $i \in \{1, 2, \dots, n\}$. Then $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the uniform distribution on the interval $[a, a + h]$, and moreover, $X_{(k)} = a + hU_{(k)}$. So the distribution of $X_{(k)}$ follows from the previous result. Parts (a) and (b) follow from standard results for the beta distribution.

We return to the standard uniform distribution and consider the range of the random sample.

Suppose that \mathbf{X} is a random sample of size n from the standard uniform distribution. The sample range R has the beta distribution with left parameter $n - 1$ and right parameter 2. The probability density function g is given by

$$g(r) = n(n-1)r^{n-2}(1-r), \quad 0 \leq r \leq 1 \quad (6.6.23)$$

Proof

From the [result above](#), the joint PDF of $(X_{(1)}, X_{(n)})$ is $f_{1,n}(x, y) = n(n-1)(y-x)^{n-2}$ for $0 \leq x \leq y \leq 1$. Hence, for $r \in [0, 1]$,

$$\mathbb{P}(R > r) = \mathbb{P}(X_{(n)} - X_{(1)} > r) = \int_0^{1-r} \int_{x+r}^1 n(n-1)(y-x)^{n-2} dy dx = (n-1)r^n - nr^{n-1} + 1 \quad (6.6.24)$$

It follows that the CDF of R is $G(r) = nr^{n-1} - (n-1)r^n$ for $0 \leq r \leq 1$. Taking the derivative with respect to r and simplifying gives the PDF $g(r) = n(n-1)r^{n-2}(1-r)$ for $0 \leq r \leq 1$. We can tell from the form of g that the distribution is beta with left parameter $n - 1$ and right parameter 2.

Once again, it's easy to extend this result to a general uniform distribution.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the uniform distribution on $[a, a+h]$ where $a \in \mathbb{R}$ and $h \in (0, \infty)$. The sample range $R = X_{(n)} - X_{(1)}$ has the beta distribution with left parameter $n-1$, right parameter 2, and scale parameter h . In particular,

1. $\mathbb{E}(R) = h \frac{n-1}{n+1}$
2. $\text{var}(R) = h^2 \frac{2(n-1)}{(n+1)^2(n+2)}$

Proof

Suppose again that $\mathbf{U} = (U_1, U_2, \dots, U_n)$ is a random sample of size n from the standard uniform distribution, and let $X_i = a + hU_i$ for $i \in \{1, 2, \dots, n\}$. Then $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the uniform distribution on the interval $[a, a+h]$, and moreover, $X_{(k)} = a + hU_{(k)}$. Hence $X_{(n)} - X_{(1)} = h(U_{(n)} - U_{(1)})$ so the distribution of R follows from the previous result. Parts (a) and (b) follow from standard results for the beta distribution.

The joint distribution of the order statistics for a sample from the uniform distribution is easy to get.

Suppose that (X_1, X_2, \dots, X_n) is a random sample of size n from the uniform distribution on the interval $[a, a+h]$, where $a \in \mathbb{R}$ and $h \in (0, \infty)$. Then $(X_{(1)}, X_{(2)}, \dots, X_{(n)})$ is uniformly distributed on $\{\mathbf{x} \in [a, a+h]^n : a \leq x_1 \leq x_2 \leq \dots \leq x_n < a+h\}$.

Proof

This follows easily from the fact that (X_1, X_2, \dots, X_n) is uniformly distributed on $[a, a+h]^n$. From the [result above](#), the joint PDF of the order statistics is $g(x_1, x_2, \dots, x_n) = n!/h^n$ for $(x_1, x_2, \dots, x_n) \in [a, a+h]^n$ with $a \leq x_1 \leq x_2 \leq \dots \leq x_n \leq a+h$.

The Exponential Distribution

Recall that the exponential distribution with rate parameter $\lambda > 0$ has probability density function

$$f(x) = \lambda e^{-\lambda x}, \quad 0 \leq x < \infty \quad (6.6.25)$$

The exponential distribution is widely used to model failure times and other random times under certain ideal conditions. In particular, the exponential distribution governs the times between arrivals in the Poisson process.

Suppose that \mathbf{X} is a random sample of size n from the exponential distribution with rate parameter λ . The probability density function of the k th order statistic $X_{(k)}$ is

$$f_k(x) = \frac{n!}{(k-1)!(n-k)!} \lambda (1 - e^{-\lambda x})^{k-1} e^{-\lambda(n-k+1)x}, \quad 0 \leq x < \infty \quad (6.6.26)$$

In particular, the minimum of the variables $X_{(1)}$ also has an exponential distribution, but with rate parameter $n\lambda$.

Proof

The PDF of $X_{(k)}$ follows from the [theorem above](#) since $F(x) = 1 - e^{-\lambda x}$ for $0 \leq x < \infty$. Substituting $k=1$ gives $f_1(x) = n\lambda e^{-n\lambda x}$ for $0 \leq x < \infty$.

In the order statistic experiment, select the standard exponential distribution and $n=5$. Vary k from 1 to 5 and note the shape of the probability density function of $X_{(k)}$. For each value of k , run the simulation 1000 times and compare the empirical density function to the true probability density function.

Suppose again that \mathbf{X} is a random sample of size n from the exponential distribution with rate parameter λ . The sample range R has the same distribution as the maximum of a random sample of size $n-1$ from the exponential distribution. The probability density function is

$$h(t) = (n-1)\lambda(1 - e^{-\lambda t})^{n-2} e^{-\lambda t}, \quad 0 \leq t < \infty \quad (6.6.27)$$

Proof

By the [result above](#), $(X_{(1)}, X_{(n)})$ has joint PDF $f_{1,n}(x, y) = n(n-1)\lambda^2(e^{-\lambda x} - e^{-\lambda y})^{n-2} e^{-\lambda x} e^{-\lambda y}$ for $0 \leq x \leq y < \infty$. Hence for $0 \leq t < \infty$,

$$\mathbb{P}(R \leq t) = \mathbb{P}(X_{(n)} - X_{(1)} \leq t) = \int_0^\infty \int_x^{x+t} n(n-1)\lambda^2(e^{-\lambda x} - e^{-\lambda y})^{n-2} e^{-\lambda x} e^{-\lambda y} dy dx \quad (6.6.28)$$

Substituting $u = e^{-\lambda y}$, $du = -\lambda e^{-\lambda y} dy$ into the inside integral and evaluating gives

$$\mathbb{P}(R \leq t) = \int_0^{\infty} n\lambda e^{-n\lambda x} (1 - e^{-\lambda t})^{n-1} dx = (1 - e^{-\lambda t})^{n-1} \quad (6.6.29)$$

Differentiating with respect to t gives the the PDF. Comparing with our [previous result](#), we see that this is the PDF of the maximum of a sample of size $n - 1$ from the exponential distribution.

Suppose again that \mathbf{X} is a random sample of size n from the exponential distribution with rate parameter λ . The joint probability density function of the order statistics $(X_{(1)}, X_{(2)}, \dots, X_{(n)})$ is

$$g(x_1, x_2, \dots, x_n) = n!\lambda^n e^{-\lambda(x_1 + x_2 + \dots + x_n)}, \quad 0 \leq x_1 \leq x_2 \leq \dots \leq x_n < \infty \quad (6.6.30)$$

Proof

This follows from the [result above](#) and simple algebra.

Dice

Four fair dice are rolled. Find the probability density function of each of the order statistics.

Answer

x	1	2	3	4	5	6
$f_1(x)$	$\frac{671}{1296}$	$\frac{369}{1296}$	$\frac{175}{1296}$	$\frac{65}{1296}$	$\frac{15}{1296}$	$\frac{1}{1296}$
$f_2(x)$	$\frac{171}{1296}$	$\frac{357}{1296}$	$\frac{363}{1296}$	$\frac{261}{1296}$	$\frac{123}{1296}$	$\frac{21}{1296}$
$f_3(x)$	$\frac{21}{1296}$	$\frac{123}{1296}$	$\frac{261}{1296}$	$\frac{363}{1296}$	$\frac{357}{1296}$	$\frac{171}{1296}$
$f_4(x)$	$\frac{1}{1296}$	$\frac{15}{1296}$	$\frac{65}{1296}$	$\frac{175}{1296}$	$\frac{369}{1296}$	$\frac{671}{1296}$

In the dice experiment, select the order statistic and die distribution given in parts (a)–(d) below. Increase the number of dice from 1 to 20, noting the shape of the probability density function at each stage. Now with $n = 4$, run the simulation 1000 times, and note the apparent convergence of the relative frequency function to the probability density function.

1. Maximum score with fair dice.
2. Minimum score with fair dice.
3. Maximum score with ace-six flat dice.
4. Minimum score with ace-six flat dice.

Four fair dice are rolled. Find the joint probability density function of the four order statistics.

Answer

The joint probability density function g is defined on $\{(x_1, x_2, x_3, x_4) \in \{1, 2, 3, 4, 5, 6\}^4 : x_1 \leq x_2 \leq x_3 \leq x_4\}$

1. $g(x_1, x_2, x_3, x_4) = \frac{1}{1296}$ if the coordinates are all the same (there are 6 such vectors).
2. $g(x_1, x_2, x_3, x_4) = \frac{4}{1296}$ if there are two distinct coordinates, one value occurring 3 times and the other value once (there are 30 such vectors).
3. $g(x_1, x_2, x_3, x_4) = \frac{6}{1296}$ if there are two distinct coordinates in (x_1, x_2, x_3, x_4) , each value occurring 2 times (there are 15 such vectors).
4. $g(x_1, x_2, x_3, x_4) = \frac{12}{1296}$ if there are three distinct coordinates, one value occurring twice and the other values once (there are 60 such vectors).
5. $g(x_1, x_2, x_3, x_4) = \frac{24}{1296}$ if the coordinates are distinct (there are 15 such vectors).

Four fair dice are rolled. Find the probability density function of the sample range.

Answer

R has probability density function h given by $h(0) = \frac{6}{1296}$, $h(1) = \frac{70}{1296}$, $h(2) = \frac{300}{1296}$, $h(3) = \frac{300}{1296}$, $h(4) = \frac{318}{1296}$, $h(5) = \frac{302}{1296}$

Probability Plot Simulations

In the probability plot experiment, set the sampling distribution to normal distribution with mean 5 and standard deviation 2. Set the sample size to $n = 20$. For each of the following test distributions, run the experiment 50 times and note the geometry of the probability

plot:

1. Standard normal
2. Uniform on the interval $[0, 1]$
3. Exponential with parameter 1

In the probability plot experiment, set the sampling distribution to the uniform distribution on $[4, 10]$. Set the sample size to $n = 20$. For each of the following test distributions, run the experiment 50 times and note the geometry of the probability plot:

1. Standard normal
2. Uniform on the interval $[0, 1]$
3. Exponential with parameter 1

In the probability plot experiment, Set the sampling distribution to the exponential distribution with parameter 3. Set the sample size to $n = 20$. For each of the following test distributions, run the experiment 50 times and note the geometry of the probability plot:

1. Standard normal
2. Uniform on the interval $[0, 1]$
3. Exponential with parameter 1

Data Analysis Exercises

Statistical software should be used for the problems in this subsection.

Consider the petal length and species variables in Fisher's iris data.

1. Classify the variables by type and level of measurement.
2. Compute the five number summary and draw the boxplot for petal length.
3. Compute the five number summary and draw the boxplot for petal length by species.
4. Draw the normal probability plot for petal length.

Answers

1. petal length: continuous, ratio. type: discrete, nominal
2. (10, 15, 44, 51, 69)
3. type 0: (10, 14, 15, 16, 19) type 1: (45, 51, 55.5, 59, 69) type 2: (30, 40, 44, 47, 56)

Consider the erosion variable in the Challenger data set.

1. Classify the variable by type and level of measurement.
2. Compute the five number summary and draw the boxplot.
3. Identify any outliers.

Answer

1. continuous, ratio
2. (0, 0, 0, 0, 53)
3. All of the positive values 28, 40, 48, and 53 are outliers.

A stem and leaf plot of Michelson's velocity of light data is given below. In this example, the last digit (which is always 0) has been left out, for convenience. Also, note that there are two sets of leaves for each stem, one corresponding to leaves from 0 to 4 (so actually from 00 to 40) and the other corresponding to leaves from 5 to 9 (so actually from 50 to 90). Thus, the minimum value is 620 and the numbers in the second 7 row are 750, 760, 760, and so forth.

6	2
6	5
7	222444
7	566666788999
8	000001111111111223344444444
9	0011233444
9	55566667888
10	000
10	7

Classify the variable by type and level of measurement.

1. Compute the five number summary and draw the boxplot.
2. Compute the five number summary for the velocity in km/hr. The transformation is $y = x + 299\,000$.
3. Draw the normal probability plot.

Answer

1. continuous, interval
2. (620, 805, 850, 895, 1071)
3. (299 620, 299 805, 299 850, 299 895, 300 071)

Consider Short's parallax of the sun data.

1. Classify the variable by type and level of measurement.
2. Compute the five number summary and draw the boxplot.
3. Compute the five number summary and draw the boxplot if the variable is converted to degrees. There are 3600 seconds in a degree.
4. Compute the five number summary and draw the boxplot if the variable is converted to radians. There are $\pi/180$ radians in a degree.
5. Draw the normal probability plot.

Answer

1. continuous, ratio
2. (5.76, 8.34, 8.50, 9.02, 10.57)
3. (0.00160, 0.00232, 0.00236, 0.00251, 0.00294)
4. (0.0000278, 0.0000404, 0.0000412, 0.0000437, 0.0000512)

Consider Cavendish's density of the earth data.

1. Classify the variable by type and level of measurement.
2. Compute the five number summary and draw the boxplot.
3. Draw the normal probability plot.

Answer

1. continuous, ratio
2. (4.88, 5.30, 5.46, 5.61, 5.85)

Consider the M&M data.

1. Classify the variables by type and level of measurement.
2. Compute the five number summary and draw the boxplot for each color count.
3. Construct a stem and leaf plot for the total number of candies.
4. Compute the five number summary and draw the boxplot for the total number of candies.
5. Compute the five number summary and draw the boxplot for net weight.

Answer

1. color counts: discrete ratio. net weight: continuous ratio.
2. red: (3, 5.5, 9, 14, 20) green: (2, 5, 7, 9, 17) blue: (1, 4, 6.5, 10, 19) orange: (0, 3.5, 6, 10.5, 13) yellow: (3, 8, 13.5, 18, 26) brown: (4, 8, 12.5, 18, 20)

3.

5	0												
5	3												
5	4	5	5	5	5								
5	6	6	6	6	7	7	7						
5	8	8	8	8	8	8	8	8	8	9	9	9	
6	0	0	1	1									

4. (50, 55.5, 58, 60, 61)
5. (46.22, 48.28, 49.07, 50.23, 52.06)

Consider the body weight, species, and gender variables in the Cicada data.

1. Classify the variables by type and level of measurement.
2. Compute the five number summary and draw the boxplot for body weight.
3. Compute the five number summary and draw the boxplot for body weight by species.
4. Compute the five number summary and draw the boxplot for body weight by gender.

Answer

1. body weight: continuous, ratio. species: discrete, nominal. gender: discrete, nominal.
2. (0.08, 0.13, 0.17, 0.22, 0.39)
3. species 0: (0.08, 0.13, 0.16, 0.21, 0.27) species 1: (0.08, 0.14, 0.18, 0.23, 0.31) species 2: (0.12, 0.12, 0.215, 0.29, 0.39)
4. female: (0.08, 0.17, 0.21, 0.25, 0.31) male: (0.08, 0.12, 0.14, 0.16, 0.39)

Consider Pearson's height data.

1. Classify the variables by type and level of measurement.
2. Compute the five number summary and sketch the boxplot for the height of the father.
3. Compute the five number summary and sketch the boxplot for the height of the son.

Answer

1. continuous ratio
2. (59.0, 65.8, 67.8, 69.6, 75.4)
3. (58.5, 66.9, 68.6, 70.5, 78.4)

This page titled [6.6: Order Statistics](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.7: Sample Correlation and Regression

Descriptive Theory

Recall the basic model of statistics: we have a population of objects of interest, and we have various measurements (variables) that we make on these objects. We select objects from the population and record the variables for the objects in the sample; these become our data. Our first discussion is from a purely descriptive point of view. That is, we do not assume that the data are generated by an underlying probability distribution. But as always, remember that the data themselves define a probability distribution, namely the *empirical distribution* that assigns equal probability to each data point.

Suppose that x and y are real-valued variables for a population, and that $((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n))$ is an observed sample of size n from (x, y) . We will let $\mathbf{x} = (x_1, x_2, \dots, x_n)$ denote the sample from x and $\mathbf{y} = (y_1, y_2, \dots, y_n)$ the sample from y . In this section, we are interested in statistics that are *measures of association* between the \mathbf{x} and \mathbf{y} , and in finding the line (or other curve) that best fits the data.

Recall that the sample means are

$$m(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n x_i, \quad m(\mathbf{y}) = \frac{1}{n} \sum_{i=1}^n y_i \quad (6.7.1)$$

and the sample variances are

$$s^2(\mathbf{x}) = \frac{1}{n-1} \sum_{i=1}^n [x_i - m(\mathbf{x})]^2, \quad s^2(\mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n [y_i - m(\mathbf{y})]^2 \quad (6.7.2)$$

Scatterplots

Often, the first step in *exploratory data analysis* is to draw a graph of the points; this is called a *scatterplot* and can give a visual sense of the statistical relationship between the variables.




Figure 6.7.1: A scatterplot

In particular, we are interested in whether the cloud of points seems to show a linear trend or whether some nonlinear curve might fit the cloud of points. We are interested in the extent to which one variable x can be used to predict the other variable y .

Definitions

Our next goal is to define statistics that measure the association between the x and y data.

The *sample covariance* is defined to be

$$s(\mathbf{x}, \mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n [x_i - m(\mathbf{x})][y_i - m(\mathbf{y})] \quad (6.7.3)$$

Assuming that the data vectors are not constant, so that the standard deviations are positive, the *sample correlation* is defined to be

$$r(\mathbf{x}, \mathbf{y}) = \frac{s(\mathbf{x}, \mathbf{y})}{s(\mathbf{x})s(\mathbf{y})} \quad (6.7.4)$$

Note that the sample covariance is an average of the product of the deviations of the x and y data from their means. Thus, the physical unit of the sample covariance is the product of the units of x and y . Correlation is a standardized version of covariance. In particular, correlation is dimensionless (has no physical units), since the covariance in the numerator and the product of the standard deviations in the denominator have the same units (the product of the units of x and y). Note also that covariance and correlation have the same *sign*: positive, negative, or zero. In the first case, the data \mathbf{x} and \mathbf{y} are said to be *positively correlated*; in the second case \mathbf{x} and \mathbf{y} are said to be *negatively correlated*; and in the third case \mathbf{x} and \mathbf{y} are said to be *uncorrelated*.

To see that the sample covariance is a measure of association, recall first that the point $(m(\mathbf{x}), m(\mathbf{y}))$ is a measure of the center of the bivariate data. Indeed, if each point is the location of a unit mass, then $(m(\mathbf{x}), m(\mathbf{y}))$ is the *center of mass* as defined in physics. Horizontal and vertical lines through this center point divide the plane into four quadrants. The product deviation $[x_i - m(\mathbf{x})][y_i - m(\mathbf{y})]$ is positive in the first and third quadrants and negative in the second and fourth quadrants. After we study linear regression [below](#), we will have a much deeper sense of what covariance measures.




Figure 6.7.2: Scatterplot with means

You may be perplexed that we average the product deviations by dividing by $n - 1$ rather than n . The best explanation is that in the probability model discussed [below](#), the sample covariance is an unbiased estimator of the distribution covariance. However, the mode of averaging can also be understood in terms of *degrees of freedom*, as was done for sample variance. Initially, we have $2n$ degrees of freedom in the bivariate data. We lose two by computing the sample means $m(\mathbf{x})$ and $m(\mathbf{y})$. Of the remaining $2n - 2$ degrees of freedom, we lose $n - 1$ by computing the product deviations. Thus, we are left with $n - 1$ degrees of freedom total. As is typical in statistics, we average not by dividing by the number of terms in the sum but rather by the number of degrees of freedom in those terms. However, from a purely descriptive point of view, it would also be reasonable to divide by n .

Recall that there is a natural probability distribution associated with the data, namely the empirical distribution that gives probability $\frac{1}{n}$ to each data point (x_i, y_i) . (Thus, if these points are distinct this is the discrete uniform distribution on the data.) The sample means are simply the expected values of this bivariate distribution, and except for a constant multiple (dividing by $n - 1$ rather than n), the sample variances are simply the variances of this bivariate distribution. Similarly, except for a constant multiple (again dividing by $n - 1$ rather than n), the sample covariance is the covariance of the bivariate distribution and the sample correlation is the correlation of the bivariate distribution. All of the following results in our discussion of descriptive statistics are actually special cases of more general results for probability distributions.

Properties of Covariance

The next few exercises establish some essential properties of sample covariance. As usual, bold symbols denote samples of a fixed size n from the corresponding population variables (that is, vectors of length n), while symbols in regular type denote real numbers. Our first result is a formula for sample covariance that is sometimes better than the definition for computational purposes. To state the result succinctly, let $\mathbf{xy} = (x_1 y_1, x_2 y_2, \dots, x_n y_n)$ denote the sample from the product variable xy .

The sample covariance can be computed as follows:

$$s(\mathbf{x}, \mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n x_i y_i - \frac{n}{n-1} m(\mathbf{x}) m(\mathbf{y}) = \frac{n}{n-1} [m(\mathbf{xy}) - m(\mathbf{x}) m(\mathbf{y})] \quad (6.7.5)$$

Proof

Note that

$$\sum_{i=1}^n [(x_i - m(\mathbf{x}))][y_i - m(\mathbf{y})] = \sum_{i=1}^n [x_i y_i - x_i m(\mathbf{y}) - y_i m(\mathbf{x}) + m(\mathbf{x}) m(\mathbf{y})] \quad (6.7.6)$$

$$= \sum_{i=1}^n x_i y_i - m(\mathbf{y}) \sum_{i=1}^n x_i - m(\mathbf{x}) \sum_{i=1}^n y_i + n m(\mathbf{x}) m(\mathbf{y}) \quad (6.7.7)$$

$$= \sum_{i=1}^n x_i y_i - n m(\mathbf{y}) m(\mathbf{x}) - n m(\mathbf{x}) m(\mathbf{y}) + n m(\mathbf{x}) m(\mathbf{y}) \quad (6.7.8)$$

$$= \sum_{i=1}^n x_i y_i - n m(\mathbf{x}) m(\mathbf{y}) \quad (6.7.9)$$

The following theorem gives another formula for the sample covariance, one that does not require the computation of intermediate statistics.

The sample covariance can be computed as follows:

$$s(\mathbf{x}, \mathbf{y}) = \frac{1}{2n(n-1)} \sum_{i=1}^n \sum_{j=1}^n (x_i - x_j)(y_i - y_j) \quad (6.7.10)$$

Proof

Note that

$$\sum_{i=1}^n \sum_{j=1}^n (x_i - x_j)(y_i - y_j) = \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^n [x_i - m(\mathbf{x}) + m(\mathbf{x}) - x_j][y_i - m(\mathbf{y}) + m(\mathbf{y}) - y_j] \quad (6.7.11)$$

$$= \sum_{i=1}^n \sum_{j=1}^n [(x_i - m(\mathbf{x}))[y_i - m(\mathbf{y})] + [x_i - m(\mathbf{x})][m(\mathbf{y}) - y_j] + [m(\mathbf{x}) - x_j][y_i - m(\mathbf{y})] + [m(\mathbf{x}) - x_j][m(\mathbf{y}) - y_j]] \quad (6.7.12)$$

We compute the sums term by term. The first is

$$n \sum_{i=1}^n [x_i - m(\mathbf{x})][y_i - m(\mathbf{y})] \quad (6.7.13)$$

The second two sums are 0. The last sum is

$$n \sum_{j=1}^n [m(\mathbf{x}) - x_j][m(\mathbf{y}) - y_j] = n \sum_{i=1}^n [x_i - m(\mathbf{x})][y_i - m(\mathbf{y})] \quad (6.7.14)$$

Dividing the entire sum by $2n(n-1)$ results in $\text{cov}(\mathbf{x}, \mathbf{y})$.

As the name suggests, sample covariance generalizes sample variance.

$$s(\mathbf{x}, \mathbf{x}) = s^2(\mathbf{x}).$$

In light of the previous theorem, we can now see that the [first computational formula](#) and the [second computational formula](#) above generalize the computational formulas for sample variance. Clearly, sample covariance is *symmetric*.

$$s(\mathbf{x}, \mathbf{y}) = s(\mathbf{y}, \mathbf{x}).$$

Sample covariance is linear in the first argument with the second argument fixed.

If \mathbf{x} , \mathbf{y} , and \mathbf{z} are data vectors from population variables x , y , and z , respectively, and if c is a constant, then

1. $s(\mathbf{x} + \mathbf{y}, \mathbf{z}) = s(\mathbf{x}, \mathbf{z}) + s(\mathbf{y}, \mathbf{z})$
2. $s(c\mathbf{x}, \mathbf{y}) = cs(\mathbf{x}, \mathbf{y})$

Proof

1. Recall that $m(\mathbf{x} + \mathbf{y}) = m(\mathbf{x}) + m(\mathbf{y})$. Hence

$$s(\mathbf{x} + \mathbf{y}, \mathbf{z}) = \frac{1}{n-1} \sum_{i=1}^n [x_i + y_i - m(\mathbf{x} + \mathbf{y})][z_i - m(\mathbf{z})] \quad (6.7.15)$$

$$= \frac{1}{n-1} \sum_{i=1}^n ([x_i - m(\mathbf{x})] + [y_i - m(\mathbf{y})])[z_i - m(\mathbf{z})] \quad (6.7.16)$$

$$= \frac{1}{n-1} \sum_{i=1}^n [x_i - m(\mathbf{x})][z_i - m(\mathbf{z})] + \frac{1}{n-1} \sum_{i=1}^n [y_i - m(\mathbf{y})][z_i - m(\mathbf{z})] \quad (6.7.17)$$

$$= s(\mathbf{x}, \mathbf{z}) + s(\mathbf{y}, \mathbf{z}) \quad (6.7.18)$$

2. Recall that $m(c\mathbf{x}) = cm(\mathbf{x})$. Hence

$$s(c\mathbf{x}, \mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n [cx_i - m(c\mathbf{x})][y_i - m(\mathbf{y})] \quad (6.7.19)$$

$$= \frac{1}{n-1} \sum_{i=1}^n [cx_i - cm(\mathbf{x})][y_i - m(\mathbf{y})] = cs(\mathbf{x}, \mathbf{y}) \quad (6.7.20)$$

By symmetry, sample covariance is also linear in the second argument with the first argument fixed, and hence is *bilinear*. The general version of the bilinear property is given in the following theorem:

Suppose that \mathbf{x}_i is a data vector from a population variable x_i for $i \in \{1, 2, \dots, k\}$ and that \mathbf{y}_j is a data vector from a population variable y_j for $j \in \{1, 2, \dots, l\}$. Suppose also that a_1, a_2, \dots, a_k and b_1, b_2, \dots, b_l are constants. Then

$$s\left(\sum_{i=1}^k a_i \mathbf{x}_i, \sum_{j=1}^l b_j \mathbf{y}_j\right) = \sum_{i=1}^k \sum_{j=1}^l a_i b_j s(\mathbf{x}_i, \mathbf{y}_j) \quad (6.7.21)$$

A special case of the bilinear property provides a nice way to compute the sample variance of a sum.

$$s^2(\mathbf{x} + \mathbf{y}) = s^2(\mathbf{x}) + 2s(\mathbf{x}, \mathbf{y}) + s^2(\mathbf{y}).$$

Proof

From the preceding results,

$$s^2(\mathbf{x} + \mathbf{y}) = s(\mathbf{x} + \mathbf{y}, \mathbf{x} + \mathbf{y}) = s(\mathbf{x}, \mathbf{x}) + s(\mathbf{x}, \mathbf{y}) + s(\mathbf{y}, \mathbf{x}) + s(\mathbf{y}, \mathbf{y}) \quad (6.7.22)$$

$$= s^2(\mathbf{x}) + 2s(\mathbf{x}, \mathbf{y}) + s^2(\mathbf{y}) \quad (6.7.23)$$

The generalization of this result to sums of three or more vectors is completely straightforward: namely, the sample variance of a sum is the sum of all of the pairwise sample covariances. Note that the sample variance of a sum can be greater than, less than, or equal to the sum of the sample variances, depending on the sign and magnitude of the pure covariance term. In particular, if the vectors are pairwise uncorrelated, then the variance of the sum is the sum of the variances.

If \mathbf{c} is a constant data set then $s(\mathbf{x}, \mathbf{c}) = 0$.

Proof

This follows directly from the definition. If $c_i = c$ for each i , then $m(\mathbf{c}) = c$ and hence $c_i - m(\mathbf{c}) = 0$ for each i .

Combining the result in the last exercise with the [bilinear property](#), we see that covariance is unchanged if constants are added to the data sets. That is, if \mathbf{c} and \mathbf{d} are constant vectors then $s(\mathbf{x} + \mathbf{c}, \mathbf{y} + \mathbf{d}) = s(\mathbf{x}, \mathbf{y})$.

Properties of Correlation

A few simple properties of correlation are given next. Most of these follow easily from the corresponding properties of covariance. First, recall that the standard scores of x_i and y_i are, respectively,

$$u_i = \frac{x_i - m(\mathbf{x})}{s(\mathbf{x})}, \quad v_i = \frac{y_i - m(\mathbf{y})}{s(\mathbf{y})} \quad (6.7.24)$$

The standard scores from a data set are dimensionless quantities that have mean 0 and variance 1.

The correlation between \mathbf{x} and \mathbf{y} is the covariance of their standard scores \mathbf{u} and \mathbf{v} . That is, $r(\mathbf{x}, \mathbf{y}) = s(\mathbf{u}, \mathbf{v})$.

Proof

In vector notation, note that

$$\mathbf{u} = \frac{1}{s(\mathbf{x})} [\mathbf{x} - m(\mathbf{x})], \quad \mathbf{v} = \frac{1}{s(\mathbf{y})} [\mathbf{y} - m(\mathbf{y})] \quad (6.7.25)$$

Hence the result follows immediately from properties of covariance:

$$s(\mathbf{u}, \mathbf{v}) = \frac{1}{s(\mathbf{x})s(\mathbf{y})} s(\mathbf{x}, \mathbf{y}) = r(\mathbf{x}, \mathbf{y}) \quad (6.7.26)$$

Correlation is symmetric.

$$r(\mathbf{x}, \mathbf{y}) = r(\mathbf{y}, \mathbf{x}).$$

Unlike covariance, correlation is unaffected by multiplying one of the data sets by a positive constant (recall that this can always be thought of as a change of scale in the underlying variable). On the other hand, multiplying a data set by a negative constant changes the sign of the correlation.

If $c \neq 0$ is a constant then

1. $r(c\mathbf{x}, \mathbf{y}) = r(\mathbf{x}, \mathbf{y})$ if $c > 0$
2. $r(c\mathbf{x}, \mathbf{y}) = -r(\mathbf{x}, \mathbf{y})$ if $c < 0$

Proof

By definition and from the [scaling property of covariance](#),

$$r(c\mathbf{x}, \mathbf{y}) = \frac{s(c\mathbf{x}, \mathbf{y})}{s(c\mathbf{x})s(\mathbf{y})} = \frac{cs(\mathbf{x}, \mathbf{y})}{|c|s(\mathbf{x})s(\mathbf{y})} = \frac{c}{|c|}r(\mathbf{x}, \mathbf{y}) \quad (6.7.27)$$

and of course, $c/|c| = 1$ if $c > 0$ and $c/|c| = -1$ if $c < 0$.

Like covariance, correlation is unaffected by adding constants to the data sets. Adding a constant to a data set often corresponds to a *change of location*.

If \mathbf{c} and \mathbf{d} are constant vectors then $r(\mathbf{x} + \mathbf{c}, \mathbf{y} + \mathbf{d}) = r(\mathbf{x}, \mathbf{y})$.

Proof

This result follows directly from the corresponding properties of covariance and standard deviation:

$$r(\mathbf{x} + \mathbf{c}, \mathbf{y} + \mathbf{d}) = \frac{s(\mathbf{x} + \mathbf{c}, \mathbf{y} + \mathbf{d})}{s(\mathbf{x} + \mathbf{c})s(\mathbf{y} + \mathbf{d})} = \frac{s(\mathbf{x}, \mathbf{y})}{s(\mathbf{x})s(\mathbf{y})} = r(\mathbf{x}, \mathbf{y}) \quad (6.7.28)$$

The last couple of properties reinforce the fact that correlation is a standardized measure of association that is not affected by changing the units of measurement. In the first Challenger data set, for example, the variables of interest are temperature at time of launch (in degrees Fahrenheit) and O-ring erosion (in millimeters). The correlation between these variables is of critical importance. If we were to measure temperature in degrees Celsius and O-ring erosion in inches, the correlation between the two variables would be unchanged.

The most important properties of correlation arise from studying the line that best fits the data, our next topic.

Linear Regression

We are interested in finding the line $y = a + bx$ that best fits the sample points $((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n))$. This is a basic and important problem in many areas of mathematics, not just statistics. We think of x as the *predictor variable* and y as the *response variable*. Thus, the term *best* means that we want to find the line (that is, find the coefficients a and b) that minimizes the average of the squared errors between the actual y values in our data and the predicted y values:

$$\text{mse}(a, b) = \frac{1}{n-1} \sum_{i=1}^n [y_i - (a + bx_i)]^2 \quad (6.7.29)$$

Note that the minimizing value of (a, b) would be the same if the function were simply the sum of the squared errors, or if we averaged by dividing by n rather than $n-1$, or if we used the square root of any of these functions. Of course that actual *minimum value* of the function would be different if we changed the function, but again, not the point (a, b) where the minimum occurs. Our particular choice of mse as the error function is best for statistical purposes. Finding (a, b) that minimize mse is a standard problem in calculus.

The graph of mse is a paraboloid opening upward. The function mse is minimized when

$$b(\mathbf{x}, \mathbf{y}) = \frac{s(\mathbf{x}, \mathbf{y})}{s^2(\mathbf{x})} \quad (6.7.30)$$

$$a(\mathbf{x}, \mathbf{y}) = m(\mathbf{y}) - b(\mathbf{x}, \mathbf{y})m(\mathbf{x}) = m(\mathbf{y}) - \frac{s(\mathbf{x}, \mathbf{y})}{s^2(\mathbf{x})}m(\mathbf{x}) \quad (6.7.31)$$

Proof

We can tell from the algebraic form of mse that the graph is a paraboloid opening upward. To find the unique point that minimizes mse, note that

$$\frac{\partial}{\partial a} \text{mse}(a, b) = \frac{1}{n-1} \sum 2[y_i - (a + bx_i)](-1) = \frac{2}{n-1} \left[-\sum_{i=1}^n y_i + na + b \sum_{i=1}^n x_i \right] \quad (6.7.32)$$

$$\frac{\partial}{\partial b} \text{mse}(a, b) = \frac{1}{n-1} \sum 2[y_i - (a + bx_i)](-x_i) = \frac{2}{n-1} \left[-\sum_{i=1}^n x_i y_i + a \sum_{i=1}^n x_i + b \sum_{i=1}^n x_i^2 \right] \quad (6.7.33)$$

Solving $\frac{\partial}{\partial a} \text{mse}(a, b) = 0$, gives $a = m(\mathbf{y}) - bm(\mathbf{x})$. Substituting this into $\frac{\partial}{\partial b} \text{mse}(a, b) = 0$ and solving for b gives

$$b = \frac{n[m(\mathbf{x}\mathbf{y}) - m(\mathbf{x})m(\mathbf{y})]}{n[m(\mathbf{x}^2) - m^2(\mathbf{x})]} \quad (6.7.34)$$

Dividing the numerator and denominator in the last expression by $n-1$ and using the computational formula above, we see that $b = s(\mathbf{x}, \mathbf{y})/s^2(\mathbf{x})$.

Of course, the optimal values of a and b are *statistics*, that is, functions of the data. Thus the *sample regression line* is

$$\mathbf{y} = m(\mathbf{y}) + \frac{s(\mathbf{x}, \mathbf{y})}{s^2(\mathbf{x})} [\mathbf{x} - m(\mathbf{x})] \quad (6.7.35)$$




Figure 6.7.3: Scatterplot with regression line

Note that the regression line passes through the point $(m(\mathbf{x}), m(\mathbf{y}))$, the center of the sample of points.




Figure 6.7.4: The regression line passes through the center

The minimum mean square error is

$$\text{mse} [a(\mathbf{x}, \mathbf{y}), b(\mathbf{x}, \mathbf{y})] = s(\mathbf{y})^2 [1 - r^2(\mathbf{x}, \mathbf{y})] \quad (6.7.36)$$

Proof

This follows from substituting $a(\mathbf{x}, \mathbf{y})$ $b(\mathbf{x}, \mathbf{y})$ into mse and simplifying.

Sample correlation and covariance satisfy the following properties.

1. $-1 \leq r(\mathbf{x}, \mathbf{y}) \leq 1$
2. $-s(\mathbf{x})s(\mathbf{y}) \leq s(\mathbf{x}, \mathbf{y}) \leq s(\mathbf{x})s(\mathbf{y})$
3. $r(\mathbf{x}, \mathbf{y}) = -1$ if and only if the sample points lie on a line with negative slope.
4. $r(\mathbf{x}, \mathbf{y}) = 1$ if and only if the sample points lie on a line with positive slope.

Proof

Note that $\text{mse} \geq 0$ and hence from the previous theorem, we must have $r^2(\mathbf{x}, \mathbf{y}) \leq 1$. This is equivalent to part (a), which in turn, from the definition of sample correlation, is equivalent to part (b). For parts (c) and (d), note that $\text{mse}(a, b) = 0$ if and only if $y_i = a + bx_i$ for each i , and moreover, $b(\mathbf{x}, \mathbf{y})$ has the same sign as $r(\mathbf{x}, \mathbf{y})$.

Thus, we now see in a deeper way that the sample covariance and correlation measure the degree of linearity of the sample points. Recall from our discussion of measures of center and spread that the constant a that minimizes

$$\text{mse}(a) = \frac{1}{n-1} \sum_{i=1}^n (y_i - a)^2 \quad (6.7.37)$$

is the sample mean $m(\mathbf{y})$, and the minimum value of the mean square error is the sample variance $s^2(\mathbf{y})$. Thus, the difference between this value of the mean square error and the one above, namely $s^2(\mathbf{y})r^2(\mathbf{x}, \mathbf{y})$ is the reduction in the variability of the y data when the linear term in x is added to the predictor. The fractional reduction is $r^2(\mathbf{x}, \mathbf{y})$,

and hence this statistics is called the (sample) *coefficient of determination*. Note that if the data vectors \mathbf{x} and \mathbf{y} are uncorrelated, then x has no value as a predictor of y ; the regression line in this case is the horizontal line $y = m(\mathbf{y})$ and the mean square error is $s^2(\mathbf{y})$.

The choice of predictor and response variables is important.

The sample regression line with predictor variable x and response variable y is not the same as the sample regression line with predictor variable y and response variable x , except in the extreme case $r(\mathbf{x}, \mathbf{y}) = \pm 1$ where the sample points all lie on a line.

Residuals

The difference between the actual y value of a data point and the value predicted by the regression line is called the *residual* of that data point. Thus, the residual corresponding to (x_i, y_i) is $d_i = y_i - \hat{y}_i$ where \hat{y}_i is the regression line at x_i :

$$\hat{y}_i = m(\mathbf{y}) + \frac{s(\mathbf{x}, \mathbf{y})}{s(\mathbf{x})^2} [x_i - m(\mathbf{x})] \quad (6.7.38)$$

Note that the predicted value \hat{y}_i and the residual d_i are *statistics*, that is, functions of the data (\mathbf{x}, \mathbf{y}) , but we are suppressing this in the notation for simplicity.

The residuals sum to 0: $\sum_{i=1}^n d_i = 0$.

Proof

This follows from the definition, and is a restatement of the fact that the regression line passes through the center of the data set $(m(\mathbf{x}), m(\mathbf{y}))$.

Various plots of the residuals can help one understand the relationship between the x and y data. Some of the more common are given in the following definition:

Residual plots

1. A plot of (i, d_i) for $i \in \{1, 2, \dots, n\}$, that is, a plot of indices versus residuals.
2. A plot of (x_i, d_i) for $i \in \{1, 2, \dots, n\}$, that is, a plot of x values versus residuals.
3. A plot of (d_i, y_i) for $i \in \{1, 2, \dots, n\}$, that is, a plot of residuals versus actual y values.
4. A plot of (d_i, \hat{y}_i) for $i \in \{1, 2, \dots, n\}$, that is a plot of residuals versus predicted y values.
5. A histogram of the residuals (d_1, d_2, \dots, d_n) .

Sums of Squares

For our next discussion, we will re-interpret the minimum mean square error formula above. Here are the new definitions:

Sums of squares

1. $\text{sst}(\mathbf{y}) = \sum_{i=1}^n [y_i - m(\mathbf{y})]^2$ is the *total sum of squares*.
2. $\text{ssr}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n [\hat{y}_i - m(\mathbf{y})]^2$ is the *regression sum of squares*
3. $\text{sse}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n (y_i - \hat{y}_i)^2$ is the *error sum of squares*.

Note that $\text{sst}(\mathbf{y})$ is simply $n - 1$ times the variance $s^2(\mathbf{y})$ and is the total of the sums of the squares of the deviations of the y values from the mean of the y values. Similarly, $\text{sse}(\mathbf{x}, \mathbf{y})$ is simply $n - 1$ times the minimum mean square error given above. Of course, $\text{sst}(\mathbf{y})$ has $n - 1$ degrees of freedom, while $\text{sse}(\mathbf{x}, \mathbf{y})$ has $n - 2$ degrees of freedom and $\text{ssr}(\mathbf{x}, \mathbf{y})$ a single degree of freedom. The total sum of squares is the sum of the regression sum of squares and the error sum of squares:

The sums of squares are related as follows:

1. $\text{ssr}(\mathbf{x}, \mathbf{y}) = r^2(\mathbf{x}, \mathbf{y}) \text{sst}(\mathbf{y})$
2. $\text{sst}(\mathbf{y}) = \text{ssr}(\mathbf{x}, \mathbf{y}) + \text{sse}(\mathbf{x}, \mathbf{y})$

Proof

By definition of sst and r , we see that $r^2(\mathbf{x}, \mathbf{y}) \text{sst}(\mathbf{y}) = s^2(\mathbf{x}, \mathbf{y}) / s^2(\mathbf{x})$. But from the regression equation,

$$[\hat{y}_i - m(\mathbf{y})]^2 = \frac{s^2(\mathbf{x}, \mathbf{y})}{s^2(\mathbf{x})} [x_i - m(\mathbf{x})]^2 \quad (6.7.39)$$

Summing over i gives

$$\text{ssr}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n [\hat{y}_i - m(\mathbf{y})]^2 = \frac{s^2(\mathbf{x}, \mathbf{y})}{s^2(\mathbf{x})} \quad (6.7.40)$$

Hence $\text{ssr}(\mathbf{x}, \mathbf{y}) = r^2(\mathbf{x}, \mathbf{y}) \text{sst}(\mathbf{y})$. Finally, multiplying the result above by $n - 1$ gives $\text{sse}(\mathbf{x}, \mathbf{y}) = \text{sst}(\mathbf{y}) - r^2(\mathbf{x}, \mathbf{y}) \text{sst}(\mathbf{y}) = \text{sst}(\mathbf{y}) - \text{ssr}(\mathbf{x}, \mathbf{y})$.

Note that $r^2(\mathbf{x}, \mathbf{y}) = \text{ssr}(\mathbf{x}, \mathbf{y}) / \text{sst}(\mathbf{y})$, so once again, $r^2(\mathbf{x}, \mathbf{y})$ is the coefficient of determination—the proportion of the variability in the y data explained by the x data. We can average sse by dividing by its degrees of freedom and then take the square root to obtain a standard error:

The *standard error of estimate* is

$$\text{se}(\mathbf{x}, \mathbf{y}) = \sqrt{\frac{\text{sse}(\mathbf{x}, \mathbf{y})}{n - 2}} \quad (6.7.41)$$

This really is a *standard* error in the same sense as a *standard* deviation. It's an average of the errors of sorts, but in the root mean square sense.

Finally, it's important to note that linear regression is a much more powerful idea than might first appear, and in fact the term *linear* can be a bit misleading. By applying various transformations to y or x or both, we can fit a variety of two-parameter curves to the given data $((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n))$. Some of the most common transformations are explored in the exercises below.

Probability Theory

We continue our discussion of sample covariance, correlation, and regression but now from the more interesting point of view that the variables are random. Specifically, suppose that we have a basic random experiment, and that X and Y are real-valued random variables for the experiment. Equivalently, (X, Y) is a random vector taking values in \mathbb{R}^2 . Let $\mu = \mathbb{E}(X)$ and $\nu = \mathbb{E}(Y)$ denote the distribution means, $\sigma^2 = \text{var}(X)$ and $\tau^2 = \text{var}(Y)$ the distribution variances, and let $\delta = \text{cov}(X, Y)$ denote the distribution covariance, so that the distribution correlation is

$$\rho = \text{cor}(X, Y) = \frac{\text{cov}(X, Y)}{\text{sd}(X)\text{sd}(Y)} = \frac{\delta}{\sigma\tau} \quad (6.7.42)$$

We will also need some higher order moments. Let $\sigma_4 = \mathbb{E}[(X - \mu)^4]$, $\tau_4 = \mathbb{E}[(Y - \nu)^4]$, and $\delta_2 = \mathbb{E}[(X - \mu)^2(Y - \nu)^2]$. Naturally, we assume that all of these moments are finite.

Now suppose that we run the basic experiment n times. This creates a compound experiment with a sequence of independent random vectors $((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n))$ each with the same distribution as (X, Y) . In statistical terms, this is a random sample of size n from the distribution of (X, Y) . The statistics discussed in previous section are well defined but now they are all random variables. We use the notation established previously, except that we use our usual convention of denoting random variables with capital letters. Of course, the deterministic properties and relations established [above](#) still hold. Note that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of X and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ is a random sample of size n from the distribution of Y . The main purpose of this subsection is to study the relationship between various statistics from \mathbf{X} and \mathbf{Y} , and to study statistics that are natural estimators of the distribution covariance and correlation.

The Sample Means

Recall that the sample means are

$$M(\mathbf{X}) = \frac{1}{n} \sum_{i=1}^n X_i, \quad M(\mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n Y_i \quad (6.7.43)$$

From the sections on the law of large numbers and the central limit theorem, we know a great deal about the distributions of $M(\mathbf{X})$ and $M(\mathbf{Y})$ *individually*. But we need to know more about the *joint distribution*.

The covariance and correlation between $M(\mathbf{X})$ and $M(\mathbf{Y})$ are

1. $\text{cov}[M(\mathbf{X}), M(\mathbf{Y})] = \delta/n$
2. $\text{cor}[M(\mathbf{X}), M(\mathbf{Y})] = \rho$

Proof

Part (a) follows from the bilinearity of the covariance operator:

$$\text{cov}\left(\frac{1}{n} \sum_{i=1}^n X_i, \frac{1}{n} \sum_{j=1}^n Y_j\right) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \text{cov}(X_i, Y_j) \quad (6.7.44)$$

By independence, the terms in the last sum are 0 if $i \neq j$. For $i = j$ the terms are $\text{cov}(X, Y) = \delta$. There are n such terms so $\text{cov}[M(\mathbf{X}), M(\mathbf{Y})] = \delta/n$. For part (b), recall that $\text{var}[M(\mathbf{X})] = \sigma^2/n$ and $\text{var}[M(\mathbf{Y})] = \tau^2/n$. Hence

$$\text{cor}[M(\mathbf{X}), M(\mathbf{Y})] = \frac{\delta/n}{(\sigma/\sqrt{n})(\tau/\sqrt{n})} = \frac{\delta}{\sigma\tau} = \rho \quad (6.7.45)$$

Note that the correlation between the sample means is the same as the correlation of the underlying sampling distribution. In particular, the correlation does not depend on the sample size n .

The Sample Variances

Recall that special versions of the sample variances, in the unlikely event that the distribution means are known, are

$$W^2(\mathbf{X}) = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2, \quad W^2(\mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n (Y_i - \nu)^2 \quad (6.7.46)$$

Once again, we have studied these statistics individually, so our emphasis now is on the joint distribution.

The covariance and correlation between $W^2(\mathbf{X})$ and $W^2(\mathbf{Y})$ are

1. $\text{cov}[W^2(\mathbf{X}), W^2(\mathbf{Y})] = (\delta_2 - \sigma^2\tau^2)/n$
2. $\text{cor}[W^2(\mathbf{X}), W^2(\mathbf{Y})] = (\delta_2 - \sigma^2\tau^2)/\sqrt{(\sigma_4 - \sigma^4)(\tau_4 - \tau^4)}$

Proof

For part (a), we use the bilinearity of the covariance operator to obtain

$$\text{cov}[W^2(\mathbf{X}), W^2(\mathbf{Y})] = \text{cov}\left(\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2, \frac{1}{n} \sum_{j=1}^n (Y_j - \nu)^2\right) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \text{cov}[(X_i - \mu)^2, (Y_j - \nu)^2] \quad (6.7.47)$$

By independence, the terms in the last sum are 0 when $i \neq j$. When $i = j$ the terms are

$$\text{cov}[(X - \mu)^2(Y - \nu)^2] = \mathbb{E}[(X - \mu)^2(Y - \nu)^2] - \mathbb{E}[(X - \mu)^2]\mathbb{E}[(Y - \nu)^2] = \delta_2 - \sigma^2\tau^2 \quad (6.7.48)$$

There are n such terms, so $\text{cov}[W^2(\mathbf{X}), W^2(\mathbf{Y})] = (\delta_2 - \sigma^2\tau^2)/n$. Part (b) follows from part (a) and the variances of $W^2(\mathbf{X})$ and $W^2(\mathbf{Y})$ from the section on Sample Variance.

Note that the correlation does not depend on the sample size n . Next, recall that the standard versions of the sample variances are

$$S^2(\mathbf{X}) = \frac{1}{n-1} \sum_{i=1}^n [X_i - M(\mathbf{X})]^2, \quad S^2(\mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n [Y_i - M(\mathbf{Y})]^2 \quad (6.7.49)$$

The covariance and correlation of the sample variances are

1. $\text{cov}[S^2(\mathbf{X}), S^2(\mathbf{Y})] = (\delta_2 - \sigma^2 \tau^2) / n + 2\delta^2 / [n(n-1)]$
2. $\text{cor}[S^2(\mathbf{X}), S^2(\mathbf{Y})] = [(n-1)(\delta_2 - \sigma^2 \tau^2) + 2\delta^2] / \sqrt{[(n-1)\sigma_4 - (n-3)\sigma^4][(n-1)\tau_4 - (n-3)\tau^4]}$

Proof

Recall that

$$S^2(\mathbf{X}) = \frac{1}{2n(n-1)} \sum_{i=1}^n \sum_{j=1}^n (X_i - X_j)^2, \quad S^2(\mathbf{Y}) = \frac{1}{2n(n-1)} \sum_{k=1}^n \sum_{l=1}^n (Y_k - Y_l)^2 \quad (6.7.50)$$

Hence using the bilinearity of the covariance operator we have

$$\text{cov}[S^2(\mathbf{X}), S^2(\mathbf{Y})] = \frac{1}{4n^2(n-1)^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{l=1}^n \text{cov}[(X_i - X_j)^2, (Y_k - Y_l)^2] \quad (6.7.51)$$

We compute the covariances in this sum by considering disjoint cases:

- $\text{cov}[(X_i - X_j)^2, (Y_k - Y_l)^2] = 0$ if $i = j$ or if $k = l$, and there are $2n^3 - n^2$ such terms.
- $\text{cov}[(X_i - X_j)^2, (Y_k - Y_l)^2] = 0$ by independence if i, j, k, l are distinct, and there are $n(n-1)(n-2)(n-3)$ such terms.
- $\text{cov}[(X_i - X_j)^2, (Y_k - Y_l)^2] = 2\delta_2 - 2\sigma^2 \tau^2 + 4\delta^2$ if $i \neq j$ and $\{k, l\} = \{i, j\}$, and there are $2n(n-1)$ such terms.
- $\text{cov}[(X_i - X_j)^2, (Y_k - Y_l)^2] = \delta_2 - \sigma^2 \tau^2$ if $i \neq j, k \neq l$, and $\#\{i, j\} \cap \{k, l\} = 1$, and there are $4n(n-1)(n-2)$ such terms.

Substituting and simplifying gives the result in (a). For (b), we use the definition of correlation and the formulas for $\text{var}[S^2(\mathbf{X})]$ and $\text{var}[S^2(\mathbf{Y})]$ from the section on the sample variance.

Asymptotically, the correlation between the sample variances is the same as the correlation between the special sample variances given above:

$$\text{cor}[S^2(\mathbf{X}), S^2(\mathbf{Y})] \rightarrow \frac{\delta_2 - \sigma^2 \tau^2}{\sqrt{(\sigma_4 - \sigma^4)(\tau_4 - \tau^4)}} \text{ as } n \rightarrow \infty \quad (6.7.52)$$

Sample Covariance

Suppose first that the distribution means μ and ν are known. As noted earlier, this is almost always an unrealistic assumption, but is still a good place to start because the analysis is very simple and the results we obtain will be useful below. A natural estimator of the distribution covariance $\delta = \text{cov}(X, Y)$ in this case is the *special sample covariance*

$$W(\mathbf{X}, \mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)(Y_i - \nu) \quad (6.7.53)$$

Note that the special sample covariance generalizes the special sample variance: $W(\mathbf{X}, \mathbf{X}) = W^2(\mathbf{X})$.

$W(\mathbf{X}, \mathbf{Y})$ is the sample mean for a random sample of size n from the distribution of $(X - \mu)(Y - \nu)$ and satisfies the following properties:

1. $\mathbb{E}[W(\mathbf{X}, \mathbf{Y})] = \delta$
2. $\text{var}[W(\mathbf{X}, \mathbf{Y})] = \frac{1}{n}(\delta_2 - \delta^2)$
3. $W(\mathbf{X}, \mathbf{Y}) \rightarrow \delta$ as $n \rightarrow \infty$ with probability 1

Proof

These results follow directly from the section on the Law of Large Numbers. For part (b), note that

$$\text{var}[(X - \mu)(Y - \nu)] = \mathbb{E}[(X - \mu)^2(Y - \nu)^2] - (\mathbb{E}[(X - \mu)(Y - \nu)])^2 = \delta_2 - \delta^2 \quad (6.7.54)$$

As an estimator of δ , part (a) means that $W(\mathbf{X}, \mathbf{Y})$ is *unbiased* and part (b) means that $W(\mathbf{X}, \mathbf{Y})$ is *consistent*.

Consider now the more realistic assumption that the distribution means μ and ν are unknown. A natural approach in this case is to average $[(X_i - M(\mathbf{X}))][Y_i - M(\mathbf{Y})]$ over $i \in \{1, 2, \dots, n\}$. But rather than dividing by n in our average, we should divide by whatever constant gives an unbiased estimator of δ . As shown in the next theorem, this constant turns out to be $n-1$, leading to the standard *sample covariance*:

$$S(\mathbf{X}, \mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n [X_i - M(\mathbf{X})][Y_i - M(\mathbf{Y})] \quad (6.7.55)$$

$\mathbb{E}[S(\mathbf{X}, \mathbf{Y})] = \delta$.

Proof

Expanding as above we have,

$$\sum_{i=1}^n [X_i - M(\mathbf{X})][Y_i - M(\mathbf{Y})] = \sum_{i=1}^n X_i Y_i - nM(\mathbf{X})M(\mathbf{Y}) \quad (6.7.56)$$

But $\mathbb{E}[X_i Y_i] = \text{cov}(X_i, Y_i) + \mathbb{E}(X_i)\mathbb{E}(Y_i) = \delta + \mu\nu$. Similarly, from the covariance of the sample means and the unbiased property, $\mathbb{E}[M(\mathbf{X})M(\mathbf{Y})] = \text{cov}[M(\mathbf{X}), M(\mathbf{Y})] + \mathbb{E}[M(\mathbf{X})]\mathbb{E}[M(\mathbf{Y})] = \delta/n + \mu\nu$. So taking expected values in the displayed equation above gives

$$\mathbb{E}\left(\sum_{i=1}^n [X_i - M(\mathbf{X})][Y_i - M(\mathbf{Y})]\right) = n(\delta + \mu\nu) - n(\delta/n + \mu\nu) = (n-1)\delta \quad (6.7.57)$$

$S(\mathbf{X}, \mathbf{Y}) \rightarrow \delta$ as $n \rightarrow \infty$ with probability 1.

Proof

Once again, we have

$$S(\mathbf{X}, \mathbf{Y}) = \frac{n}{n-1} [M(\mathbf{XY}) - M(\mathbf{X})M(\mathbf{Y})] \quad (6.7.58)$$

where $M(\mathbf{XY})$ denotes the sample mean for the sample of the products $(X_1Y_1, X_2Y_2, \dots, X_nY_n)$. By the strong law of large numbers, $M(\mathbf{X}) \rightarrow \mu$ as $n \rightarrow \infty$, $M(\mathbf{Y}) \rightarrow \nu$ as $n \rightarrow \infty$, and $M(\mathbf{XY}) \rightarrow \mathbb{E}(XY) = \delta + \mu\nu$ as $n \rightarrow \infty$, each with probability 1. So the result follows by letting $n \rightarrow \infty$ in the displayed equation.

Of course, the *sample correlation* is

$$R(\mathbf{X}, \mathbf{Y}) = \frac{S(\mathbf{X}, \mathbf{Y})}{S(\mathbf{X})S(\mathbf{Y})} \quad (6.7.59)$$

Since the sample correlation $R(\mathbf{X}, \mathbf{Y})$ is a nonlinear function of the sample covariance and sample standard deviations, it will not in general be an unbiased estimator of the distribution correlation ρ . In most cases, it would be difficult to even compute the mean and variance of $R(\mathbf{X}, \mathbf{Y})$. Nonetheless, we can show convergence of the sample correlation to the distribution correlation.

$R(\mathbf{X}, \mathbf{Y}) \rightarrow \rho$ as $n \rightarrow \infty$ with probability 1.

Proof

This follows immediately from the strong law of large numbers and previous results. From the result [above](#) $S(\mathbf{X}, \mathbf{Y}) \rightarrow \delta$ as $n \rightarrow \infty$, and from the section on the sample variance, $S(\mathbf{X}) \rightarrow \sigma$ as $n \rightarrow \infty$ and $S(\mathbf{Y}) \rightarrow \tau$ as $n \rightarrow \infty$, each with probability 1. Hence $R(\mathbf{X}, \mathbf{Y}) \rightarrow \delta/\sigma\tau = \rho$ as $n \rightarrow \infty$ with probability 1.

Our next theorem gives a formula for the variance of the sample covariance, not to be confused with the covariance of the sample variances given [above](#)!

The variance of the sample covariance is

$$\text{var}[S(\mathbf{X}, \mathbf{Y})] = \frac{1}{n} \left(\delta_2 + \frac{1}{n-1} \sigma^2 \tau^2 - \frac{n-2}{n-1} \delta^2 \right) \quad (6.7.60)$$

Proof

Recall first that

$$S(\mathbf{X}, \mathbf{Y}) = \frac{1}{2n(n-1)} \sum_{i=1}^n \sum_{j=1}^n (X_i - X_j)(Y_i - Y_j) \quad (6.7.61)$$

Hence using the bilinearity of the covariance operator we have

$$\text{var}[S(\mathbf{X}, \mathbf{Y})] = \frac{1}{4n^2(n-1)^2} \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \sum_{l=1}^n \text{cov}[(X_i - X_j)(Y_i - Y_j), (X_k - X_l)(Y_k - Y_l)] \quad (6.7.62)$$

We compute the covariances in this sum by considering disjoint cases:

- $\text{cov}[(X_i - X_j)(Y_i - Y_j), (X_k - X_l)(Y_k - Y_l)] = 0$ if $i = j$ or if $k = l$, and there are $2n^3 - n^2$ such terms.
- $\text{cov}[(X_i - X_j)(Y_i - Y_j), (X_k - X_l)(Y_k - Y_l)] = 0$ if i, j, k, l are distinct, and there are $n(n-1)(n-2)(n-3)$ such terms.
- $\text{cov}[(X_i - X_j)(Y_i - Y_j), (X_k - X_l)(Y_k - Y_l)] = 2\delta_2 + 2\sigma^2\tau^2$ if $i \neq j$ and $\{k, l\} = \{i, j\}$, and there are $2n(n-1)$ such terms.
- $\text{cov}[(X_i - X_j)(Y_i - Y_j), (X_k - X_l)(Y_k - Y_l)] = \delta_2 - \delta^2$ if $i \neq j, k \neq l$, and $\#\{i, j\} \cap \{k, l\} = 1$, and there are $4n(n-1)(n-2)$ such terms.

Substituting and simplifying gives the result

It's not surprising that the variance of the standard sample covariance (where we don't know the distribution means) is greater than the variance of the special sample covariance (where we do know the distribution means).

$\text{var}[S(\mathbf{X}, \mathbf{Y})] > \text{var}[W(\mathbf{X}, \mathbf{Y})]$

Proof

From results above, and some simple algebra,

$$\text{var}[S(\mathbf{X}, \mathbf{Y})] - \text{var}[W(\mathbf{X}, \mathbf{Y})] = \frac{1}{n(n-1)} (\delta^2 + \sigma^2\tau^2) > 0 \quad (6.7.63)$$

But note that the difference goes to 0 as $n \rightarrow \infty$.

$\text{var}[S(\mathbf{X}, \mathbf{Y})] \rightarrow 0$ as $n \rightarrow \infty$. Thus, the sample covariance is a *consistent* estimator of the distribution covariance.

Regression

In our first discussion [above](#), we studied regression from a deterministic, descriptive point of view. The results obtained applied only to the sample. Statistically more interesting and deeper questions arise when the data come from a random experiment, and we try to draw inferences about the underlying distribution from the sample regression. There are two models that commonly arise. One is where the response variable is random, but the predictor variable is deterministic. The other is the model we consider here, where the predictor variable and the response variable are both random, so that the data form a random sample from a bivariate distribution.

Thus, suppose again that we have a basic random vector (X, Y) for an experiment. Recall that in the section on (distribution) correlation and regression, we showed that the best linear predictor of Y given X , in the sense of minimizing mean square error, is the random variable

$$L(Y | X) = \mathbb{E}(Y) + \frac{\text{cov}(X, Y)}{\text{var}(X)} [X - \mathbb{E}(X)] = \nu + \frac{\delta}{\sigma^2} (X - \mu) \quad (6.7.64)$$

so that the *distribution regression line* is given by

$$y = L(Y | X = x) = \nu + \frac{\delta}{\sigma^2}(x - \mu) \quad (6.7.65)$$

Moreover, the (minimum) value of the mean square error is $\mathbb{E}\{[Y - L(Y | X)]\} = \text{var}(Y)[1 - \text{cor}^2(X, Y)] = r^2(1 - \rho^2)$.




Figure 6.7.5: The distribution regression line

Of course, in real applications, we are unlikely to know the distribution parameters μ , ν , σ^2 , and δ . If we want to estimate the distribution regression line, a natural approach would be to consider a random sample $((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n))$ from the distribution of (X, Y) and compute the sample regression line. Of course, the results are exactly the same as in the discussion [above](#), except that all of the relevant quantities are random variables. The sample regression line is

$$y = M(Y) + \frac{S(X, Y)}{S^2(X)}[x - M(X)] \quad (6.7.66)$$

The mean square error is $S^2(Y)[1 - R^2(X, Y)]$ and the coefficient of determination is $R^2(X, Y)$.




Figure 6.7.6: The distribution and sample regression lines

The fact that the sample regression line and mean square error are completely analogous to the distribution regression line and mean square error is mathematically elegant and reassuring. Again, the coefficients of the sample regression line can be viewed as estimators of the respective coefficients in the distribution regression line.

The coefficients of the sample regression line converge to the coefficients of the distribution regression line with probability 1.

1. $\frac{S(X, Y)}{S^2(X)} \rightarrow \frac{\delta}{\sigma^2}$ as $n \rightarrow \infty$
2. $M(Y) - \frac{S(X, Y)}{S^2(X)}M(X) \rightarrow \nu - \frac{\delta}{\sigma^2}\mu$ as $n \rightarrow \infty$

Proof

This follows from the strong law of large numbers and previous results. with probability 1, $S(X, Y) \rightarrow \delta$ as $n \rightarrow \infty$, $S^2(X) \rightarrow \sigma^2$ as $n \rightarrow \infty$, $M(X) \rightarrow \mu$ as $n \rightarrow \infty$, and $M(Y) \rightarrow \nu$ as $n \rightarrow \infty$.

Of course, if the linear relationship between X and Y is not strong, as measured by the sample correlation, then transformation applied to one or both variables may help. Again, some typical transformations are explored in the exercises [below](#).

Exercises

Basic Properties

Suppose that x and y are population variables, and \mathbf{x} and \mathbf{y} samples of size n from x and y respectively. Suppose also that $m(\mathbf{x}) = 3$, $m(\mathbf{y}) = -1$, $s^2(\mathbf{x}) = 4$, $s^2(\mathbf{y}) = 9$, and $s(\mathbf{x}, \mathbf{y}) = 5$. Find each of the following:

1. $r(\mathbf{x}, \mathbf{y})$
2. $m(2\mathbf{x} + 3\mathbf{y})$
3. $s^2(2\mathbf{x} + 3\mathbf{y})$
4. $s(2\mathbf{x} + 3\mathbf{y} - 1, 4\mathbf{x} + 2\mathbf{y} - 3)$

Suppose that x is the temperature (in degrees Fahrenheit) and y the resistance (in ohms) for a certain type of electronic component after 10 hours of operation. For a sample of 30 components, $m(\mathbf{x}) = 113$, $s(\mathbf{x}) = 18$, $m(\mathbf{y}) = 100$, $s(\mathbf{y}) = 10$, $r(\mathbf{x}, \mathbf{y}) = 0.6$.

1. Classify x and y by type and level of measurement.
2. Find the sample covariance.
3. Find the equation of the regression line.

Suppose now that temperature is converted to degrees Celsius (the transformation is $\frac{5}{9}(x - 32)$).

4. Find the sample means.

5. Find the sample standard deviations.
6. Find the sample covariance and correlation.
7. Find the equation of the regression line.

Answer

1. continuous, interval
2. $m = 45^\circ$, $s = 10^\circ$

Suppose that x is the length and y the width (in inches) of a leaf in a certain type of plant. For a sample of 50 leaves $m(\mathbf{x}) = 10$, $s(\mathbf{x}) = 2$, $m(\mathbf{y}) = 4$, $s(\mathbf{y}) = 1$, and $r(\mathbf{x}, \mathbf{y}) = 0.8$.

1. Classify x and y by type and level of measurement.
2. Find the sample covariance.
3. Find the equation of the regression line with x as the predictor variable and y as the response variable.

Suppose now that x and y are converted to inches (0.3937 inches per centimeter).

4. Find the sample means.
5. Find the sample standard deviations.
6. Find the sample covariance and correlation.
7. Find the equation of the regression line.

Answer

1. continuous, ratio
2. $m = 25.4$, $s = 5.08$

Scatterplot Exercises

Click in the interactive scatterplot, in various places, and watch how the means, standard deviations, correlation, and regression line change.

Click in the interactive scatterplot to define 20 points and try to come as close as possible to each of the following sample correlations:

1. 0
2. 0.5
3. -0.5
4. 0.7
5. -0.7
6. 0.9
7. -0.9.

Click in the interactive scatterplot to define 20 points. Try to generate a scatterplot in which the regression line has

1. slope 1, intercept 1
2. slope 3, intercept 0
3. slope -2, intercept 1

Simulation Exercises

Run the bivariate uniform experiment 2000 times in each of the following cases. Compare the sample means to the distribution means, the sample standard deviations to the distribution standard deviations, the sample correlation to the distribution correlation, and the sample regression line to the distribution regression line.

1. The uniform distribution on the square
2. The uniform distribution on the triangle.
3. The uniform distribution on the circle.

Run the bivariate normal experiment 2000 times for various values of the distribution standard deviations and the distribution correlation. Compare the sample means to the distribution means, the sample standard deviations to the distribution standard deviations, the sample correlation to the distribution correlation, and the sample regression line to the distribution regression line.

Transformations

Consider the function $y = a + bx^2$.

1. Sketch the graph for some representative values of a and b .
2. Note that y is a linear function of x^2 , with intercept a and slope b .
3. Hence, to fit this curve to sample data, simply apply the standard regression procedure to the data from the variables x^2 and y .

Consider the function $y = \frac{1}{a+bx}$.

1. Sketch the graph for some representative values of a and b .
2. Note that $\frac{1}{y}$ is a linear function of x , with intercept a and slope b .
3. Hence, to fit this curve to our sample data, simply apply the standard regression procedure to the data from the variables x and $\frac{1}{y}$.

Consider the function $y = \frac{x}{a+bx}$.

1. Sketch the graph for some representative values of a and b .
2. Note that $\frac{1}{y}$ is a linear function of $\frac{1}{x}$, with intercept b and slope a .

- Hence, to fit this curve to sample data, simply apply the standard regression procedure to the data from the variables $\frac{1}{x}$ and $\frac{1}{y}$.
- Note again that the names of the intercept and slope are reversed from the standard formulas.

Consider the function $y = ae^{bx}$.

- Sketch the graph for some representative values of a and b .
- Note that $\ln(y)$ is a linear function of x , with intercept $\ln(a)$ and slope b .
- Hence, to fit this curve to sample data, simply apply the standard regression procedure to the data from the variables x and $\ln(y)$.
- After solving for the intercept $\ln(a)$, recover the statistic $a = e^{\ln(a)}$.

Consider the function $y = ax^b$.

- Sketch the graph for some representative values of a and b .
- Note that $\ln(y)$ is a linear function of $\ln(x)$, with intercept $\ln(a)$ and slope b .
- Hence, to fit this curve to sample data, simply apply the standard regression procedure to the data from the variables $\ln(x)$ and $\ln(y)$.
- After solving for the intercept $\ln(a)$, recover the statistic $a = e^{\ln(a)}$.

Computational Exercises

All statistical software packages will perform regression analysis. In addition to the regression line, most packages will typically report the coefficient of determination $r^2(\mathbf{x}, \mathbf{y})$, the sums of squares $\text{sst}(\mathbf{y})$, $\text{ssr}(\mathbf{x}, \mathbf{y})$, $\text{sse}(\mathbf{x}, \mathbf{y})$, and the standard error of estimate $\text{se}(\mathbf{x}, \mathbf{y})$. Most packages will also draw the scatterplot, with the regression line superimposed, and will draw the various graphs of residuals discussed above. Many packages also provide easy ways to transform the data. Thus, there is very little reason to perform the computations by hand, except with a small data set to master the definitions and formulas. In the following problem, do the computations and draw the graphs with minimal technological aids.

Suppose that x is the number of math courses completed and y the number of science courses completed for a student at Enormous State University (ESU). A sample of 10 ESU students gives the following data: $((1, 1), (3, 3), (6, 4), (2, 1), (8, 5), (2, 2), (4, 3), (6, 4), (4, 3), (4, 4))$

- Classify x and y by type and level of measurement.
- Sketch the scatterplot.

Construct a table with rows corresponding to cases and columns corresponding to i , x_i , y_i , $x_i - m(\mathbf{x})$, $y_i - m(\mathbf{y})$, $[x_i - m(\mathbf{x})]^2$, $[y_i - m(\mathbf{y})]^2$, $[x_i - m(\mathbf{x})][y_i - m(\mathbf{y})]$, \hat{y}_i , $\hat{y}_i - m(\mathbf{y})$, $[\hat{y}_i - m(\mathbf{y})]^2$, $y_i - \hat{y}_i$, and $(y_i - \hat{y}_i)^2$. Add a row at the bottom for totals and means. Use precision arithmetic.

- Complete the first 8 columns.
- Find the sample correlation and the coefficient of determination.
- Find the sample regression equation.
- Complete the table.
- Verify the identities for the sums of squares.

Answer

i	x_i	y_i	$x_i - m(\mathbf{x})$	$y_i - m(\mathbf{y})$	$[x_i - m(\mathbf{x})]^2$	$[y_i - m(\mathbf{y})]^2$	$[x_i - m(\mathbf{x})][y_i - m(\mathbf{y})]$	\hat{y}_i	$\hat{y}_i - m(\mathbf{y})$	$[\hat{y}_i - m(\mathbf{y})]^2$	$y_i - \hat{y}_i$	$(y_i - \hat{y}_i)^2$
1	1	1	-3	-2	9	4	6	9/7	-12/7	144/49	-2/7	4/49
2	3	3	-1	0	1	0	0	17/7	-4/7	16/49	4/7	16/49
3	6	4	2	1	4	1	2	29/7	8/7	64/49	-1/7	1/49
4	2	1	-2	-2	4	4	4	13/7	-8/7	64/49	-6/7	36/49
5	8	5	4	2	16	4	8	37/7	16/7	256/49	-2/7	4/49
6	2	2	-2	-1	4	1	2	13/7	-8/7	64/49	1/7	1/49
7	4	3	0	0	0	0	0	3	0	0	0	0
8	6	4	2	1	4	1	2	29/7	8/7	64/49	-1/7	1/49
9	4	3	0	0	0	0	0	3	0	0	0	0
10	4	4	0	1	0	1	0	3	0	0	1	1
Total	40	30	0	0	42	16	24	30	0	96/7	0	16/7
Mean	4	3	0	0	14/3	16/9	8/3	3	0	96/7	0	2/7

- discrete, ratio
- $r = 2\sqrt{3/14} \approx 0.926$, $r^2 = 6/7$
- $y = 3 + \frac{4}{7}(x - 4)$
- $16 = 96/7 + 16/7$

The following two exercise should help you review some of the probability topics in this section.

Suppose that (X, Y) has a continuous distribution with probability density function $f(x, y) = 15x^2y$ for $0 \leq x \leq y \leq 1$. Find each of the following:

- $\mu = \mathbb{E}(X)$ and $\nu = \mathbb{E}(Y)$
- $\sigma^2 = \text{var}(X)$ and $\tau^2 = \text{var}(Y)$
- $\sigma_3 = \mathbb{E}[(X - \mu)^3]$ and $\tau_3 = \mathbb{E}[(Y - \nu)^3]$
- $\sigma_4 = \mathbb{E}[(X - \mu)^4]$ and $\tau_4 = \mathbb{E}[(Y - \nu)^4]$
- $\delta = \text{cov}(X, Y)$, $\rho = \text{cor}(X, Y)$, and $\delta_2 = \mathbb{E}[(X - \mu)^2(Y - \nu)^2]$

6. $L(Y | X)$ and $L(X | Y)$

Answer

1. $5/8$ $5/6$
2. $17/448$ $5/252$
3. $-5/1792$ $-5/1512$
4. $305/860165/3024$
5. $5/336$ $\sqrt{5/17}$ $1/768$
6. $L(Y | X) = \frac{10}{17} + \frac{20}{51}X$, $L(X | Y) = \frac{3}{4}Y$

Suppose now that $((X_1, Y_1), (X_2, Y_2), \dots, (X_9, Y_9))$ is a random sample of size 9 from the distribution in the previous exercise. Find each of the following:

1. $\mathbb{E}[M(\mathbf{X})]$ and $\text{var}[M(\mathbf{X})]$
2. $\mathbb{E}[M(\mathbf{Y})]$ and $\text{var}[M(\mathbf{Y})]$
3. $\text{cov}[M(\mathbf{X}), M(\mathbf{Y})]$ and $\text{cor}[M(\mathbf{X}), M(\mathbf{Y})]$
4. $\mathbb{E}[W^2(\mathbf{X})]$ and $\text{var}[W^2(\mathbf{X})]$
5. $\mathbb{E}[W^2(\mathbf{Y})]$ and $\text{var}[W^2(\mathbf{Y})]$
6. $\mathbb{E}[S^2(\mathbf{X})]$ and $\text{var}[S^2(\mathbf{X})]$
7. $\mathbb{E}[S^2(\mathbf{Y})]$ and $\text{var}[S^2(\mathbf{Y})]$
8. $\mathbb{E}[W(\mathbf{X}, \mathbf{Y})]$ and $\text{var}[W(\mathbf{X}, \mathbf{Y})]$
9. $\mathbb{E}[S(\mathbf{X}, \mathbf{Y})]$ and $\text{var}[S(\mathbf{X}, \mathbf{Y})]$

Answer

1. $5/8$ $17/4032$
2. $5/6$ $5/2268$
3. $5/3024$ $\sqrt{5/17}$
4. $17/448$ $317/1354752$
5. $5/252$ $5/35721$
6. $17/448$ $5935/21676032$
7. $5/252$ $115/762048$
8. $5/336$ $61/508032$
9. $5/336$ $181/1354752$

Data Analysis Exercises

Use statistical software for the following problems.

Consider the height variables in Pearson's height data.

1. Classify the variables by type and level of measurement.
2. Compute the correlation coefficient and the coefficient of determination.
3. Compute the least squares regression line, with the height of the father as the predictor variable and the height of the son as the response variable.
4. Draw the scatterplot and the regression line together.
5. Predict the height of a son whose father is 68 inches tall.
6. Compute the regression line if the heights are converted to centimeters (there are 2.54 centimeters per inch).

Answer

1. Continuous, ratio
2. $r = 0.501$, $r^2 = 0.251$
3. $y = 33.893 + 0.514x$
5. 68.85
6. $y = 86.088 + 0.514x$

Consider the petal length, petal width, and species variables in Fisher's iris data.

1. Classify the variables by type and level of measurement.
2. Compute the correlation between petal length and petal width.
3. Compute the correlation between petal length and petal width by species.

Answer

1. Species: discrete, nominal; petal length and width: continuous ratio
2. 0.9559
3. Setosa: 0.3316, Verginica: 0.3496, Versicolor: 0.6162

Consider the number of candies and net weight variables in the M&M data.

1. Classify the variable by type and level of measurement.
2. Compute the correlation coefficient and the coefficient of determination.
3. Compute the least squares regression line with number of candies as the predictor variable and net weight as the response variable.
4. Draw the scatterplot and the regression line in part (b) together.
5. Predict the net weight of a bag of M&Ms with 56 candies.
6. Naively, one might expect a much stronger correlation between the number of candies and the net weight in a bag of M&Ms. What is another source of variability in net weight?

Answer

1. Number of candies: discrete, ratio; net weight: continuous, ratio
2. $r = 0.793$, $r^2 = 0.629$
3. $y = 20.278 + 0.507x$
5. 48.657
6. Variability in the weight of individual candies.

Consider the response rate and total SAT score variables in the SAT by state data set.

1. Classify the variables by type and level of measurement.
2. Compute the correlation coefficient and the coefficient of determination.
3. Compute the least squares regression line with response rate as the predictor variable and SAT score as the response variable.
4. Draw the scatterplot and regression line together.
5. Give a possible explanation for the negative correlation.

Answer

1. Response rate: continuous, ratio. SAT score could probably be considered either discrete or continuous, but is only at the interval level of measurement, since the smallest possible scores is 400 (200 each on the verbal and math portions).
2. $r = -0.849$, $r^2 = 0.721$
3. $y = 1141.5 - 2.1x$
5. States with low response rate may be states for which the SAT is optional. In that case, the students who take the test are the better, college-bound students. Conversely, states with high response rates may be states for which the SAT is mandatory. In that case, all students including the weaker, non-college-bound students take the test.

Consider the verbal and math SAT scores (for all students) in the SAT by year data set.

1. Classify the variables by type and level of measurement.
2. Compute the correlation coefficient and the coefficient of determination.
3. Compute the least squares regression line.
4. Draw the scatterplot and regression line together.

Answer

1. Continuous perhaps, but only at the interval level of measurement because the smallest possible score on each part is 200.
2. $r = 0.614$, $r^2 = 0.377$
3. $y = 321.5 + 0.3x$

Consider the temperature and erosion variables in the first data set in the Challenger data.

1. Classify the variables by type and level of measurement.
2. Compute the correlation coefficient and the coefficient of determination.
3. Compute the least squares regression line.
4. Draw the scatter plot and the regression line together.
5. Predict the O-ring erosion with a temperature of 31°F .
6. Is the prediction in part (c) meaningful? Explain.
7. Find the regression line if temperature is converted to degrees Celsius. Recall that the conversion is $\frac{5}{9}(x - 32)$.

Answer

1. temperature: continuous, interval; erosion: continuous ratio
2. $r = -0.555$, $r^2 = 0.308$
3. $y = 106.8 - 1.414x$
5. 62.9.
6. This estimate is problematic, because 31° is far outside of the range of the sample data.
7. $y = 61.54 - 2.545x$

This page titled [6.7: Sample Correlation and Regression](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

6.8: Special Properties of Normal Samples

Random samples from normal distributions are the most important special cases of the topics in this chapter. As we will see, many of the results simplify significantly when the underlying sampling distribution is normal. In addition we will derive the distributions of a number of random variables constructed from normal samples that are of fundamental importance in inferential statistics.

The One Sample Model

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$. Recall that the term *random sample* means that \mathbf{X} is a sequence of independent, identically distributed random variables. Recall also that the normal distribution has probability density function

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (6.8.1)$$

In the notation that we have used elsewhere in this chapter, $\sigma_3 = \mathbb{E}[(X-\mu)^3] = 0$ (equivalently, the skewness of the normal distribution is 0) and $\sigma_4 = \mathbb{E}[(X-\mu)^4] = 3\sigma^4$ (equivalently, the kurtosis of the normal distribution is 3). Since the sample (and in particular the sample size n) is fixed in this subsection, it will be suppressed in the notation.

The Sample Mean

First recall that the sample mean is

$$M = \frac{1}{n} \sum_{i=1}^n X_i \quad (6.8.2)$$

M is normally distributed with mean and variance given by

1. $\mathbb{E}(M) = \mu$
2. $\text{var}(M) = \sigma^2/n$

Proof

This follows from basic properties of the normal distribution. Recall that the sum of independent normally distributed variables also has a normal distribution, and a linear transformation of a normally distributed variable is also normally distributed. The mean and variance of M hold in general, and were derived in the section on the Law of Large Numbers.

Of course, by the central limit theorem, the distribution of M is approximately normal, if n is large, even if the underlying sampling distribution is not normal. The *standard score* of M is given as follows:

$$Z = \frac{M - \mu}{\sigma/\sqrt{n}} \quad (6.8.3)$$

Z has the standard normal distribution.

The standard score Z associated with the sample mean M plays a critical role in constructing interval estimates and hypothesis tests for the distribution mean μ when the distribution standard deviation σ is known. The random variable Z will also appear in several derivations in this section.

The Sample Variance

The main goal of this subsection is to show that certain multiples of the two versions of the sample variance that we have studied have *chi-square* distributions. Recall that the chi-square distribution with $k \in \mathbb{N}_+$ degrees of freedom has probability density function

$$f(x) = \frac{1}{\Gamma(k/2)2^{k/2}} x^{k/2-1} e^{-x/2}, \quad 0 < x < \infty \quad (6.8.4)$$

and has mean k and variance $2k$. The moment generating function is

$$G(t) = \frac{1}{(1 - 2t)^{k/2}}, \quad -\infty < t < \frac{1}{2} \quad (6.8.5)$$

The most important result to remember is that the chi-square distribution with k degrees of freedom governs $\sum_{i=1}^k Z_i^2$, where (Z_1, Z_2, \dots, Z_k) is a sequence of independent, standard normal random variables.

Recall that if μ is known, a natural estimator of the variance σ^2 is the statistic

$$W^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 \quad (6.8.6)$$

Although the assumption that μ is known is almost always artificial, W^2 is very easy to analyze and it will be used in some of the derivations below. Our first result is the distribution of a simple multiple of W^2 . Let

$$U = \frac{n}{\sigma^2} W^2 \quad (6.8.7)$$

U has the chi-square distribution with n degrees of freedom.

Proof

Note that

$$\frac{n}{\sigma^2} W^2 = \sum_{i=1}^n \left(\frac{X_i - \mu}{\sigma} \right)^2 \quad (6.8.8)$$

and the terms in the sum are independent standard normal variables.

The variable U associated with the statistic W^2 plays a critical role in constructing interval estimates and hypothesis tests for the distribution standard deviation σ when the distribution mean μ is known (although again, this assumption is usually not realistic).

The mean and variance of W^2 are

1. $\mathbb{E}(W^2) = \sigma^2$
2. $\text{var}(W^2) = 2\sigma^4/n$

Proof

These results follow from the chi-square distribution of U and standard properties of expected value and variance.

As an estimator of σ^2 , part (a) means that W^2 is *unbiased* and part (b) means that W^2 is *consistent*. Of course, these moment results are special cases of the general results obtained in the section on Sample Variance. In that section, we also showed that M and W^2 are uncorrelated if the underlying sampling distribution has skewness 0 ($\sigma_3 = 0$), as is the case here.

Recall now that the standard version of the sample variance is the statistic

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M)^2 \quad (6.8.9)$$

The sample variance S^2 is the usual estimator of σ^2 when μ is unknown (which is usually the case). We showed earlier that in general, the sample mean M and the sample variance S^2 are uncorrelated if the underlying sampling distribution has skewness 0 ($\sigma_3 = 0$). It turns out that if the sampling distribution is normal, these variables are in fact independent, a very important and useful property, and at first blush, a very surprising result since S^2 appears to depend explicitly on M .

The sample mean M and the sample variance S^2 are independent.

Proof

The proof is based on the vector of deviations from the sample mean. Let

$$\mathbf{D} = (X_1 - M, X_2 - M, \dots, X_{n-1} - M) \quad (6.8.10)$$

Note that S^2 can be written as a function of \mathbf{D} since $\sum_{i=1}^n (X_i - M) = 0$. Next, M and the vector \mathbf{D} have a joint multivariate normal distribution. We showed earlier that M and $X_i - M$ are uncorrelated for each i , and hence it follows that M and \mathbf{D} are independent. Finally, since S^2 is a function of \mathbf{D} , it follows that M and S^2 are independent.

We can now determine the distribution of a simple multiple of the sample variance S^2 . Let

$$V = \frac{n-1}{\sigma^2} S^2 \quad (6.8.11)$$

V has the chi-square distribution with $n-1$ degrees of freedom.

Proof

We first show that $U = V + Z^2$ where U is the chi-square variable associated with W^2 and where Z is the standard score associated with M . To see this, note that

$$U = \frac{1}{\sigma^2} \sum_{i=1}^n (X_i - \mu)^2 = \frac{1}{\sigma^2} \sum_{i=1}^n (X_i - M + M - \mu)^2 \quad (6.8.12)$$

$$= \frac{1}{\sigma^2} \sum_{i=1}^n (X_i - M)^2 + \frac{2}{\sigma^2} \sum_{i=1}^n (X_i - M)(M - \mu) + \frac{1}{\sigma^2} \sum_{i=1}^n (M - \mu)^2 \quad (6.8.13)$$

In the right side of the last equation, the first term is V . The second term is 0 because $\sum_{i=1}^n (X_i - M) = 0$. The last term is $\frac{n}{\sigma^2} (M - \mu)^2 = Z^2$. Now, from the result [above](#), U has the chi-square distribution with n degrees of freedom, and of course Z^2 has the chi-square distribution with 1 degree of freedom. From the [previous result](#), V and Z^2 are independent. Recall that the moment generating function of a sum of independent variables is the product of the MGFs. Thus, taking moment generating functions in the equation $U = V + Z^2$ gives

$$\frac{1}{(1-2t)^{n/2}} = \mathbb{E}(e^{tV}) \frac{1}{(1-2t)^{1/2}}, \quad t < \frac{1}{2} \quad (6.8.14)$$

Solving we have $\mathbb{E}(e^{tV}) = 1/(1-2t)^{(n-1)/2}$ for $t < 1/2$ and therefore V has the chi-square distribution with $n-1$ degrees of freedom.

The variable V associated with the statistic S^2 plays a critical role in constructing interval estimates and hypothesis tests for the distribution standard deviation σ when the distribution mean μ is unknown (almost always the case).

The mean and variance of S^2 are

1. $\mathbb{E}(S^2) = \sigma^2$
2. $\text{var}(S^2) = 2\sigma^4/(n-1)$

Proof

These results follow from the chi-square distribution of V and standard properties of expected value and variance.

As before, these moment results are special cases of the general results obtained in the section on Sample Variance. Again, as an estimator of σ^2 , part (a) means that S^2 is unbiased, and part (b) means that S^2 is consistent. Note also that $\text{var}(S^2)$ is larger than $\text{var}(W^2)$ (not surprising), by a factor of $\frac{n}{n-1}$.

In the special distribution simulator, select the chi-square distribution. Vary the degree of freedom parameter and note the shape and location of the probability density function and the mean, standard deviation bar. For selected values of the parameter, run the experiment 1000 times and compare the empirical density function and moments to the true probability density function and moments.

The covariance and correlation between the special sample variance and the standard sample variance are

1. $\text{cov}(W^2, S^2) = 2\sigma^4/n$
2. $\text{cor}(W^2, S^2) = \sqrt{(n-1)/n}$

Proof

These results follow from general results obtained in the section on sample variance and the fact that $\sigma_4 = 3\sigma^4$.

Note that the correlation does not depend on the parameters μ and σ , and converges to 1 as $n \rightarrow \infty$,

The T Variable

Recall that the Student t distribution with $k \in \mathbb{N}_+$ degrees of freedom has probability density function

$$f(t) = C_k \left(1 + \frac{t^2}{k}\right)^{-(k+1)/2}, \quad t \in \mathbb{R} \quad (6.8.15)$$

where C_k is the appropriate normalizing constant. The distribution has mean 0 if $k > 1$ and variance $k/(k-2)$ if $k > 2$. In this subsection, the main point to remember is that the t distribution with k degrees of freedom is the distribution of

$$\frac{Z}{\sqrt{V/k}} \quad (6.8.16)$$

where Z has the standard normal distribution; V has the chi-square distribution with k degrees of freedom; and Z and V are independent. Our goal is to derive the distribution of

$$T = \frac{M - \mu}{S/\sqrt{n}} \quad (6.8.17)$$

Note that T is similar to the standard score Z associated with M , but with the sample standard deviation S replacing the distribution standard deviation σ . The variable T plays a critical role in constructing interval estimates and hypothesis tests for the distribution mean μ when the distribution standard deviation σ is unknown.

As usual, let Z denote the standard score associated with the sample mean M , and let V denote the chi-square variable associated with the sample variance S^2 . Then

$$T = \frac{Z}{\sqrt{V/(n-1)}} \quad (6.8.18)$$

and hence T has the student t distribution with $n-1$ degrees of freedom.

Proof

In the definition of T , divide the numerator and denominator by σ/\sqrt{n} . The numerator is then $(M - \mu)/(\sigma/\sqrt{n}) = Z$ and the denominator is $S/\sigma = \sqrt{V/(n-1)}$. Since Z and V are independent, Z has the standard normal distribution, and V has the chi-square distribution with $n-1$ degrees of freedom, it follows that T has the student t distribution with $n-1$ degrees of freedom.

In the special distribution simulator, select the t distribution. Vary the degree of freedom parameter and note the shape and location of the probability density function and the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the empirical density function and moments to the distribution density function and moments.

The Two Sample Model

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_m)$ is a random sample of size m from the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma \in (0, \infty)$, and that $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ is a random sample of size n from the normal distribution with mean $\nu \in \mathbb{R}$ and standard deviation $\tau \in (0, \infty)$. Finally, suppose that \mathbf{X} and \mathbf{Y} are independent. Of course, all of the results above in the [one sample model](#) apply to \mathbf{X} and \mathbf{Y} separately, but now we are interested in statistics that are helpful in inferential procedures that compare the two normal distributions. We will use the basic notation established above, but we will indicate the dependence on the sample.

The two-sample (or more generally the multi-sample model) occurs naturally when a basic variable in the statistical experiment is *filtered* according to one or more other variable (often nominal variables). For example, in the cicada data, the weights of the male cicadas and the weights of the female cicadas may fit observations from the two-sample normal model. The basic variable *weight* is

filtered by the variable *gender*. If weight is filtered by gender and species, we might have observations from the 6-sample normal model.

The Difference in the Sample Means

We know from our work [above](#) that $M(\mathbf{X})$ and $M(\mathbf{Y})$ have normal distributions. Moreover, these sample means are independent because the underlying samples \mathbf{X} and \mathbf{Y} are independent. Hence, it follows from a basic property of the normal distribution that any linear combination of $M(\mathbf{X})$ and $M(\mathbf{Y})$ will be normally distributed as well. For inferential procedures that compare the distribution means μ and ν , the linear combination that is most important is the difference.

$M(\mathbf{X}) - M(\mathbf{Y})$ has a normal distribution with mean and variance given by

1. $\mathbb{E}[(M(\mathbf{X}) - M(\mathbf{Y}))] = \mu - \nu$
2. $\text{var}[(M(\mathbf{X}) - M(\mathbf{Y}))] = \sigma^2/m + \tau^2/n$

Hence the standard score

$$Z = \frac{[(M(\mathbf{X}) - M(\mathbf{Y}))] - (\mu - \nu)}{\sqrt{\sigma^2/m + \tau^2/n}} \quad (6.8.19)$$

has the standard normal distribution. This standard score plays a fundamental role in constructing interval estimates and hypothesis test for the difference $\mu - \nu$ when the distribution standard deviations σ and τ are known.

Ratios of Sample Variances

Next we will show that the ratios of certain multiples of the sample variances (both versions) of \mathbf{X} and \mathbf{Y} have *F distributions*. Recall that the *F* distribution with $j \in \mathbb{N}_+$ degrees of freedom in the numerator and $k \in \mathbb{N}_+$ degrees of freedom in the denominator is the distribution of

$$\frac{U/j}{V/k} \quad (6.8.20)$$

where U has the chi-square distribution with j degrees of freedom; V has the chi-square distribution with k degrees of freedom; and U and V are independent. The *F* distribution is named in honor of Ronald Fisher and has probability density function

$$f(x) = C_{j,k} \frac{x^{(j-2)/2}}{[1 + (j/k)x]^{(j+k)/2}}, \quad 0 < x < \infty \quad (6.8.21)$$

where $C_{j,k}$ is the appropriate normalizing constant. The mean is $\frac{k}{k-2}$ if $k > 2$, and the variance is $2\left(\frac{k}{k-2}\right)^2 \frac{j+k-2}{j(k-4)}$ if $k > 4$.

The random variable given below has the *F* distribution with m degrees of freedom in the numerator and n degrees of freedom in the denominator:

$$\frac{W^2(\mathbf{X})/\sigma^2}{W^2(\mathbf{Y})/\tau^2} \quad (6.8.22)$$

Proof

Using the notation in the subsection on the special sample variances, note that $W^2(\mathbf{X})/\sigma^2 = U(\mathbf{X})/m$ and $W^2(\mathbf{Y})/\tau^2 = U(\mathbf{Y})/n$. The result then follows immediately since $U(\mathbf{X})$ and $U(\mathbf{Y})$ are independent chi-square variables with m and n degrees of freedom, respectively.

The random variable given below has the *F* distribution with $m - 1$ degrees of freedom in the numerator and $n - 1$ degrees of freedom in the denominator:

$$\frac{S^2(\mathbf{X})/\sigma^2}{S^2(\mathbf{Y})/\tau^2} \quad (6.8.23)$$

Proof

Using the notation in the subsection on the standard sample variances, note that $S^2(\mathbf{X})/\sigma^2 = V(\mathbf{X})/(m-1)$ and $S^2(\mathbf{Y})/\tau^2 = V(\mathbf{Y})/(n-1)$. The result then follows immediately since $V(\mathbf{X})$ and $V(\mathbf{Y})$ are independent chi-square variables with $m-1$ and $n-1$ degrees of freedom, respectively.

These variables are useful for constructing interval estimates and hypothesis tests of the ratio of the standard deviations σ/τ . The choice of the F variable depends on whether the means μ and ν are known or unknown. Usually, of course, the means are unknown and so the statistic in [above](#) is used.

In the special distribution simulator, select the F distribution. Vary the degrees of freedom parameters and note the shape and location of the probability density function and the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the empirical density function and moments to the true distribution density function and moments.

The T Variable

Our final construction in the two sample normal model will result in a variable that has the student t distribution. This variable plays a fundamental role in constructing interval estimates and hypothesis test for the difference $\mu - \nu$ when the distribution standard deviations σ and τ are unknown. The construction requires the additional assumption that the distribution standard deviations are the same: $\sigma = \tau$. This assumption is reasonable if there is an inherent variability in the measurement variables that does not change even when different treatments are applied to the objects in the population.

Note first that the standard score associated with the difference in the sample means becomes

$$Z = \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - (\nu - \mu)}{\sigma \sqrt{1/m + 1/n}} \quad (6.8.24)$$

To construct our desired variable, we first need an estimate of σ^2 . A natural approach is to consider a weighted average of the sample variances $S^2(\mathbf{X})$ and $S^2(\mathbf{Y})$, with the degrees of freedom as the weight factors (this is called the *pooled estimate* of σ^2). Thus, let

$$S^2(\mathbf{X}, \mathbf{Y}) = \frac{(m-1)S^2(\mathbf{X}) + (n-1)S^2(\mathbf{Y})}{m+n-2} \quad (6.8.25)$$

The random variable V given below has the chi-square distribution with $m+n-2$ degrees of freedom:

$$V = \frac{(m-1)S^2(\mathbf{X}) + (n-1)S^2(\mathbf{Y})}{\sigma^2} \quad (6.8.26)$$

Proof

The variable can be expressed as the sum of independent chi-square variables.

The variables $M(\mathbf{Y}) - M(\mathbf{X})$ and $S^2(\mathbf{X}, \mathbf{Y})$ are independent.

Proof

The following pairs of variables are independent: $(M(\mathbf{X}), S^2(\mathbf{X}))$ and $(M(\mathbf{Y}), S^2(\mathbf{Y}))$; $M(\mathbf{X})$ and $S^2(\mathbf{X})$; $M(\mathbf{Y})$ and $S^2(\mathbf{Y})$

The random variable T given below has the student t distribution with $m+n-2$ degrees of freedom.

$$T = \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - (\nu - \mu)}{S(\mathbf{X}, \mathbf{Y}) \sqrt{1/m + 1/n}} \quad (6.8.27)$$

Proof

The random variable can be written as $Z/\sqrt{V/(m+n-2)}$ where Z is the standard normal variable given above and V is the chi-square variable given above. Moreover, Z and V are independent by the previous result.

The Bivariate Sample Model

Suppose now that $((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n))$ is a random sample of size n from the bivariate normal distribution with means $\mu \in \mathbb{R}$ and $\nu \in \mathbb{R}$, standard deviations $\sigma \in (0, \infty)$ and $\tau \in (0, \infty)$, and correlation $\rho \in [0, 1]$. Of course, $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the normal distribution with mean μ and standard deviation σ , and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ is a random sample of size n from the normal distribution with mean ν and standard deviation τ , so the results above in the [one sample model](#) apply to \mathbf{X} and \mathbf{Y} individually. Thus our interest in this section is in the relation between various \mathbf{X} and \mathbf{Y} statistics and properties of sample covariance.

The bivariate (or more generally multivariate) model occurs naturally when considering two (or more) variables in the statistical experiment. For example, the heights of the fathers and the heights of the sons in Pearson's height data may well fit observations from the bivariate normal model.

In the notation that we have used previously, recall that $\sigma^3 = \mathbb{E}[(X - \mu)^3] = 0$, $\sigma_4 = \mathbb{E}[(X - \mu)^4] = 3\sigma^4$, $\tau_3 = \mathbb{E}[(Y - \nu)^3] = 0$, $\tau_4 = \mathbb{E}[(Y - \nu)^4] = 3\tau^4$, $\delta = \text{cov}(X, Y) = \sigma\tau\rho$, and $\delta_2 = \mathbb{E}[(X - \mu)^2(Y - \nu)^2] = \sigma^2\tau^2(1 + 2\rho^2)$.

The data vector $((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n))$ has a multivariate normal distribution.

1. The mean vector has a block form, with each block being (μ, ν) .
2. The variance-covariance matrix has a block-diagonal form, with each block being $\begin{bmatrix} \sigma^2 & \sigma\tau\rho \\ \sigma\tau\rho & \tau^2 \end{bmatrix}$.

Proof

This follows from standard results for the multivariate normal distribution. Of course the blocks in parts (a) and (b) are simply the mean and variance-covariance matrix of a single observation (X, Y) .

Sample Means

$(M(\mathbf{X}), M(\mathbf{Y}))$ has a bivariate normal distribution. The covariance and correlation are

1. $\text{cov}[M(\mathbf{X}), M(\mathbf{Y})] = \sigma\tau\rho/n$
2. $\text{cor}[M(\mathbf{X}), M(\mathbf{Y})] = \rho$

Proof

The bivariate normal distribution follows from previous result since $(M(\mathbf{X}), M(\mathbf{Y}))$ can be obtained from the data vector by a linear transformation. Parts (a) and (b) follow from our previous general results.

Of course, we know the individual means and variances of $M(\mathbf{X})$ and $M(\mathbf{Y})$ from the [one-sample model](#) above. Hence we know the complete distribution of $(M(\mathbf{X}), M(\mathbf{Y}))$.

Sample Variances

The covariance and correlation between the special sample variances are

1. $\text{cov}[W^2(\mathbf{X}), W^2(\mathbf{Y})] = 2\sigma^2\tau^2\rho^2/n$
2. $\text{cor}[W^2(\mathbf{X}), W^2(\mathbf{Y})] = \rho^2$

Proof

These results follow from our previous general results and the special form of δ_2 , σ_4 , and τ_4 .

The covariance and correlation between the standard sample variances are

1. $\text{cov}[S^2(\mathbf{X}), S^2(\mathbf{Y})] = 2\sigma^2\tau^2\rho^2/(n-1)$
2. $\text{cor}[S^2(\mathbf{X}), S^2(\mathbf{Y})] = \rho^2$

Proof

These results follow from our previous general results and the special form of δ , δ_2 , σ_4 , and τ_4 .

Sample Covariance

If μ and ν are known (again usually an artificial assumption), a natural estimator of the distribution covariance δ is the special version of the sample covariance

$$W(\mathbf{X}, \mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)(Y_i - \nu) \quad (6.8.28)$$

The mean and variance of $W(\mathbf{X}, \mathbf{Y})$ are

1. $\mathbb{E}[W(\mathbf{X}, \mathbf{Y})] = \sigma\tau\rho$
2. $\text{var}[W(\mathbf{X}, \mathbf{Y})] = \sigma^2\tau^2(1 + \rho^2)/n$

Proof

These results follow from our previous general results and the special form of δ and δ_2 .

If μ and ν are unknown (again usually the case), then a natural estimator of the distribution covariance δ is the standard sample covariance

$$S(\mathbf{X}, \mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n [X_i - M(\mathbf{X})][Y_i - M(\mathbf{Y})] \quad (6.8.29)$$

The mean and variance of the sample variance are

1. $\mathbb{E}[S(\mathbf{X}, \mathbf{Y})] = \sigma\tau\rho$
2. $\text{var}[S(\mathbf{X}, \mathbf{Y})] = \sigma^2\tau^2(1 + \rho^2)/(n-1)$

Proof

These results follow from our previous general results and the special form of δ and δ_2 .

Computational Exercises

We use the basic notation established above for samples \mathbf{X} and \mathbf{Y} , and for the statistics M , W^2 , S^2 , T , and so forth.

Suppose that the net weights (in grams) of 25 bags of M&Ms form a random sample \mathbf{X} from the normal distribution with mean 50 and standard deviation 4. Find each of the following:

1. The mean and standard deviation of M .
2. The mean and standard deviation of W^2 .
3. The mean and standard deviation of S^2 .
4. The mean and standard deviation of T .
5. $\mathbb{P}(M > 49, S^2 < 20)$.
6. $\mathbb{P}(-1 < T < 1)$.

Answer

1. 50, 4/5
2. 16, $16\sqrt{2}/5$
3. 16, $8/\sqrt{3}$
4. 0, $2\sqrt{3/11}$
5. 0.7291
6. 0.6727

Suppose that the SAT math scores from 16 Alabama students form a random sample \mathbf{X} from the normal distribution with mean 550 and standard deviation 20, while the SAT math scores from 25 Georgia students form a random sample \mathbf{Y} from the normal distribution with mean 540 and standard deviation 15. The two samples are independent. Find each of the following:

1. The mean and standard deviation of $M(\mathbf{X})$.
2. The mean and standard deviation of $M(\mathbf{Y})$.
3. The mean and standard deviation of $M(\mathbf{X}) - M(\mathbf{Y})$.

4. $\mathbb{P}[M(\mathbf{X}) > M(\mathbf{Y})]$.
5. The mean and standard deviation of $S^2(\mathbf{X})$.
6. The mean and standard deviation of $S^2(\mathbf{Y})$.
7. The mean and standard deviation of $S^2(\mathbf{X})/S^2(\mathbf{Y})$
8. $\mathbb{P}[S(\mathbf{X}) > S(\mathbf{Y})]$.

Answer

1. 550, 5
2. 540, 3
3. 10, $\sqrt{34}$
4. 0.9568
5. 400, $80\sqrt{10/3}$
6. 225, $75\sqrt{3}/2$
7. $64/33$, $\frac{32}{165}\sqrt{74/3}$
8. 0.8750

This page titled [6.8: Special Properties of Normal Samples](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

7: Point Estimation

Point estimation refers to the process of estimating a parameter from a probability distribution, based on observed data from the distribution. It is one of the core topics in mathematical statistics. In this chapter, we will explore the most common methods of point estimation: the method of moments, the method of maximum likelihood, and Bayes' estimators. We also study important properties of estimators, including sufficiency and completeness, and the basic question of whether an estimator is the best possible one.

[7.1: Estimators](#)

[7.2: The Method of Moments](#)

[7.3: Maximum Likelihood](#)

[7.4: Bayesian Estimation](#)

[7.5: Best Unbiased Estimators](#)

[7.6: Sufficient, Complete and Ancillary Statistics](#)

This page titled [7: Point Estimation](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.1: Estimators

The Basic Statistical Model

As usual, our starting point is a random experiment with an underlying sample space and a probability measure \mathbb{P} . In the basic statistical model, we have an observable random variable \mathbf{X} taking values in a set S . Recall that in general, this variable can have quite a complicated structure. For example, if the experiment is to sample n objects from a population and record various measurements of interest, then the data vector has the form

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (7.1.1)$$

where X_i is the vector of measurements for the i th object. The most important special case is when (X_1, X_2, \dots, X_n) are independent and identically distributed (IID). In this case \mathbf{X} is a random sample of size n from the distribution of an underlying measurement variable X .

Statistics

Recall also that a *statistic* is an observable function of the outcome variable of the random experiment: $\mathbf{U} = \mathbf{u}(\mathbf{X})$ where \mathbf{u} is a known function from S into another set T . Thus, a statistic is simply a random variable derived from the observation variable \mathbf{X} , with the assumption that \mathbf{U} is also observable. As the notation indicates, \mathbf{U} is typically also vector-valued. Note that the original data vector \mathbf{X} is itself a statistic, but usually we are interested in statistics derived from \mathbf{X} . A statistic \mathbf{U} may be computed to answer an inferential question. In this context, if the dimension of \mathbf{U} (as a vector) is smaller than the dimension of \mathbf{X} (as is usually the case), then we have achieved *data reduction*. Ideally, we would like to achieve significant data reduction with no loss of information about the inferential question at hand.

Parameters

In the technical sense, a *parameter* θ is a function of the *distribution* of \mathbf{X} , taking values in a *parameter space* T . Typically, the distribution of \mathbf{X} will have $k \in \mathbb{N}_+$ real parameters of interest, so that θ has the form $\theta = (\theta_1, \theta_2, \dots, \theta_k)$ and thus $T \subseteq \mathbb{R}^k$. In many cases, one or more of the parameters are unknown, and must be estimated from the data variable \mathbf{X} . This is one of the most important and basic of all statistical problems, and is the subject of this chapter. If \mathbf{U} is a statistic, then the distribution of \mathbf{U} will depend on the parameters of \mathbf{X} , and thus so will distributional constructs such as means, variances, covariances, probability density functions and so forth. We usually suppress this dependence notationally to keep our mathematical expressions from becoming too unwieldy, but it's very important to realize that the underlying dependence is present. Remember that the critical idea is that by observing a value \mathbf{u} of a statistic \mathbf{U} we (hopefully) gain information about the unknown parameters.

Estimators

Suppose now that we have an unknown real parameter θ taking values in a parameter space $T \subseteq \mathbb{R}$. A real-valued statistic $U = u(\mathbf{X})$ that is used to estimate θ is called, appropriately enough, an *estimator* of θ . Thus, the estimator is a random variable and hence has a distribution, a mean, a variance, and so on (all of which, as noted above, will generally depend on θ). When we actually run the experiment and observe the data \mathbf{x} , the observed value $u = u(\mathbf{x})$ (a single number) is the *estimate* of the parameter θ . The following definitions are basic.

Suppose that U is a statistic used as an estimator of a parameter θ with values in $T \subseteq \mathbb{R}$. For $\theta \in T$,

1. $U - \theta$ is the *error*.
2. $\text{bias}(U) = E(U - \theta) = E(U) - \theta$ is the *bias* of U
3. $\text{mse}(U) = E[(U - \theta)^2]$ is the *mean square error* of U

Thus the error is the difference between the estimator and the parameter being estimated, so of course the error is a random variable. The bias of U is simply the expected error, and the mean square error (the name says it all) is the expected square of the error. Note that bias and mean square error are functions of $\theta \in T$. The following definitions are a natural complement to the definition of bias.

Suppose again that U is a statistic used as an estimator of a parameter θ with values in $T \subseteq \mathbb{R}$.

1. U is *unbiased* if $\text{bias}(U) = 0$, or equivalently $E(U) = \theta$, for all $\theta \in T$.

2. U is *negatively biased* if $\text{bias}(U) \leq 0$, or equivalently $\mathbb{E}(U) \leq \theta$, for all $\theta \in T$.
3. U is *positively biased* if $\text{bias}(U) \geq 0$, or equivalently $\mathbb{E}(U) \geq \theta$, for all $\theta \in T$.

Thus, for an unbiased estimator, the expected value of the estimator is the parameter being estimated, clearly a desirable property. On the other hand, a positively biased estimator overestimates the parameter, on average, while a negatively biased estimator underestimates the parameter on average. Our definitions of negative and positive bias are *weak* in the sense that the weak inequalities \leq and \geq are used. There are corresponding strong definitions, of course, using the strong inequalities $<$ and $>$. Note, however, that none of these definitions may apply. For example, it might be the case that $\text{bias}(U) < 0$ for some $\theta \in T$, $\text{bias}(U) = 0$ for other $\theta \in T$, and $\text{bias}(U) > 0$ for yet other $\theta \in T$.

$$\text{mse}(U) = \text{var}(U) + \text{bias}^2(U)$$

Proof

This follows from basic properties of expected value and variance:

$$\mathbb{E}[(U - \theta)^2] = \text{var}(U - \theta) + [\mathbb{E}(U - \theta)]^2 = \text{var}(U) + \text{bias}^2(U) \quad (7.1.2)$$

In particular, if the estimator is unbiased, then the mean square error of U is simply the variance of U .

Ideally, we would like to have unbiased estimators with small mean square error. However, this is not always possible, and the result in (3) shows the delicate relationship between bias and mean square error. In the next section we will see an example with two estimators of a parameter that are multiples of each other; one is unbiased, but the other has smaller mean square error. However, if we have two unbiased estimators of θ , we naturally prefer the one with the smaller variance (mean square error).

Suppose that U and V are unbiased estimators of a parameter θ with values in $T \subseteq \mathbb{R}$.

1. U is *more efficient than* V if $\text{var}(U) \leq \text{var}(V)$.
2. The *relative efficiency* of U with respect to V is

$$\text{eff}(U, V) = \frac{\text{var}(V)}{\text{var}(U)} \quad (7.1.3)$$

Asymptotic Properties

Suppose again that we have a real parameter θ with possible values in a parameter space T . Often in a statistical experiment, we observe an infinite sequence of random variables over time, $\mathbf{X} = (X_1, X_2, \dots)$, so that at time n we have observed $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$. In this setting we often have a general formula that defines an estimator of θ for each sample size n . Technically, this gives a *sequence* of real-valued estimators of θ : $\mathbf{U} = (U_1, U_2, \dots)$ where U_n is a real-valued function of \mathbf{X}_n for each $n \in \mathbb{N}_+$. In this case, we can discuss the asymptotic properties of the estimators as $n \rightarrow \infty$. Most of the definitions are natural generalizations of the ones above.

The sequence of estimators $\mathbf{U} = (U_1, U_2, \dots)$ is *asymptotically unbiased* if $\text{bias}(U_n) \rightarrow 0$ as $n \rightarrow \infty$ for every $\theta \in T$, or equivalently, $\mathbb{E}(U_n) \rightarrow \theta$ as $n \rightarrow \infty$ for every $\theta \in T$.

Suppose that $\mathbf{U} = (U_1, U_2, \dots)$ and $\mathbf{V} = (V_1, V_2, \dots)$ are two sequences of estimators that are asymptotically unbiased. The *asymptotic relative efficiency* of \mathbf{U} to \mathbf{V} is

$$\lim_{n \rightarrow \infty} \text{eff}(U_n, V_n) = \lim_{n \rightarrow \infty} \frac{\text{var}(V_n)}{\text{var}(U_n)} \quad (7.1.4)$$

assuming that the limit exists.

Naturally, we expect our estimators to improve, as the sample size n increases, and in some sense to converge to the parameter as $n \rightarrow \infty$. This general idea is known as *consistency*. Once again, for the remainder of this discussion, we assume that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of estimators for a real-valued parameter θ , with values in the parameter space T .

Consistency

1. \mathbf{U} is consistent if $U_n \rightarrow \theta$ as $n \rightarrow \infty$ in probability for each $\theta \in T$. That is, $\mathbb{P}(|U_n - \theta| > \epsilon) \rightarrow 0$ as $n \rightarrow \infty$ for every $\epsilon > 0$ and $\theta \in T$.
2. \mathbf{U} is mean-square consistent if $\text{mse}(U_n) = \mathbb{E}[(U_n - \theta)^2] \rightarrow 0$ as $n \rightarrow \infty$ for $\theta \in T$.

Here is the connection between the two definitions:

If \mathbf{U} is mean-square consistent then \mathbf{U} is consistent.

Proof

From Markov's inequality,

$$\mathbb{P}(|U_n - \theta| > \epsilon) = \mathbb{P}[(U_n - \theta)^2 > \epsilon^2] \leq \frac{\mathbb{E}[(U_n - \theta)^2]}{\epsilon^2} \rightarrow 0 \text{ as } n \rightarrow \infty \quad (7.1.5)$$

That mean-square consistency implies simple consistency is simply a statistical version of the theorem that states that mean-square convergence implies convergence in probability. Here is another nice consequence of mean-square consistency.

If \mathbf{U} is mean-square consistent then \mathbf{U} is asymptotically unbiased.

Proof

This result follows from the fact that mean absolute error is smaller than root mean square error, which in turn is special case of a general result for norms. See the advanced section on vector spaces for more details. So, using this result and the ordinary triangle inequality for expected value we have

$$|\mathbb{E}(U_n - \theta)| \leq \mathbb{E}(|U_n - \theta|) \leq \sqrt{\mathbb{E}[(U_n - \theta)^2]} \rightarrow 0 \text{ as } n \rightarrow \infty \quad (7.1.6)$$

Hence $\mathbb{E}(U_n) \rightarrow \theta$ as $n \rightarrow \infty$ for $\theta \in T$.

In the next several subsections, we will review several basic estimation problems that were studied in the chapter on Random Samples.

Estimation in the Single Variable Model

Suppose that X is a basic real-valued random variable for an experiment, with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$. We sample from the distribution of X to produce a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent variables, each with the distribution of X . For each $n \in \mathbb{N}_+$, $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of X .

Estimating the Mean

This subsection is a review of some results obtained in the section on the Law of Large Numbers in the chapter on Random Samples. Recall that a natural estimator of the distribution mean μ is the sample mean, defined by

$$M_n = \frac{1}{n} \sum_{i=1}^n X_i, \quad n \in \mathbb{N}_+ \quad (7.1.7)$$

Properties of $\mathbf{M} = (M_1, M_2, \dots)$ as a sequence of estimators of μ .

1. $\mathbb{E}(M_n) = \mu$ so M_n is unbiased for $n \in \mathbb{N}_+$
2. $\text{var}(M_n) = \sigma^2/n$ for $n \in \mathbb{N}_+$ so \mathbf{M} is consistent.

The consistency of \mathbf{M} is simply the weak law of large numbers. Moreover, there are a number of important special cases of the results in (10). See the section on Sample Mean for the details.

Special cases of the sample mean

1. Suppose that $X = \mathbf{1}_A$, the indicator variable for an event A that has probability $\mathbb{P}(A)$. Then the sample mean for a random sample of size $n \in \mathbb{N}_+$ from the distribution of X is the *relative frequency* or *empirical probability* of A , denoted $P_n(A)$. Hence $P_n(A)$ is an unbiased estimator of $\mathbb{P}(A)$ for $n \in \mathbb{N}_+$ and $(P_n(A) : n \in \mathbb{N}_+)$ is consistent..

2. Suppose that F denotes the distribution function of a real-valued random variable Y . Then for fixed $y \in \mathbb{R}$, the *empirical distribution function* $F_n(y)$ is simply the sample mean for a random sample of size $n \in \mathbb{N}_+$ from the distribution of the indicator variable $X = \mathbf{1}(Y \leq y)$. Hence $F_n(y)$ is an unbiased estimator of $F(y)$ for $n \in \mathbb{N}_+$ and $(F_n(y) : n \in \mathbb{N}_+)$ is consistent.
3. Suppose that U is a random variable with a discrete distribution on a countable set S and f denotes the probability density function of U . Then for fixed $u \in S$, the *empirical probability density function* $f_n(u)$ is simply the sample mean for a random sample of size $n \in \mathbb{N}_+$ from the distribution of the indicator variable $X = \mathbf{1}(U = u)$. Hence $f_n(u)$ is an unbiased estimator of $f(u)$ for $n \in \mathbb{N}_+$ and $(f_n(u) : n \in \mathbb{N}_+)$ is consistent.

Estimating the Variance

This subsection is a review of some results obtained in the section on the Sample Variance in the chapter on Random Samples. We also assume that the fourth central moment $\sigma_4 = \mathbb{E}[(X - \mu)^4]$ is finite. Recall that σ_4/σ^4 is the kurtosis of X . Recall first that if μ is known (almost always an artificial assumption), then a natural estimator of σ^2 is a special version of the sample variance, defined by

$$W_n^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2, \quad n \in \mathbb{N}_+ \quad (7.1.8)$$

Properties of $\mathbf{W}^2 = (W_1^2, W_2^2, \dots)$ as a sequence of estimators of σ^2 .

1. $\mathbb{E}(W_n^2) = \sigma^2$ so W_n^2 is unbiased for $n \in \mathbb{N}_+$
2. $\text{var}(W_n^2) = \frac{1}{n}(\sigma_4 - \sigma^4)$ for $n \in \mathbb{N}_+$ so \mathbf{W}^2 is consistent.

Proof

\mathbf{W}^2 corresponds to sampling from the distribution of $(X - \mu)^2$. This distribution has mean σ^2 and variance $\sigma_4 - \sigma^4$, so the results follow immediately from theorem (10).

If μ is unknown (the more reasonable assumption), then a natural estimator of the distribution variance is the standard version of the sample variance, defined by

$$S_n^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M_n)^2, \quad n \in \{2, 3, \dots\} \quad (7.1.9)$$

Properties of $\mathbf{S}^2 = (S_2^2, S_3^2, \dots)$ as a sequence of estimators of σ^2

1. $\mathbb{E}(S_n^2) = \sigma^2$ so S_n^2 is unbiased for $n \in \{2, 3, \dots\}$
2. $\text{var}(S_n^2) = \frac{1}{n} \left(\sigma_4 - \frac{n-3}{n-1} \sigma^4 \right)$ for $n \in \{2, 3, \dots\}$ so \mathbf{S}^2 is consistent sequence.

Naturally, we would like to compare the sequences \mathbf{W}^2 and \mathbf{S}^2 as estimators of σ^2 . But again remember that \mathbf{W}^2 only makes sense if μ is known.

Comparison of \mathbf{W}^2 and \mathbf{S}^2

1. $\text{var}(W_n^2) < \text{var}(S_n^2)$ for $n \in \{2, 3, \dots\}$.
2. The asymptotic relative efficiency of \mathbf{W}^2 to \mathbf{S}^2 is 1.

So by (a) W_n^2 is better than S_n^2 for $n \in \{2, 3, \dots\}$, assuming that μ is known so that we can actually use W_n^2 . This is perhaps not surprising, but by (b) S_n^2 works just about as well as W_n^2 for a large sample size n . Of course, the sample standard deviation S_n is a natural estimator of the distribution standard deviation σ . Unfortunately, this estimator is biased. Here is a more general result:

Suppose that θ is a parameter with possible values in $T \subseteq (0, \infty)$ (with at least two points) and that U is a statistic with values in T . If U^2 is an unbiased estimator of θ^2 then U is a negatively biased estimator of θ .

Proof

Note that

$$\text{var}(U) = \mathbb{E}(U^2) - [\mathbb{E}(U)]^2 = \theta^2 - [\mathbb{E}(U)]^2, \quad \theta \in T \quad (7.1.10)$$

Since T has at least two points, U cannot be deterministic so $\text{var}(U) > 0$. It follows that $[\mathbb{E}(U)]^2 < \theta^2$ so $\mathbb{E}(U) < \theta$ for $\theta \in T$.

Thus, we should not be too obsessed with the unbiased property. For most sampling distributions, there will be no statistic U with the property that U is an unbiased estimator of σ and U^2 is an unbiased estimator of σ^2 .

Estimation in the Bivariate Model

In this subsection we review some of the results obtained in the section on the Correlation and Regression in the chapter on Random Samples

Suppose that X and Y are real-valued random variables for an experiment, so that (X, Y) has a bivariate distribution in \mathbb{R}^2 . Let $\mu = \mathbb{E}(X)$ and $\sigma^2 = \text{var}(X)$ denote the mean and variance of X , and let $\nu = \mathbb{E}(Y)$ and $\tau^2 = \text{var}(Y)$ denote the mean and variance of Y . For the bivariate parameters, let $\delta = \text{cov}(X, Y)$ denote the distribution covariance and $\rho = \text{cor}(X, Y)$ the distribution correlation. We need one higher-order moment as well: let $\delta_2 = \mathbb{E}[(X - \mu)^2(Y - \nu)^2]$, and as usual, we assume that all of the parameters exist. So the general parameter spaces are $\mu, \nu \in \mathbb{R}$, $\sigma^2, \tau^2 \in (0, \infty)$, $\delta \in \mathbb{R}$, and $\rho \in [0, 1]$. Suppose now that we sample from the distribution of (X, Y) to generate a sequence of independent variables $((X_1, Y_1), (X_2, Y_2), \dots)$, each with the distribution of (X, Y) . As usual, we will let $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ and $\mathbf{Y}_n = (Y_1, Y_2, \dots, Y_n)$; these are random samples of size n from the distributions of X and Y , respectively.

Since we now have two underlying variables, we need to enhance our notation somewhat. It will help to define the deterministic versions of our statistics. So if $\mathbf{x} = (x_1, x_2, \dots)$ and $\mathbf{y} = (y_1, y_2, \dots)$ are sequences of real numbers and $n \in \mathbb{N}_+$, we define the mean and special covariance functions by

$$m_n(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n x_i$$

$$w_n(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)(y_i - \nu)$$

If $n \in \{2, 3, \dots\}$ we define the variance and standard covariance functions by

$$s_n^2(\mathbf{x}) = \frac{1}{n-1} \sum_{i=1}^n [x_i - m_n(\mathbf{x})]^2$$

$$s_n(\mathbf{x}, \mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^n [x_i - m_n(\mathbf{x})][y_i - m_n(\mathbf{y})]$$

It should be clear from context whether we are using the one argument or two argument version of s_n . On this point, note that $s_n(\mathbf{x}, \mathbf{x}) = s_n^2(\mathbf{x})$.

Estimating the Covariance

If μ and ν are known (almost always an artificial assumption), then a natural estimator of the distribution covariance δ is a special version of the sample covariance, defined by

$$W_n = w_n(\mathbf{X}, \mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)(Y_i - \nu), \quad n \in \mathbb{N}_+ \quad (7.1.11)$$

Properties of $\mathbf{W} = (W_1, W_2, \dots)$ as a sequence of estimators of δ .

1. $\mathbb{E}(W_n) = \delta$ so W_n is unbiased for $n \in \mathbb{N}_+$.
2. $\text{var}(W_n) = \frac{1}{n}(\delta_2 - \delta^2)$ for $n \in \mathbb{N}_+$ so \mathbf{W} is consistent.

Proof

We've done this proof before, but it's so basic that it's worth repeating. Note that \mathbf{W} corresponds to sampling from the distribution of $(X - \mu)(Y - \nu)$. This distribution has mean δ and variance $\delta_2 - \delta^2$, so the results follow immediately from Theorem (10).

If μ and ν are unknown (usually the more reasonable assumption), then a natural estimator of the distribution covariance δ is the standard version of the sample covariance, defined by

$$S_n = s_n(\mathbf{X}, \mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n [X_i - m_n(\mathbf{X})][Y_i - m_n(\mathbf{Y})], \quad n \in \{2, 3, \dots\} \quad (7.1.12)$$

Properties of $\mathbf{S} = (S_2, S_3, \dots)$ as a sequence of estimators of δ .

1. $\mathbb{E}(S_n) = \delta$ so S_n is unbiased for $n \in \{2, 3, \dots\}$.
2. $\text{var}(S_n) = \frac{1}{n} \left(\delta_2 + \frac{1}{n-1} \sigma^2 \tau^2 - \frac{n-2}{n-1} \delta^2 \right)$ for $n \in \{2, 3, \dots\}$ so \mathbf{S} is consistent.

Once again, since we have two competing sequences of estimators of δ , we would like to compare them.

Comparison of \mathbf{W} and \mathbf{S} as estimators of δ :

1. $\text{var}(W_n) < \text{var}(S_n)$ for $n \in \{2, 3, \dots\}$.
2. The asymptotic relative efficiency of \mathbf{W} to \mathbf{S} is 1.

Thus, U_n is better than V_n for $n \in \{2, 3, \dots\}$, assuming that μ and ν are known so that we can actually use W_n . But for large n , V_n works just about as well as U_n .

Estimating the Correlation

A natural estimator of the distribution correlation ρ is the sample correlation

$$R_n = \frac{s_n(\mathbf{X}, \mathbf{Y})}{s_n(\mathbf{X})s_n(\mathbf{Y})}, \quad n \in \{2, 3, \dots\} \quad (7.1.13)$$

Note that this statistics is a nonlinear function of the sample covariance and the two sample standard deviations. For most distributions of (X, Y) , we have no hope of computing the bias or mean square error of this estimator. If we *could* compute the expected value, we would probably find that the estimator is biased. On the other hand, even though we cannot compute the mean square error, a simple application of the law of large numbers shows that $R_n \rightarrow \rho$ as $n \rightarrow \infty$ with probability 1. Thus, $\mathbf{R} = (R_2, R_3, \dots)$ is at least consistent.

Estimating the regression coefficients

Recall that the distribution regression line, with X as the predictor variable and Y as the response variable, is $y = a + bx$ where

$$a = \mathbb{E}(Y) - \frac{\text{cov}(X, Y)}{\text{var}(X)} \mathbb{E}(X), \quad b = \frac{\text{cov}(X, Y)}{\text{var}(X)} \quad (7.1.14)$$

On the other hand, the sample regression line, based on the sample of size $n \in \{2, 3, \dots\}$, is $y = A_n + B_n x$ where

$$A_n = m_n(\mathbf{Y}) - \frac{s_n(\mathbf{X}, \mathbf{Y})}{s_n^2(\mathbf{X})} m_n(\mathbf{X}), \quad B_n = \frac{s_n(\mathbf{X}, \mathbf{Y})}{s_n^2(\mathbf{X})} \quad (7.1.15)$$

Of course, the statistics A_n and B_n are natural estimators of the parameters a and b , respectively, and in a sense are derived from our previous estimators of the distribution mean, variance, and covariance. Once again, for most distributions of (X, Y) , it would be difficult to compute the bias and mean square errors of these estimators. But applications of the law of large numbers show that with probability 1, $A_n \rightarrow a$ and $B_n \rightarrow b$ as $n \rightarrow \infty$, so at least $\mathbf{A} = (A_2, A_3, \dots)$ and $\mathbf{B} = (B_2, B_3, \dots)$ are consistent.

Exercises and Special Cases

The Poisson Distribution

Let's consider a simple example that illustrates some of the ideas above. Recall that the *Poisson distribution* with parameter $\lambda \in (0, \infty)$ has probability density function g given by

$$g(x) = e^{-\lambda} \frac{\lambda^x}{x!}, \quad x \in \mathbb{N} \quad (7.1.16)$$

The Poisson distribution is often used to model the number of random “points” in a region of time or space, and is studied in more detail in the chapter on the Poisson process. The parameter λ is proportional to the size of the region of time or space; the proportionality constant is the average *rate* of the random points. The distribution is named for Simeon Poisson.

Suppose that X has the Poisson distribution with parameter λ . Hence

1. $\mu = \mathbb{E}(X) = \lambda$
2. $\sigma^2 = \text{var}(X) = \lambda$
3. $\sigma_4 = \mathbb{E}[(X - \lambda)^4] = 3\lambda^2 + \lambda$

Proof

Recall the permutation notation $x^{(n)} = x(x-1)\cdots(x-n+1)$ for $x \in \mathbb{R}$ and $n \in \mathbb{N}$. The expected value $\mathbb{E}[X^{(n)}]$ is the *factorial moment* of X of order n . It's easy to see that the factorial moments are $\mathbb{E}[X^{(n)}] = \lambda^n$ for $n \in \mathbb{N}$. The results follow from this.

Suppose now that we sample from the distribution of X to produce a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots)$, each having the Poisson distribution with unknown parameter $\lambda \in (0, \infty)$. Again, $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is a random sample of size $n \in \mathbb{N}_+$ from the distribution for each $n \in \mathbb{N}$. From the previous exercise, λ is both the mean and the variance of the distribution, so that we could use either the sample mean M_n or the sample variance S_n^2 as an estimator of λ . Both are unbiased, so which is better? Naturally, we use mean square error as our criterion.

Comparison of M to S^2 as estimators of λ .

1. $\text{var}(M_n) = \frac{\lambda}{n}$ for $n \in \mathbb{N}_+$.
2. $\text{var}(S_n^2) = \frac{\lambda}{n} \left(1 + 2\lambda \frac{n}{n-1}\right)$ for $n \in \{2, 3, \dots\}$.
3. $\text{var}(M_n) < \text{var}(S_n^2)$ so M_n for $n \in \{2, 3, \dots\}$.
4. The asymptotic relative efficiency of M to S^2 is $1 + 2\lambda$.

So our conclusion is that the sample mean M_n is a better estimator of the parameter λ than the sample variance S_n^2 for $n \in \{2, 3, \dots\}$, and the difference in quality increases with λ .

Run the Poisson experiment 100 times for several values of the parameter. In each case, compute the estimators M and S^2 . Which estimator seems to work better?

The emission of elementary particles from a sample of radioactive material in a time interval is often assumed to follow the Poisson distribution. Thus, suppose that the alpha emissions data set is a sample from a Poisson distribution. Estimate the rate parameter λ .

1. using the sample mean
2. using the sample variance

Answer

1. 8.367
2. 8.649

Simulation Exercises

In the sample mean experiment, set the sampling distribution to gamma. Increase the sample size with the scroll bar and note graphically and numerically the unbiased and consistent properties. Run the experiment 1000 times and compare the sample mean to the distribution mean.

Run the normal estimation experiment 1000 times for several values of the parameters.

1. Compare the empirical bias and mean square error of M with the theoretical values.
2. Compare the empirical bias and mean square error of S^2 and of W^2 to their theoretical values. Which estimator seems to work better?

In matching experiment, the random variable is the number of matches. Run the simulation 1000 times and compare

1. the sample mean to the distribution mean.
2. the empirical density function to the probability density function.

Run the exponential experiment 1000 times and compare the sample standard deviation to the distribution standard deviation.

Data Analysis Exercises

For Michelson's velocity of light data, compute the sample mean and sample variance.

Answer

852.4, 6242.67

For Cavendish's density of the earth data, compute the sample mean and sample variance.

Answer

5.448, 0.048817

For Short's parallax of the sun data, compute the sample mean and sample variance.

Answer

8.616, 0.561032

Consider the Cicada data.

1. Compute the sample mean and sample variance of the body length variable.
2. Compute the sample mean and sample variance of the body weight variable.
3. Compute the sample covariance and sample correlation between the body length and body weight variables.

Answer

1. 24.0, 3.92
2. 0.180, 0.003512
3. 0.0471, 0.4012

Consider the M&M data.

1. Compute the sample mean and sample variance of the net weight variable.
2. Compute the sample mean and sample variance of the total number of candies.
3. Compute the sample covariance and sample correlation between the number of candies and the net weight.

Answer

1. 57.1, 5.68
2. 49.215, 2.3163
3. 2.878, 0.794

Consider the Pearson data.

1. Compute the sample mean and sample variance of the height of the father.
2. Compute the sample mean and sample variance of the height of the son.
3. Compute the sample covariance and sample correlation between the height of the father and height of the son.

Answer

1. 67.69, 7.5396
2. 68.68, 7.9309
3. 3.875, 0.501

The estimators of the mean, variance, and covariance that we have considered in this section have been natural in a sense. However, for other parameters, it is not clear how to even find a reasonable estimator in the first place. In the next several sections, we will consider the problem of constructing estimators. Then we return to the study of the mathematical properties of estimators, and consider the question of when we can know that an estimator is the best possible, given the data.

This page titled [7.1: Estimators](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.2: The Method of Moments

Basic Theory

The Method

Suppose that we have a basic random experiment with an observable, real-valued random variable X . The distribution of X has k unknown real-valued parameters, or equivalently, a parameter vector $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_k)$ taking values in a parameter space, a subset of \mathbb{R}^k . As usual, we repeat the experiment n times to generate a random sample of size n from the distribution of X .

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (7.2.1)$$

Thus, \mathbf{X} is a sequence of independent random variables, each with the distribution of X . The *method of moments* is a technique for constructing estimators of the parameters that is based on matching the *sample moments* with the corresponding *distribution moments*. First, let

$$\mu^{(j)}(\boldsymbol{\theta}) = \mathbb{E}(X^j), \quad j \in \mathbb{N}_+ \quad (7.2.2)$$

so that $\mu^{(j)}(\boldsymbol{\theta})$ is the j th moment of X about 0. Note that we are emphasizing the dependence of these moments on the vector of parameters $\boldsymbol{\theta}$. Note also that $\mu^{(1)}(\boldsymbol{\theta})$ is just the mean of X , which we usually denote simply by μ . Next, let

$$M^{(j)}(\mathbf{X}) = \frac{1}{n} \sum_{i=1}^n X_i^j, \quad j \in \mathbb{N}_+ \quad (7.2.3)$$

so that $M^{(j)}(\mathbf{X})$ is the j th sample moment about 0. Equivalently, $M^{(j)}(\mathbf{X})$ is the sample mean for the random sample $(X_1^j, X_2^j, \dots, X_n^j)$ from the distribution of X^j . Note that we are emphasizing the dependence of the sample moments on the sample \mathbf{X} . Note also that $M^{(1)}(\mathbf{X})$ is just the ordinary sample mean, which we usually just denote by M (or by M_n if we wish to emphasize the dependence on the sample size). From our previous work, we know that $M^{(j)}(\mathbf{X})$ is an unbiased and consistent estimator of $\mu^{(j)}(\boldsymbol{\theta})$ for each j . Here's how the method works:

To construct the *method of moments* estimators (W_1, W_2, \dots, W_k) for the parameters $(\theta_1, \theta_2, \dots, \theta_k)$ respectively, we consider the equations

$$\mu^{(j)}(W_1, W_2, \dots, W_k) = M^{(j)}(X_1, X_2, \dots, X_n) \quad (7.2.4)$$

consecutively for $j \in \mathbb{N}_+$ until we are able to solve for (W_1, W_2, \dots, W_k) in terms of $(M^{(1)}, M^{(2)}, \dots)$.

The equations for $j \in \{1, 2, \dots, k\}$ give k equations in k unknowns, so there is hope (but no guarantee) that the equations can be solved for (W_1, W_2, \dots, W_k) in terms of $(M^{(1)}, M^{(2)}, \dots, M^{(k)})$. In fact, sometimes we need equations with $j > k$. [Exercise 28](#) below gives a simple example. The method of moments can be extended to parameters associated with bivariate or more general multivariate distributions, by matching sample product moments with the corresponding distribution product moments. The method of moments also sometimes makes sense when the sample variables (X_1, X_2, \dots, X_n) are not independent, but at least are identically distributed. The [hypergeometric model](#) below is an example of this.

Of course, the method of moments estimators depend on the sample size $n \in \mathbb{N}_+$. We have suppressed this so far, to keep the notation simple. But in the applications below, we put the notation back in because we want to discuss asymptotic behavior.

Estimates for the Mean and Variance

Estimating the mean and variance of a distribution are the simplest applications of the method of moments. Throughout this subsection, we assume that we have a basic real-valued random variable X with $\mu = \mathbb{E}(X) \in \mathbb{R}$ and $\sigma^2 = \text{var}(X) \in (0, \infty)$. Occasionally we will also need $\sigma_4 = \mathbb{E}[(X - \mu)^4]$, the fourth central moment. We sample from the distribution of X to produce a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent variables, each with the distribution of X . For each $n \in \mathbb{N}_+$, $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of X . We start by estimating the mean, which is essentially trivial by this method.

Suppose that the mean μ is unknown. The method of moments estimator of μ based on \mathbf{X}_n is the sample mean

$$M_n = \frac{1}{n} \sum_{i=1}^n X_i \quad (7.2.5)$$

1. $\mathbb{E}(M_n) = \mu$ so M_n is unbiased for $n \in \mathbb{N}_+$
2. $\text{var}(M_n) = \sigma^2/n$ for $n \in \mathbb{N}_+$ so $\mathbf{M} = (M_1, M_2, \dots)$ is consistent.

Proof

It does not get any more basic than this. The method of moments works by matching the distribution mean with the sample mean. The fact that $\mathbb{E}(M_n) = \mu$ and $\text{var}(M_n) = \sigma^2/n$ for $n \in \mathbb{N}_+$ are properties that we have seen several times before.

Estimating the variance of the distribution, on the other hand, depends on whether the distribution mean μ is known or unknown. First we will consider the more realistic case when the mean is also unknown. Recall that for $n \in \{2, 3, \dots\}$, the sample variance based on \mathbf{X}_n is

$$S_n^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M_n)^2 \quad (7.2.6)$$

Recall also that $\mathbb{E}(S_n^2) = \sigma^2$ so S_n^2 is unbiased for $n \in \{2, 3, \dots\}$, and that $\text{var}(S_n^2) = \frac{1}{n} \left(\sigma_4 - \frac{n-3}{n-1} \sigma^4 \right)$ so $\mathbf{S}^2 = (S_2^2, S_3^2, \dots)$ is consistent.

Suppose that the mean μ and the variance σ^2 are both unknown. For $n \in \mathbb{N}_+$, the method of moments estimator of σ^2 based on \mathbf{X}_n is

$$T_n^2 = \frac{1}{n} \sum_{i=1}^n (X_i - M_n)^2 \quad (7.2.7)$$

1. $\text{bias}(T_n^2) = -\sigma^2/n$ for $n \in \mathbb{N}_+$ so $\mathbf{T}^2 = (T_1^2, T_2^2, \dots)$ is asymptotically unbiased.
2. $\text{mse}(T_n^2) = \frac{1}{n^3} [(n-1)^2 \sigma_4 - (n^2 - 5n + 3) \sigma^4]$ for $n \in \mathbb{N}_+$ so \mathbf{T}^2 is consistent.

Proof

As before, the method of moments estimator of the distribution mean μ is the sample mean M_n . On the other hand, $\sigma^2 = \mu^{(2)} - \mu^2$ and hence the method of moments estimator of σ^2 is $T_n^2 = M_n^{(2)} - M_n^2$, which simplifies to the result above. Note that $T_n^2 = \frac{n-1}{n} S_n^2$ for $n \in \{2, 3, \dots\}$.

1. Note that $\mathbb{E}(T_n^2) = \frac{n-1}{n} \mathbb{E}(S_n^2) = \frac{n-1}{n} \sigma^2$, so $\text{bias}(T_n^2) = \frac{n-1}{n} \sigma^2 - \sigma^2 = -\frac{1}{n} \sigma^2$.
2. Recall that $\text{mse}(T_n^2) = \text{var}(T_n^2) + \text{bias}^2(T_n^2)$. But $\text{var}(T_n^2) = \left(\frac{n-1}{n} \right)^2 \text{var}(S_n^2)$. The result follows from substituting $\text{var}(S_n^2)$ given above and $\text{bias}(T_n^2)$ in part (a).

Hence T_n^2 is negatively biased and on average underestimates σ^2 . Because of this result, T_n^2 is referred to as the *biased sample variance* to distinguish it from the ordinary (unbiased) sample variance S_n^2 .

Next let's consider the usually unrealistic (but mathematically interesting) case where the mean is known, but not the variance.

Suppose that the mean μ is known and the variance σ^2 unknown. For $n \in \mathbb{N}_+$, the method of moments estimator of σ^2 based on \mathbf{X}_n is

$$W_n^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 \quad (7.2.8)$$

1. $\mathbb{E}(W_n^2) = \sigma^2$ so W_n^2 is unbiased for $n \in \mathbb{N}_+$
2. $\text{var}(W_n^2) = \frac{1}{n} (\sigma_4 - \sigma^4)$ for $n \in \mathbb{N}_+$ so $\mathbf{W}^2 = (W_1^2, W_2^2, \dots)$ is consistent.

Proof

These results follow since W_n^2 is the sample mean corresponding to a random sample of size n from the distribution of $(X - \mu)^2$.

We compared the sequence of estimators S^2 with the sequence of estimators W^2 in the introductory section on Estimators. Recall that $\text{var}(W_n^2) < \text{var}(S_n^2)$ for $n \in \{2, 3, \dots\}$ but $\text{var}(S_n^2)/\text{var}(W_n^2) \rightarrow 1$ as $n \rightarrow \infty$. There is no simple, general relationship between $\text{mse}(T_n^2)$ and $\text{mse}(S_n^2)$ or between $\text{mse}(T_n^2)$ and $\text{mse}(W_n^2)$, but the asymptotic relationship is simple.

$\text{mse}(T_n^2)/\text{mse}(W_n^2) \rightarrow 1$ and $\text{mse}(T_n^2)/\text{mse}(S_n^2) \rightarrow 1$ as $n \rightarrow \infty$

Proof

In light of the previous remarks, we just have to prove one of these limits. The first limit is simple, since the coefficients of σ_4 and σ^4 in $\text{mse}(T_n^2)$ are asymptotically $1/n$ as $n \rightarrow \infty$.

It also follows that if both μ and σ^2 are unknown, then the method of moments estimator of the standard deviation σ is $T = \sqrt{T^2}$. In the unlikely event that μ is known, but σ^2 unknown, then the method of moments estimator of σ is $W = \sqrt{W^2}$.

Estimating Two Parameters

There are several important special distributions with two parameters; some of these are included in the [computational exercises](#) below. With two parameters, we can derive the method of moments estimators by matching the distribution mean and variance with the sample mean and variance, rather than matching the distribution mean and second moment with the sample mean and second moment. This alternative approach sometimes leads to easier equations. To setup the notation, suppose that a distribution on \mathbb{R} has parameters a and b . We sample from the distribution to produce a sequence of independent variables $\mathbf{X} = (X_1, X_2, \dots)$, each with the common distribution. For $n \in \mathbb{N}_+$, $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution. Let M_n , $M_n^{(2)}$, and T_n^2 denote the sample mean, second-order sample mean, and biased sample variance corresponding to \mathbf{X}_n , and let $\mu(a, b)$, $\mu^{(2)}(a, b)$, and $\sigma^2(a, b)$ denote the mean, second-order mean, and variance of the distribution.

If the method of moments estimators U_n and V_n of a and b , respectively, can be found by solving the first two equations

$$\mu(U_n, V_n) = M_n, \quad \mu^{(2)}(U_n, V_n) = M_n^{(2)} \quad (7.2.9)$$

then U_n and V_n can also be found by solving the equations

$$\mu(U_n, V_n) = M_n, \quad \sigma^2(U_n, V_n) = T_n^2 \quad (7.2.10)$$

Proof

Recall that $\sigma^2(a, b) = \mu^{(2)}(a, b) - \mu^2(a, b)$. In addition, $T_n^2 = M_n^{(2)} - M_n^2$. Hence the equations $\mu(U_n, V_n) = M_n$, $\sigma^2(U_n, V_n) = T_n^2$ are equivalent to the equations $\mu(U_n, V_n) = M_n$, $\mu^{(2)}(U_n, V_n) = M_n^{(2)}$.

Because of this result, the biased sample variance T_n^2 will appear in many of the estimation problems for special distributions that we consider below.

Special Distributions

The Normal Distribution

The *normal distribution* with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$ is a continuous distribution on \mathbb{R} with probability density function g given by

$$g(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (7.2.11)$$

This is one of the most important distributions in probability and statistics, primarily because of the central limit theorem. The normal distribution is studied in more detail in the chapter on Special Distributions.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the normal distribution with mean μ and variance σ^2 . From our general work [above](#), we know that if μ is unknown then the sample mean M is the method of moments estimator of μ , and if in addition, σ^2 is unknown then the method of moments estimator of σ^2 is T^2 . On the other hand, in the unlikely event that μ is known then W^2 is the method of moments estimator of σ^2 . Our goal is to see how the comparisons above simplify for the normal distribution.

Mean square errors of S_n^2 and T_n^2 .

1. $\text{mse}(T^2) = \frac{2n-1}{n^2} \sigma^4$
2. $\text{mse}(S^2) = \frac{2}{n-1} \sigma^4$
3. $\text{mse}(T^2) < \text{mse}(S^2)$ for $n \in \{2, 3, \dots\}$

Proof

Recall that for the normal distribution, $\sigma_4 = 3\sigma^4$. Substituting this into the general results gives parts (a) and (b). Part (c) follows from (a) and (b). Of course the asymptotic relative efficiency is still 1, from our previous theorem.

Thus, S^2 and T^2 are multiplies of one another; S^2 is unbiased, but when the sampling distribution is normal, T^2 has smaller mean square error. Surprisingly, T^2 has smaller mean square error even than W^2 .

Mean square errors of T^2 and W^2 .

1. $\text{mse}(W^2) = \frac{2}{n} \sigma^4$
2. $\text{mse}(T^2) < \text{mse}(W^2)$ for $n \in \{2, 3, \dots\}$

Proof

Again, since the sampling distribution is normal, $\sigma_4 = 3\sigma^4$. Substituting this into the general formula for $\text{var}(W_n^2)$ gives part (a).

Run the normal estimation experiment 1000 times for several values of the sample size n and the parameters μ and σ . Compare the empirical bias and mean square error of S^2 and of T^2 to their theoretical values. Which estimator is better in terms of bias? Which estimator is better in terms of mean square error?

Next we consider estimators of the standard deviation σ . As noted in the general discussion above, $T = \sqrt{T^2}$ is the method of moments estimator when μ is unknown, while $W = \sqrt{W^2}$ is the method of moments estimator in the unlikely event that μ is known. Another natural estimator, of course, is $S = \sqrt{S^2}$, the usual sample standard deviation. The following sequence, defined in terms of the gamma function turns out to be important in the analysis of all three estimators.

Consider the sequence

$$a_n = \sqrt{\frac{2}{n} \frac{\Gamma[(n+1)/2]}{\Gamma(n/2)}}, \quad n \in \mathbb{N}_+ \quad (7.2.12)$$

Then $0 < a_n < 1$ for $n \in \mathbb{N}_+$ and $a_n \uparrow 1$ as $n \uparrow \infty$.

First, assume that μ is known so that W_n is the method of moments estimator of σ .

For $n \in \mathbb{N}_+$,

1. $\mathbb{E}(W) = a_n \sigma$
2. $\text{bias}(W) = (a_n - 1) \sigma$
3. $\text{var}(W) = (1 - a_n^2) \sigma^2$
4. $\text{mse}(W) = 2(1 - a_n) \sigma^2$

Proof

Recall that $U^2 = nW^2/\sigma^2$ has the chi-square distribution with n degrees of freedom, and hence U has the chi distribution with n degrees of freedom. Solving gives

$$W = \frac{\sigma}{\sqrt{n}} U \quad (7.2.13)$$

From the formulas for the mean and variance of the chi distribution we have

$$\begin{aligned}\mathbb{E}(W) &= \frac{\sigma}{\sqrt{n}} \mathbb{E}(U) = \frac{\sigma}{\sqrt{n}} \sqrt{2} \frac{\Gamma[(n+1)/2]}{\Gamma(n/2)} = \sigma a_n \\ \text{var}(W) &= \frac{\sigma^2}{n} \text{var}(U) = \frac{\sigma^2}{n} \{n - [\mathbb{E}(U)]^2\} = \sigma^2 (1 - a_n^2)\end{aligned}$$

Thus W is negatively biased as an estimator of σ but asymptotically unbiased and consistent. Of course we know that in general (regardless of the underlying distribution), W^2 is an unbiased estimator of σ^2 and so W is negatively biased as an estimator of σ . In the normal case, since a_n involves no unknown parameters, the statistic W/a_n is an unbiased estimator of σ . Next we consider the usual sample standard deviation S .

For $n \in \{2, 3, \dots\}$,

1. $\mathbb{E}(S) = a_{n-1} \sigma$
2. $\text{bias}(S) = (a_{n-1} - 1) \sigma$
3. $\text{var}(S) = (1 - a_{n-1}^2) \sigma^2$
4. $\text{mse}(S) = 2(1 - a_{n-1}) \sigma^2$

Proof

Recall that $V^2 = (n-1)S^2/\sigma^2$ has the chi-square distribution with $n-1$ degrees of freedom, and hence V has the chi distribution with $n-1$ degrees of freedom. The proof now proceeds just as in the previous theorem, but with $n-1$ replacing n .

As with W , the statistic S is negatively biased as an estimator of σ but asymptotically unbiased, and also consistent. Since a_{n-1} involves no unknown parameters, the statistic S/a_{n-1} is an unbiased estimator of σ . Note also that, in terms of bias and mean square error, S with sample size n behaves like W with sample size $n-1$. Finally we consider T , the method of moments estimator of σ when μ is unknown.

For $n \in \{2, 3, \dots\}$,

1. $\mathbb{E}(T) = \sqrt{\frac{n-1}{n}} a_{n-1} \sigma$
2. $\text{bias}(T) = \left(\sqrt{\frac{n-1}{n}} a_{n-1} - 1 \right) \sigma$
3. $\text{var}(T) = \frac{n-1}{n} (1 - a_{n-1}^2) \sigma^2$
4. $\text{mse}(T) = \left(2 - \frac{1}{n} - 2 \sqrt{\frac{n-1}{n}} a_{n-1} \right) \sigma^2$

Proof

The results follow easily from the previous theorem since $T_n = \sqrt{\frac{n-1}{n}} S_n$.

The Bernoulli Distribution

Recall that an *indicator variable* is a random variable X that takes only the values 0 and 1. The distribution of X is known as the *Bernoulli distribution*, named for Jacob Bernoulli, and has probability density function g given by

$$g(x) = p^x (1-p)^{1-x}, \quad x \in \{0, 1\} \quad (7.2.14)$$

where $p \in (0, 1)$ is the *success parameter*. The mean of the distribution is p and the variance is $p(1-p)$.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Bernoulli distribution with unknown success parameter p . Since the mean of the distribution is p , it follows from our general work [above](#) that the method of moments estimator of p is M , the sample mean. In this case, the sample \mathbf{X} is a sequence of Bernoulli trials, and M has a scaled version of the binomial distribution with parameters n and p :

$$\mathbb{P}\left(M = \frac{k}{n}\right) = \binom{n}{k} p^k (1-p)^{n-k}, \quad k \in \{0, 1, \dots, n\} \quad (7.2.15)$$

Note that since $X^k = X$ for every $k \in \mathbb{N}_+$, it follows that $\mu^{(k)} = p$ and $M^{(k)} = M$ for every $k \in \mathbb{N}_+$. So any of the method of moments equations would lead to the sample mean M as the estimator of p . Although very simple, this is an important application,

since Bernoulli trials are found embedded in all sorts of estimation problems, such as empirical probability density functions and empirical distribution functions.

The Geometric Distribution

The *geometric distribution* on \mathbb{N}_+ with success parameter $p \in (0, 1)$ has probability density function g given by

$$g(x) = p(1-p)^{x-1}, \quad x \in \mathbb{N}_+ \quad (7.2.16)$$

The geometric distribution on \mathbb{N}_+ governs the number of trials needed to get the first success in a sequence of Bernoulli trials with success parameter p . The mean of the distribution is $\mu = 1/p$.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the geometric distribution on \mathbb{N}_+ with unknown success parameter p . The method of moments estimator of p is

$$U = \frac{1}{M} \quad (7.2.17)$$

Proof

The method of moments equation for U is $1/U = M$.

The geometric distribution on \mathbb{N} with success parameter $p \in (0, 1)$ has probability density function

$$g(x) = p(1-p)^x, \quad x \in \mathbb{N} \quad (7.2.18)$$

This version of the geometric distribution governs the number of failures before the first success in a sequence of Bernoulli trials. The mean of the distribution is $\mu = (1-p)/p$.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the geometric distribution on \mathbb{N} with unknown parameter p . The method of moments estimator of p is

$$U = \frac{1}{M+1} \quad (7.2.19)$$

Proof

The method of moments equation for U is $(1-U)/U = M$.

The Negative Binomial Distribution

More generally, the *negative binomial distribution* on \mathbb{N} with shape parameter $k \in (0, \infty)$ and success parameter $p \in (0, 1)$ has probability density function

$$g(x) = \binom{x+k-1}{k-1} p^k (1-p)^x, \quad x \in \mathbb{N} \quad (7.2.20)$$

If k is a positive integer, then this distribution governs the number of failures before the k th success in a sequence of Bernoulli trials with success parameter p . However, the distribution makes sense for general $k \in (0, \infty)$. The negative binomial distribution is studied in more detail in the chapter on Bernoulli Trials. The mean of the distribution is $k(1-p)/p$ and the variance is $k(1-p)/p^2$. Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the negative binomial distribution on \mathbb{N} with shape parameter k and success parameter p

If k and p are unknown, then the corresponding method of moments estimators U and V are

$$U = \frac{M^2}{T^2 - M}, \quad V = \frac{M}{T^2} \quad (7.2.21)$$

Proof

Matching the distribution mean and variance to the sample mean and variance gives the equations

$$U \frac{1-V}{V} = M, \quad U \frac{1-V}{V^2} = T^2 \quad (7.2.22)$$

As usual, the results are nicer when one of the parameters is known.

Suppose that k is known but p is unknown. The method of moments estimator V_k of p is

$$V_k = \frac{k}{M + k} \quad (7.2.23)$$

Proof

Matching the distribution mean to the sample mean gives the equation

$$k \frac{1 - V_k}{V_k} = M \quad (7.2.24)$$

Suppose that k is unknown but p is known. The method of moments estimator of k is

$$U_p = \frac{p}{1 - p} M \quad (7.2.25)$$

1. $\mathbb{E}(U_p) = k$ so U_p is unbiased.
2. $\text{var}(U_p) = \frac{k}{n(1-p)}$ so U_p is consistent.

Proof

Matching the distribution mean to the sample mean gives the equation $U_p \frac{1-p}{p} = M$.

1. $E(U_p) = \frac{p}{1-p} E(M)$ and $E(M) = \frac{1-p}{p} k$
2. $\text{var}(U_p) = \left(\frac{p}{1-p}\right)^2 \text{var}(M)$ and $\text{var}(M) = \frac{1}{n} \text{var}(X) = \frac{1-p}{np^2}$

The Poisson Distribution

The *Poisson distribution* with parameter $r \in (0, \infty)$ is a discrete distribution on \mathbb{N} with probability density function g given by

$$g(x) = e^{-r} \frac{r^x}{x!}, \quad x \in \mathbb{N} \quad (7.2.26)$$

The mean and variance are both r . The distribution is named for Simeon Poisson and is widely used to model the number of “random points” in a region of time or space. The parameter r is proportional to the size of the region, with the proportionality constant playing the role of the average *rate* at which the points are distributed in time or space. The Poisson distribution is studied in more detail in the chapter on the Poisson Process.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Poisson distribution with parameter r . Since r is the mean, it follows from our general work [above](#) that the method of moments estimator of r is the sample mean M .

The Gamma Distribution

The *gamma distribution* with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ is a continuous distribution on $(0, \infty)$ with probability density function g given by

$$g(x) = \frac{1}{\Gamma(k)b^k} x^{k-1} e^{-x/b}, \quad x \in (0, \infty) \quad (7.2.27)$$

The gamma probability density function has a variety of shapes, and so this distribution is used to model various types of positive random variables. The gamma distribution is studied in more detail in the chapter on Special Distributions. The mean is $\mu = kb$ and the variance is $\sigma^2 = kb^2$.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the gamma distribution with shape parameter k and scale parameter b .

Suppose that k and b are both unknown, and let U and V be the corresponding method of moments estimators. Then

$$U = \frac{M^2}{T^2}, \quad V = \frac{T^2}{M} \quad (7.2.28)$$

Proof

Matching the distribution mean and variance with the sample mean and variance leads to the equations $UV = M$, $UV^2 = T^2$. Solving gives the results.

The method of moments estimators of k and b given in the previous exercise are complicated, nonlinear functions of the sample mean M and the sample variance T^2 . Thus, computing the bias and mean square errors of these estimators are difficult problems that we will not attempt. However, we can judge the quality of the estimators empirically, through simulations.

When one of the parameters is known, the method of moments estimator of the other parameter is much simpler.

Suppose that k is unknown, but b is known. The method of moments estimator of k is

$$U_b = \frac{M}{b} \quad (7.2.29)$$

1. $\mathbb{E}(U_b) = k$ so U_b is unbiased.
2. $\text{var}(U_b) = k/n$ so U_b is consistent.

Proof

If b is known, then the method of moments equation for U_b is $bU_b = M$. Solving gives (a). Next, $\mathbb{E}(U_b) = \mathbb{E}(M)/b = kb/b = k$, so U_b is unbiased. Finally $\text{var}(U_b) = \text{var}(M)/b^2 = kb^2/(nb^2) = k/n$.

Suppose that b is unknown, but k is known. The method of moments estimator of b is

$$V_k = \frac{M}{k} \quad (7.2.30)$$

1. $\mathbb{E}(V_k) = b$ so V_k is unbiased.
2. $\text{var}(V_k) = b^2/kn$ so that V_k is consistent.

Proof

If k is known, then the method of moments equation for V_k is $kV_k = M$. Solving gives (a). Next, $\mathbb{E}(V_k) = \mathbb{E}(M)/k = kb/k = b$, so V_k is unbiased. Finally $\text{var}(V_k) = \text{var}(M)/k^2 = kb^2/(nk^2) = b^2/kn$.

Run the gamma estimation experiment 1000 times for several different values of the sample size n and the parameters k and b . Note the empirical bias and mean square error of the estimators U , V , U_b , and V_k . One would think that the estimators when one of the parameters is known should work better than the corresponding estimators when both parameters are unknown; but investigate this question empirically.

The Beta Distribution

The *beta distribution* with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$ is a continuous distribution on $(0, 1)$ with probability density function g given by

$$g(x) = \frac{1}{B(a, b)} x^{a-1} (1-x)^{b-1}, \quad 0 < x < 1 \quad (7.2.31)$$

The beta probability density function has a variety of shapes, and so this distribution is widely used to model various types of random variables that take values in bounded intervals. The beta distribution is studied in more detail in the chapter on Special Distributions. The first two moments are $\mu = \frac{a}{a+b}$ and $\mu^{(2)} = \frac{a(a+1)}{(a+b)(a+b+1)}$.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the beta distribution with left parameter a and right parameter b .

Suppose that a and b are both unknown, and let U and V be the corresponding method of moments estimators. Then

$$U = \frac{M(M - M^{(2)})}{M^{(2)} - M^2}, \quad V = \frac{(1 - M)(M - M^{(2)})}{M^{(2)} - M^2} \quad (7.2.32)$$

Proof

The method of moments equations for U and V are

$$\frac{U}{U+V} = M, \quad \frac{U(U+1)}{(U+V)(U+V+1)} = M^{(2)} \quad (7.2.33)$$

Solving gives the result.

The method of moments estimators of a and b given in the previous exercise are complicated nonlinear functions of the sample moments M and $M^{(2)}$. Thus, we will not attempt to determine the bias and mean square errors analytically, but you will have an opportunity to explore them empirically through a simulation.

Suppose that a is unknown, but b is known. Let U_b be the method of moments estimator of a . Then

$$U_b = b \frac{M}{1-M} \quad (7.2.34)$$

Proof

If b is known then the method of moments equation for U_b as an estimator of a is $U_b / (U_b + b) = M$. Solving for U_b gives the result.

Suppose that b is unknown, but a is known. Let V_a be the method of moments estimator of b . Then

$$V_a = a \frac{1-M}{M} \quad (7.2.35)$$

Proof

If a is known then the method of moments equation for V_a as an estimator of b is $a / (a + V_a) = M$. Solving for V_a gives the result.

Run the beta estimation experiment 1000 times for several different values of the sample size n and the parameters a and b . Note the empirical bias and mean square error of the estimators U , V , U_b , and V_a . One would think that the estimators when one of the parameters is known should work better than the corresponding estimators when both parameters are unknown; but investigate this question empirically.

The following problem gives a distribution with just one parameter but the second moment equation from the method of moments is needed to derive an estimator.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the *symmetric* beta distribution, in which the left and right parameters are equal to an unknown value $c \in (0, \infty)$. The method of moments estimator of c is

$$U = \frac{2M^{(2)}}{1-4M^{(2)}} \quad (7.2.36)$$

Proof

Note that the mean μ of the symmetric distribution is $\frac{1}{2}$, independently of c , and so the first equation in the method of moments is useless. However, matching the second distribution moment to the second sample moment leads to the equation

$$\frac{U+1}{2(2U+1)} = M^{(2)} \quad (7.2.37)$$

Solving gives the result.

The Pareto Distribution

The *Pareto distribution* with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ is a continuous distribution on (b, ∞) with probability density function g given by

$$g(x) = \frac{ab^a}{x^{a+1}}, \quad b \leq x < \infty \quad (7.2.38)$$

The Pareto distribution is named for Vilfredo Pareto and is a highly skewed and *heavy-tailed* distribution. It is often used to model income and certain other types of positive random variables. The Pareto distribution is studied in more detail in the chapter on Special Distributions. If $a > 2$, the first two moments of the Pareto distribution are $\mu = \frac{ab}{a-1}$ and $\mu^{(2)} = \frac{ab^2}{a-2}$.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Pareto distribution with shape parameter $a > 2$ and scale parameter $b > 0$.

Suppose that a and b are both unknown, and let U and V be the corresponding method of moments estimators. Then

$$U = 1 + \sqrt{\frac{M^{(2)}}{M^{(2)} - M^2}} \quad (7.2.39)$$

$$V = \frac{M^{(2)}}{M} \left(1 - \sqrt{\frac{M^{(2)} - M^2}{M^{(2)}}} \right) \quad (7.2.40)$$

Proof

The method of moments equations for U and V are

$$\frac{UV}{U-1} = M \quad (7.2.41)$$

$$\frac{UV^2}{U-2} = M^{(2)} \quad (7.2.42)$$

Solving for U and V gives the results.

As with our previous examples, the method of moments estimators are complicated nonlinear functions of M and $M^{(2)}$, so computing the bias and mean square error of the estimator is difficult. Instead, we can investigate the bias and mean square error empirically, through a simulation.

Run the Pareto estimation experiment 1000 times for several different values of the sample size n and the parameters a and b . Note the empirical bias and mean square error of the estimators U and V .

When one of the parameters is known, the method of moments estimator for the other parameter is simpler.

Suppose that a is unknown, but b is known. Let U_b be the method of moments estimator of a . Then

$$U_b = \frac{M}{M-b} \quad (7.2.43)$$

Proof

If b is known then the method of moment equation for U_b as an estimator of a is $bU_b/(U_b - 1) = M$. Solving for U_b gives the result.

Suppose that b is unknown, but a is known. Let V_a be the method of moments estimator of b . Then

$$V_a = \frac{a-1}{a} M \quad (7.2.44)$$

1. $\mathbb{E}(V_a) = b$ so V_a is unbiased.
2. $\text{var}(V_a) = \frac{b^2}{na(a-2)}$ so V_a is consistent.

Proof

If a is known then the method of moments equation for V_a as an estimator of b is $aV_a/(a-1) = M$. Solving for V_a gives

(a). Next, $\mathbb{E}(V_a) = \frac{a-1}{a} \mathbb{E}(M) = \frac{a-1}{a} \frac{ab}{a-1} = b$ so V_a is unbiased. Finally,

$$\text{var}(V_a) = \left(\frac{a-1}{a}\right)^2 \text{var}(M) = \frac{(a-1)^2}{a^2} \frac{ab^2}{n(a-1)^2(a-2)} = \frac{b^2}{na(a-2)}.$$

The Uniform Distribution

The (continuous) *uniform distribution* with location parameter $a \in \mathbb{R}$ and scale parameter $h \in (0, \infty)$ has probability density function g given by

$$g(x) = \frac{1}{h}, \quad x \in [a, a+h] \quad (7.2.45)$$

The distribution models a point chosen “at random” from the interval $[a, a+h]$. The mean of the distribution is $\mu = a + \frac{1}{2}h$ and the variance is $\sigma^2 = \frac{1}{12}h^2$. The uniform distribution is studied in more detail in the chapter on Special Distributions. Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the uniform distribution.

Suppose that a and h are both unknown, and let U and V denote the corresponding method of moments estimators. Then

$$U = 2M - \sqrt{3}T, \quad V = 2\sqrt{3}T \quad (7.2.46)$$

Proof

Matching the distribution mean and variance to the sample mean and variance leads to the equations $U + \frac{1}{2}V = M$ and $\frac{1}{12}V^2 = T^2$. Solving gives the result.

As usual, we get nicer results when one of the parameters is known.

Suppose that a is known and h is unknown, and let V_a denote the method of moments estimator of h . Then

$$V_a = 2(M - a) \quad (7.2.47)$$

1. $\mathbb{E}(V_a) = h$ so V is unbiased.
2. $\text{var}(V_a) = \frac{h^2}{3n}$ so V_a is consistent.

Proof

Matching the distribution mean to the sample mean leads to the equation $a + \frac{1}{2}V_a = M$. Solving gives the result.

1. $\mathbb{E}(V_a) = 2[\mathbb{E}(M) - a] = 2(a + h/2 - a) = h$
2. $\text{var}(V_a) = 4\text{var}(M) = \frac{h^2}{3n}$

Suppose that h is known and a is unknown, and let U_h denote the method of moments estimator of a . Then

$$U_h = M - \frac{1}{2}h \quad (7.2.48)$$

1. $\mathbb{E}(U_h) = a$ so U_h is unbiased.
2. $\text{var}(U_h) = \frac{h^2}{12n}$ so U_h is consistent.

Proof

Matching the distribution mean to the sample mean leads to the equation $U_h + \frac{1}{2}h = M$. Solving gives the result.

1. $\mathbb{E}(U_h) = \mathbb{E}(M) - \frac{1}{2}h = a + \frac{1}{2}h - \frac{1}{2}h = a$
2. $\text{var}(U_h) = \text{var}(M) = \frac{h^2}{12n}$

The Hypergeometric Model

Our basic assumption in the method of moments is that the sequence of observed random variables $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from a distribution. However, the method makes sense, at least in some cases, when the variables are identically distributed but dependent. In the *hypergeometric model*, we have a population of N objects with r of the objects *type 1* and the remaining $N - r$ objects *type 0*. The parameter N , the *population size*, is a positive integer. The parameter r , the *type 1 size*, is a

nonnegative integer with $r \leq N$. These are the basic parameters, and typically one or both is unknown. Here are some typical examples:

1. The objects are devices, classified as *good* or *defective*.
2. The objects are persons, classified as *female* or *male*.
3. The objects are voters, classified as *for* or *against* a particular candidate.
4. The objects are wildlife or a particular type, either *tagged* or *untagged*.

We sample n objects from the population at random, without replacement. Let X_i be the type of the i th object selected, so that our sequence of observed variables is $\mathbf{X} = (X_1, X_2, \dots, X_n)$. The variables are identically distributed indicator variables, with $P(X_i = 1) = r/N$ for each $i \in \{1, 2, \dots, n\}$, but are dependent since the sampling is without replacement. The number of type 1 objects in the sample is $Y = \sum_{i=1}^n X_i$. This statistic has the *hypergeometric distribution* with parameter N , r , and n , and has probability density function given by

$$P(Y = y) = \frac{\binom{r}{y} \binom{N-r}{n-y}}{\binom{N}{n}} = \binom{n}{y} \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad y \in \{\max\{0, N-n+r\}, \dots, \min\{n, r\}\} \quad (7.2.49)$$

The hypergeometric model is studied in more detail in the chapter on Finite Sampling Models.

As above, let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be the observed variables in the hypergeometric model with parameters N and r . Then

1. The method of moments estimator of $p = r/N$ is $M = Y/n$, the sample mean.
2. The method of moments estimator of r with N known is $U = NM = NY/n$.
3. The method of moments estimator of N with r known is $V = r/M = rn/Y$ if $Y > 0$.

Proof

These results all follow simply from the fact that $\mathbb{E}(X) = \mathbb{P}(X = 1) = r/N$.

In the voter example (3) above, typically N and r are both unknown, but we would only be interested in estimating the ratio $p = r/N$. In the reliability example (1), we might typically know N and would be interested in estimating r . In the wildlife example (4), we would typically know r and would be interested in estimating N . This example is known as the *capture-recapture* model.

Clearly there is a close relationship between the hypergeometric model and the [Bernoulli trials model](#) above. In fact, if the sampling is *with* replacement, the Bernoulli trials model would apply rather than the hypergeometric model. In addition, if the population size N is large compared to the sample size n , the hypergeometric model is well approximated by the Bernoulli trials model.

This page titled [7.2: The Method of Moments](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.3: Maximum Likelihood

Basic Theory

The Method

Suppose again that we have an observable random variable \mathbf{X} for an experiment, that takes values in a set S . Suppose also that distribution of \mathbf{X} depends on an unknown parameter θ , taking values in a parameter space Θ . Of course, our data variable \mathbf{X} will almost always be vector valued. The parameter θ may also be vector valued. We will denote the probability density function of \mathbf{X} on S by f_θ for $\theta \in \Theta$. The distribution of \mathbf{X} could be discrete or continuous.

The likelihood function is the function obtained by reversing the roles of \mathbf{x} and θ in the probability density function; that is, we view θ as the variable and \mathbf{x} as the given information (which is precisely the point of view in estimation).

The *likelihood function* at $\mathbf{x} \in S$ is the function $L_{\mathbf{x}} : \Theta \rightarrow [0, \infty)$ given by

$$L_{\mathbf{x}}(\theta) = f_\theta(\mathbf{x}), \quad \theta \in \Theta \quad (7.3.1)$$

In the method of *maximum likelihood*, we try to find the value of the parameter that maximizes the likelihood function for each value of the data vector.

Suppose that the maximum value of $L_{\mathbf{x}}$ occurs at $u(\mathbf{x}) \in \Theta$ for each $\mathbf{x} \in S$. Then the statistic $u(\mathbf{X})$ is a *maximum likelihood estimator* of θ .

The method of maximum likelihood is intuitively appealing—we try to find the value of the parameter that would have most likely produced the data we in fact observed.

Since the natural logarithm function is strictly increasing on $(0, \infty)$, the maximum value of the likelihood function, if it exists, will occur at the same points as the maximum value of the logarithm of the likelihood function.

The *log-likelihood function* at $\mathbf{x} \in S$ is the function $\ln L_{\mathbf{x}}$:

$$\ln L_{\mathbf{x}}(\theta) = \ln f_\theta(\mathbf{x}), \quad \theta \in \Theta \quad (7.3.2)$$

If the maximum value of $\ln L_{\mathbf{x}}$ occurs at $u(\mathbf{x}) \in \Theta$ for each $\mathbf{x} \in S$. Then the statistic $u(\mathbf{X})$ is a maximum likelihood estimator of θ .

The log-likelihood function is often easier to work with than the likelihood function (typically because the probability density function $f_\theta(\mathbf{x})$ has a product structure).

Vector of Parameters

An important special case is when $\theta = (\theta_1, \theta_2, \dots, \theta_k)$ is a vector of k real parameters, so that $\Theta \subseteq \mathbb{R}^k$. In this case, the maximum likelihood problem is to maximize a function of several variables. If Θ is a continuous set, the methods of calculus can be used. If the maximum value of $L_{\mathbf{x}}$ occurs at a point θ in the interior of Θ , then $L_{\mathbf{x}}$ has a local maximum at θ . Therefore, assuming that the likelihood function is differentiable, we can find this point by solving

$$\frac{\partial}{\partial \theta_i} L_{\mathbf{x}}(\theta) = 0, \quad i \in \{1, 2, \dots, k\} \quad (7.3.3)$$

or equivalently

$$\frac{\partial}{\partial \theta_i} \ln L_{\mathbf{x}}(\theta) = 0, \quad i \in \{1, 2, \dots, k\} \quad (7.3.4)$$

On the other hand, the maximum value may occur at a boundary point of Θ , or may not exist at all.

Random Samples

The most important special case is when the data variables form a random sample from a distribution.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of a random variable X taking values in R , with probability density function g_θ for $\theta \in \Theta$. Then \mathbf{X} takes values in $S = R^n$, and the likelihood and log-likelihood functions for $\mathbf{x} = (x_1, x_2, \dots, x_n) \in S$ are

$$L_{\mathbf{x}}(\theta) = \prod_{i=1}^n g_\theta(x_i), \quad \theta \in \Theta$$

$$\ln L_{\mathbf{x}}(\theta) = \sum_{i=1}^n \ln g_\theta(x_i), \quad \theta \in \Theta$$

Extending the Method and the Invariance Property

Returning to the general setting, suppose now that h is a one-to-one function from the parameter space Θ onto a set Λ . We can view $\lambda = h(\theta)$ as a new parameter taking values in the space Λ , and it is easy to re-parameterize the probability density function with the new parameter. Thus, let $\hat{f}_\lambda(\mathbf{x}) = f_{h^{-1}(\lambda)}(\mathbf{x})$ for $\mathbf{x} \in S$ and $\lambda \in \Lambda$. The corresponding likelihood function for $\mathbf{x} \in S$ is

$$\hat{L}_{\mathbf{x}}(\lambda) = L_{\mathbf{x}}[h^{-1}(\lambda)], \quad \lambda \in \Lambda \quad (7.3.5)$$

Clearly if $u(\mathbf{x}) \in \Theta$ maximizes $L_{\mathbf{x}}$ for $\mathbf{x} \in S$. Then $h[u(\mathbf{x})] \in \Lambda$ maximizes $\hat{L}_{\mathbf{x}}$ for $\mathbf{x} \in S$. It follows that if U is a maximum likelihood estimator for θ , then $V = h(U)$ is a maximum likelihood estimator for $\lambda = h(\theta)$.

If the function h is not one-to-one, the maximum likelihood function for the new parameter $\lambda = h(\theta)$ is not well defined, because we cannot parameterize the probability density function in terms of λ . However, there is a natural generalization of the method.

Suppose that $h : \Theta \rightarrow \Lambda$, and let $\lambda = h(\theta)$ denote the new parameter. Define the *likelihood function* for λ at $\mathbf{x} \in S$ by

$$\hat{L}_{\mathbf{x}}(\lambda) = \max \{L_{\mathbf{x}}(\theta) : \theta \in h^{-1}(\lambda)\}; \quad \lambda \in \Lambda \quad (7.3.6)$$

If $v(\mathbf{x}) \in \Lambda$ maximizes $\hat{L}_{\mathbf{x}}$ for each $\mathbf{x} \in S$, then $V = v(\mathbf{X})$ is a *maximum likelihood estimator* of λ .

This definition extends the maximum likelihood method to cases where the probability density function is not completely parameterized by the parameter of interest. The following theorem is known as the *invariance property*: if we can solve the maximum likelihood problem for θ then we can solve the maximum likelihood problem for $\lambda = h(\theta)$.

In the setting of the previous theorem, if U is a maximum likelihood estimator of θ , then $V = h(U)$ is a maximum likelihood estimator of λ .

Proof

As before, if $u(\mathbf{x}) \in \Theta$ maximizes $L_{\mathbf{x}}$ for $\mathbf{x} \in S$. Then $h[u(\mathbf{x})] \in \Lambda$ maximizes $\hat{L}_{\mathbf{x}}$ for $\mathbf{x} \in S$.

Examples and Special Cases

In the following subsections, we will study maximum likelihood estimation for a number of special parametric families of distributions. Recall that if $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from a distribution with mean μ and variance σ^2 , then the method of moments estimators of μ and σ^2 are, respectively,

$$M = \frac{1}{n} \sum_{i=1}^n X_i \quad (7.3.7)$$

$$T^2 = \frac{1}{n} \sum_{i=1}^n (X_i - M)^2 \quad (7.3.8)$$

Of course, M is the sample mean, and T^2 is the biased version of the sample variance. These statistics will also sometimes occur as maximum likelihood estimators. Another statistic that will occur in some of the examples below is

$$M_2 = \frac{1}{n} \sum_{i=1}^n X_i^2 \quad (7.3.9)$$

the second-order sample mean. As always, be sure to try the derivations yourself before looking at the solutions.

The Bernoulli Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Bernoulli distribution with success parameter $p \in [0, 1]$. Recall that the Bernoulli probability density function is

$$g(x) = p^x(1-p)^{1-x}, \quad x \in \{0, 1\} \quad (7.3.10)$$

Thus, \mathbf{X} is a sequence of independent indicator variables with $\mathbb{P}(X_i = 1) = p$ for each i . In the usual language of reliability, X_i is the outcome of trial i , where 1 means success and 0 means failure. Let $Y = \sum_{i=1}^n X_i$ denote the number of successes, so that the proportion of successes (the sample mean) is $M = Y/n$. Recall that Y has the binomial distribution with parameters n and p .

The sample mean M is the maximum likelihood estimator of p on the parameter space $(0, 1)$.

Proof

Note that $\ln g(x) = x \ln p + (1-x) \ln(1-p)$ for $x \in \{0, 1\}$. Hence the log-likelihood function at $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ is

$$\ln L_{\mathbf{x}}(p) = \sum_{i=1}^n [x_i \ln p + (1-x_i) \ln(1-p)], \quad p \in (0, 1) \quad (7.3.11)$$

Differentiating with respect to p and simplifying gives

$$\frac{d}{dp} \ln L_{\mathbf{x}}(p) = \frac{y}{p} - \frac{n-y}{1-p} \quad (7.3.12)$$

where $y = \sum_{i=1}^n x_i$. Thus, there is a single critical point at $p = y/n = m$. The second derivative is

$$\frac{d^2}{dp^2} \ln L_{\mathbf{x}}(p) = -\frac{y}{p^2} - \frac{n-1}{(1-p)^2} < 0 \quad (7.3.13)$$

Hence the log-likelihood function is concave downward and so the maximum occurs at the unique critical point m .

Recall that M is also the method of moments estimator of p . It's always nice when two different estimation procedures yield the same result. Next let's look at the same problem, but with a much restricted parameter space.

Suppose now that p takes values in $\{\frac{1}{2}, 1\}$. Then the maximum likelihood estimator of p is the statistic

$$U = \begin{cases} 1, & Y = n \\ \frac{1}{2}, & Y < n \end{cases} \quad (7.3.14)$$

1. $\mathbb{E}(U) = \begin{cases} 1, & p = 1 \\ \frac{1}{2} + (\frac{1}{2})^{n+1}, & p = \frac{1}{2} \end{cases}$
2. U is positively biased, but is asymptotically unbiased.
3. $\text{mse}(U) = \begin{cases} 0 & p = 1 \\ (\frac{1}{2})^{n+2}, & p = \frac{1}{2} \end{cases}$
4. U is consistent.

Proof

Note that the likelihood function at $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ is $L_{\mathbf{x}}(p) = p^y(1-p)^{n-y}$ for $p \in \{\frac{1}{2}, 1\}$ where as usual, $y = \sum_{i=1}^n x_i$. Thus $L_{\mathbf{x}}(\frac{1}{2}) = (\frac{1}{2})^y$. On the other hand, $L_{\mathbf{x}}(1) = 0$ if $y < n$ while $L_{\mathbf{x}}(1) = 1$ if $y = n$. Thus, if $y = n$ the maximum occurs when $p = 1$ while if $y < n$ the maximum occurs when $p = \frac{1}{2}$.

1. If $p = 1$ then $\mathbb{P}(U = 1) = \mathbb{P}(Y = n) = 1$, so trivially $\mathbb{E}(U) = 1$. If $p = \frac{1}{2}$,

$$\mathbb{E}(U) = 1\mathbb{P}(Y = n) + \frac{1}{2}\mathbb{P}(Y < n) = 1\left(\frac{1}{2}\right)^n + \frac{1}{2}\left[1 - \left(\frac{1}{2}\right)^n\right] = \frac{1}{2} + \left(\frac{1}{2}\right)^{n+1} \quad (7.3.15)$$

2. Note that $\mathbb{E}(U) \geq p$ and $\mathbb{E}(U) \rightarrow p$ as $n \rightarrow \infty$ both in the case that $p = 1$ and $p = \frac{1}{2}$.
3. If $p = 1$ then $U = 1$ with probability 1, so trivially $\text{mse}(U) = 0$. If $p = \frac{1}{2}$,

$$\text{mse}(U) = \left(1 - \frac{1}{2}\right)^2 \mathbb{P}(Y = n) + \left(\frac{1}{2} - \frac{1}{2}\right)^2 \mathbb{P}(Y < n) = \left(\frac{1}{2}\right)^2 \left(\frac{1}{2}\right)^n = \left(\frac{1}{2}\right)^{n+2} \quad (7.3.16)$$

4. From (c), $\text{mse}(U) \rightarrow 0$ as $n \rightarrow \infty$.

Note that the Bernoulli distribution in the last exercise would model a coin that is either fair or two-headed. The last two exercises show that the maximum likelihood estimator of a parameter, like the solution to any maximization problem, depends critically on the domain.

U is uniformly better than M on the parameter space $\{\frac{1}{2}, 1\}$.

Proof

Recall that $\text{mse}(M) = \text{var}(M) = p(1-p)/n$. If $p = 1$ then $\text{mse}(M) = \text{mse}(U) = 0$ so that both estimators give the correct answer. If $p = \frac{1}{2}$, $\text{mse}(U) = (\frac{1}{2})^{n+2} < \frac{1}{4n} = \text{mse}(M)$.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Bernoulli distribution with unknown success parameter $p \in (0, 1)$. Find the maximum likelihood estimator of $p(1-p)$, which is the variance of the sampling distribution.

Answer

By the invariance principle, the estimator is $M(1-M)$ where M is the sample mean.

The Geometric Distribution

Recall that the *geometric distribution* on \mathbb{N}_+ with success parameter $p \in (0, 1)$ has probability density function

$$g(x) = p(1-p)^{x-1}, \quad x \in \mathbb{N}_+ \quad (7.3.17)$$

The geometric distribution governs the trial number of the first success in a sequence of Bernoulli trials.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the geometric distribution with unknown parameter $p \in (0, 1)$. The maximum likelihood estimator of p is $U = 1/M$.

Proof

Note that $\ln g(x) = \ln p + (x-1)\ln(1-p)$ for $x \in \mathbb{N}_+$. Hence the log-likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}_+^n$ is

$$\ln L_{\mathbf{x}}(p) = n \ln p + (y-n) \ln(1-p), \quad p \in (0, 1) \quad (7.3.18)$$

where $y = \sum_{i=1}^n x_i$. So

$$\frac{d}{dp} \ln L(p) = \frac{n}{p} - \frac{y-n}{1-p} \quad (7.3.19)$$

The derivative is 0 when $p = n/y = 1/m$. Finally, $\frac{d^2}{dp^2} \ln L_{\mathbf{x}}(p) = -n/p^2 - (y-n)/(1-p)^2 < 0$ so the maximum occurs at the critical point.

Recall that U is also the method of moments estimator of p . It's always reassuring when two different estimation procedures produce the same estimator.

The Negative Binomial Distribution

More generally, the *negative binomial distribution* on \mathbb{N} with shape parameter $k \in (0, \infty)$ and success parameter $p \in (0, 1)$ has probability density function

$$g(x) = \binom{x+k-1}{k-1} p^k (1-p)^x, \quad x \in \mathbb{N} \quad (7.3.20)$$

If k is a positive integer, then this distribution governs the number of failures before the k th success in a sequence of Bernoulli trials with success parameter p . However, the distribution makes sense for general $k \in (0, \infty)$. The negative binomial distribution is studied in more detail in the chapter on Bernoulli Trials.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the negative binomial distribution on \mathbb{N} with known shape parameter k and unknown success parameter $p \in (0, 1)$. The maximum likelihood estimator of p is

$$U = \frac{k}{k+M} \quad (7.3.21)$$

Proof

Note that $\ln g(x) = \ln \binom{x+k-1}{k-1} + k \ln p + x \ln(1-p)$ for $x \in \mathbb{N}$. Hence the log-likelihood function corresponding to $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}^n$ is

$$\ln L_{\mathbf{x}}(p) = nk \ln p + y \ln(1-p) + C, \quad p \in (0, 1) \quad (7.3.22)$$

where $y = \sum_{i=1}^n x_i$ and $C = \sum_{i=1}^n \ln \binom{x_i+k-1}{k-1}$. Hence

$$\frac{d}{dp} \ln L_{\mathbf{x}}(p) = \frac{nk}{p} - \frac{y}{1-p} \quad (7.3.23)$$

The derivative is 0 when $p = nk/(nk+y) = k/(k+m)$ where as usual, $m = y/n$. Finally, $\frac{d^2}{dp^2} \ln L_{\mathbf{x}}(p) = -nk/p^2 - y/(1-p)^2 < 0$, so the maximum occurs at the critical point.

Once again, this is the same as the method of moments estimator of p with k known.

The Poisson Distribution

Recall that the *Poisson distribution* with parameter $r > 0$ has probability density function

$$g(x) = e^{-r} \frac{r^x}{x!}, \quad x \in \mathbb{N} \quad (7.3.24)$$

The Poisson distribution is named for Simeon Poisson and is widely used to model the number of random “points” in a region of time or space. The parameter r is proportional to the size of the region. The Poisson distribution is studied in more detail in the chapter on the Poisson process.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Poisson distribution with unknown parameter $r \in (0, \infty)$. The maximum likelihood estimator of r is the sample mean M .

Proof

Note that $\ln g(x) = -r + x \ln r - \ln(x!)$ for $x \in \mathbb{N}$. Hence the log-likelihood function corresponding to $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}^n$ is

$$\ln L_{\mathbf{x}}(r) = -nr + y \ln r - C, \quad r \in (0, \infty) \quad (7.3.25)$$

where $y = \sum_{i=1}^n x_i$ and $C = \sum_{i=1}^n \ln(x_i!)$. Hence $\frac{d}{dr} \ln L_{\mathbf{x}}(r) = -n + y/r$. The derivative is 0 when $r = y/n = m$. Finally, $\frac{d^2}{dr^2} \ln L_{\mathbf{x}}(r) = -y/r^2 < 0$, so the maximum occurs at the critical point.

Recall that for the Poisson distribution, the parameter r is both the mean and the variance. Thus M is also the method of moments estimator of r . We showed in the introductory section that M has smaller mean square error than S^2 , although both are unbiased.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Poisson distribution with parameter $r \in (0, \infty)$, and let $p = \mathbb{P}(X = 0) = e^{-r}$. Find the maximum likelihood estimator of p in two ways:

1. Directly, by finding the likelihood function corresponding to the parameter p .
2. By using the result of the [last exercise](#) and the invariance property.

Answer

e^{-M} where M is the sample mean.

The Normal Distribution

Recall that the *normal distribution* with mean μ and variance σ^2 has probability density function

$$g(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (7.3.26)$$

The normal distribution is often used to model physical quantities subject to small, random errors, and is studied in more detail in the chapter on Special Distributions

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with unknown mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$. The maximum likelihood estimators of μ and σ^2 are M and T^2 , respectively.

Proof

Note that

$$\ln g(x) = -\frac{1}{2}\ln(2\pi) - \frac{1}{2}\ln(\sigma^2) - \frac{1}{2\sigma^2}(x-\mu)^2, \quad x \in \mathbb{R} \quad (7.3.27)$$

Hence the log-likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ is

$$\ln L_{\mathbf{x}}(\mu, \sigma^2) = -\frac{n}{2}\ln(2\pi) - \frac{n}{2}\ln(\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2, \quad (\mu, \sigma^2) \in \mathbb{R} \times (0, \infty) \quad (7.3.28)$$

Taking partial derivatives gives

$$\begin{aligned} \frac{\partial}{\partial \mu} \ln L_{\mathbf{x}}(\mu, \sigma^2) &= \frac{1}{\sigma^2} \sum_{i=1}^n (x_i - \mu) = \frac{1}{\sigma^2} \left(\sum_{i=1}^n x_i - n\mu \right) \\ \frac{\partial}{\partial \sigma^2} \ln L_{\mathbf{x}}(\mu, \sigma^2) &= -\frac{n}{2\sigma^2} + \frac{1}{2\sigma^4} \sum_{i=1}^n (x_i - \mu)^2 \end{aligned}$$

The partial derivatives are 0 when $\mu = \frac{1}{n} \sum_{i=1}^n x_i$ and $\sigma^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2$. Hence the unique critical point is (m, t^2) . Finally, with a bit more calculus, the second partial derivatives evaluated at the critical point are

$$\frac{\partial^2}{\partial \mu^2} \ln L_{\mathbf{x}}(m, t^2) = -n/t^2, \quad \frac{\partial^2}{\partial \mu \partial \sigma^2} \ln L_{\mathbf{x}}(m, t^2) = 0, \quad \frac{\partial^2}{\partial (\sigma^2)^2} \ln L_{\mathbf{x}}(m, t^2) = -n/t^4 \quad (7.3.29)$$

Hence the second derivative matrix at the critical point is negative definite and so the maximum occurs at the critical point.

Of course, M and T^2 are also the method of moments estimators of μ and σ^2 , respectively.

Run the Normal estimation experiment 1000 times for several values of the sample size n , the mean μ , and the variance σ^2 . For the parameter σ^2 , compare the maximum likelihood estimator T^2 with the standard sample variance S^2 . Which estimator seems to work better in terms of mean square error?

Suppose again that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with unknown mean $\mu \in \mathbb{R}$ and unknown variance $\sigma^2 \in (0, \infty)$. Find the maximum likelihood estimator of $\mu^2 + \sigma^2$, which is the second moment about 0 for the sampling distribution.

Answer

By the invariance principle, the estimator is $M^2 + T^2$ where M is the sample mean and T^2 is the (biased version of the) sample variance.

The Gamma Distribution

Recall that the *gamma distribution* with shape parameter $k > 0$ and scale parameter $b > 0$ has probability density function

$$g(x) = \frac{1}{\Gamma(k) b^k} x^{k-1} e^{-x/b}, \quad 0 < x < \infty \quad (7.3.30)$$

The gamma distribution is often used to model random times and certain other types of positive random variables, and is studied in more detail in the chapter on Special Distributions

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the gamma distribution with known shape parameter k and unknown scale parameter $b \in (0, \infty)$. The maximum likelihood estimator of b is $V_k = \frac{1}{k} M$.

Proof

Note that for $x \in (0, \infty)$,

$$\ln g(x) = -\ln \Gamma(k) - k \ln b + (k-1) \ln x - \frac{x}{b} \quad (7.3.31)$$

and hence the log-likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n) \in (0, \infty)^n$ is

$$\ln L_{\mathbf{x}}(b) = -nk \ln b - \frac{y}{b} + C, \quad b \in (0, \infty) \quad (7.3.32)$$

where $y = \sum_{i=1}^n x_i$ and $C = -n \ln \Gamma(k) + (k-1) \sum_{i=1}^n \ln x_i$. It follows that

$$\frac{d}{db} \ln L_{\mathbf{x}}(b) = -\frac{nk}{b} + \frac{y}{b^2} \quad (7.3.33)$$

The derivative is 0 when $b = y/nk = 1/km$. Finally, $\frac{d^2}{db^2} \ln L_{\mathbf{x}}(b) = nk/b^2 - 2y/b^3$. At the critical point $b = y/nk$, the second derivative is $-(nk)^3/y^2 < 0$ so the maximum occurs at the critical point.

Recall that V_k is also the method of moments estimator of b when k is known. But when k is unknown, the method of moments estimator of b is $V = \frac{T^2}{M}$.

Run the gamma estimation experiment 1000 times for several values of the sample size n , shape parameter k , and scale parameter b . In each case, compare the method of moments estimator V of b when k is unknown with the method of moments and maximum likelihood estimator V_k of b when k is known. Which estimator seems to work better in terms of mean square error?

The Beta Distribution

Recall that the *beta distribution* with left parameter $a \in (0, \infty)$ and right parameter $b = 1$ has probability density function

$$g(x) = ax^{a-1}, \quad x \in (0, 1) \quad (7.3.34)$$

The beta distribution is often used to model random proportions and other random variables that take values in bounded intervals. It is studied in more detail in the chapter on Special Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the beta distribution with unknown left parameter $a \in (0, \infty)$ and right parameter $b = 1$. The maximum likelihood estimator of a is

$$W = -\frac{n}{\sum_{i=1}^n \ln X_i} = -\frac{n}{\ln(X_1 X_2 \cdots X_n)} \quad (7.3.35)$$

Proof

Note that $\ln g(x) = \ln a + (a-1) \ln x$ for $x \in (0, \infty)$. Hence the log-likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n) \in (0, \infty)^n$ is

$$\ln L_{\mathbf{x}}(a) = n \ln a + (a-1) \sum_{i=1}^n \ln x_i, \quad a \in (0, \infty) \quad (7.3.36)$$

Therefore $\frac{d}{da} \ln L_{\mathbf{x}}(a) = n/a + \sum_{i=1}^n \ln x_i$. The derivative is 0 when $a = -n / \sum_{i=1}^n \ln x_i$. Finally, $\frac{d^2}{da^2} \ln L_{\mathbf{x}}(a) = -n/a^2 < 0$, so the maximum occurs at the critical point.

Recall that when $b = 1$, the method of moments estimator of a is $U_1 = M/(1 - M)$, but when $b \in (0, \infty)$ is also unknown, the method of moments estimator of a is $U = M(M - M_2)/(M_2 - M^2)$. When $b = 1$, which estimator is better, the method of moments estimator or the maximum likelihood estimator?

In the beta estimation experiment, set $b = 1$. Run the experiment 1000 times for several values of the sample size n and the parameter a . In each case, compare the estimators U , U_1 and W . Which estimator seems to work better in terms of mean square error?

Finally, note that $1/W$ is the sample mean for a random sample of size n from the distribution of $-\ln X$. This distribution is the exponential distribution with rate a .

The Pareto Distribution

Recall that the *Pareto distribution* with shape parameter $a > 0$ and scale parameter $b > 0$ has probability density function

$$g(x) = \frac{ab^a}{x^{a+1}}, \quad b \leq x < \infty \quad (7.3.37)$$

The Pareto distribution, named for Vilfredo Pareto, is a heavy-tailed distribution often used to model income and certain other types of random variables. It is studied in more detail in the chapter on Special Distribution.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Pareto distribution with unknown shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. The maximum likelihood estimator of b is $X_{(1)} = \min\{X_1, X_2, \dots, X_n\}$, the first order statistic. The maximum likelihood estimator of a is

$$U = \frac{n}{\sum_{i=1}^n \ln X_i - n \ln X_{(1)}} = \frac{n}{\sum_{i=1}^n (\ln X_i - \ln X_{(1)})} \quad (7.3.38)$$

Proof

Note that $\ln g(x) = \ln a + a \ln b - (a+1) \ln x$ for $x \in [b, \infty)$. Hence the log-likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is

$$\ln L_{\mathbf{x}}(a, b) = n \ln a + na \ln b - (a+1) \sum_{i=1}^n \ln x_i; \quad 0 < a < \infty, \quad 0 < b \leq x_i \text{ for each } i \in \{1, 2, \dots, n\} \quad (7.3.39)$$

Equivalently, the domain is $0 < a < \infty$ and $0 < b \leq x_{(1)}$. Note that $\ln L_{\mathbf{x}}(a, b)$ is increasing in b for each a , and hence is maximized when $b = x_{(1)}$ for each a . Next,

$$\frac{d}{da} \ln L_{\mathbf{x}}(a, x_{(1)}) = \frac{n}{a} + n \ln x_{(1)} - \sum_{i=1}^n \ln x_i \quad (7.3.40)$$

The derivative is 0 when $a = n / (\sum_{i=1}^n \ln x_i - n \ln x_{(1)})$. Finally, $\frac{d^2}{da^2} \ln L_{\mathbf{x}}(a, x_{(1)}) = -n/a^2 < 0$, so the maximum occurs at the critical point.

Recall that if $a > 2$, the method of moments estimators of a and b are

$$1 + \sqrt{\frac{M_2}{M_2 - M^2}}, \quad \frac{M_2}{M} \left(1 - \sqrt{\frac{M_2 - M^2}{M_2}} \right) \quad (7.3.41)$$

Open the the Pareto estimation experiment. Run the experiment 1000 times for several values of the sample size n and the parameters a and b . Compare the method of moments and maximum likelihood estimators. Which estimators seem to work better in terms of bias and mean square error?

Often the scale parameter in the Pareto distribution is known.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Pareto distribution with unknown shape parameter $a \in (0, \infty)$ and known scale parameter $b \in (0, \infty)$. The maximum likelihood estimator of a is

$$U = \frac{n}{\sum_{i=1}^n \ln X_i - n \ln b} = \frac{n}{\sum_{i=1}^n (\ln X_i - \ln b)} \quad (7.3.42)$$

Proof

Modifying the previous proof, the log-likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is

$$\ln L_{\mathbf{x}}(a) = n \ln a + na \ln b - (a+1) \sum_{i=1}^n \ln x_i, \quad 0 < a < \infty \quad (7.3.43)$$

The derivative is

$$\frac{d}{da} \ln L_{\mathbf{x}}(a) = \frac{n}{a} + n \ln b - \sum_{i=1}^n \ln x_i \quad (7.3.44)$$

The derivative is 0 when $a = n / (\sum_{i=1}^n \ln x_i - n \ln b)$. Finally, $\frac{d^2}{da^2} \ln L_{\mathbf{x}}(a) = -n/a^2 < 0$, so the maximum occurs at the critical point.

Uniform Distributions

In this section we will study estimation problems related to the uniform distribution that are a good source of insight and counterexamples. In a sense, our first estimation problem is the continuous analogue of an estimation problem studied in the section on Order Statistics in the chapter Finite Sampling Models. Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the uniform distribution on the interval $[0, h]$, where $h \in (0, \infty)$ is an unknown parameter. Thus, the sampling distribution has probability density function

$$g(x) = \frac{1}{h}, \quad x \in [0, h] \quad (7.3.45)$$

First let's review results from the last section.

The method of moments estimator of h is $U = 2M$. The estimator U satisfies the following properties:

1. U is unbiased.
2. $\text{var}(U) = \frac{h^2}{3n}$ so U is consistent.

Now let's find the maximum likelihood estimator

The maximum likelihood estimator of h is $X_{(n)} = \max\{X_1, X_2, \dots, X_n\}$, the n th order statistic. The estimator $X_{(n)}$ satisfies the following properties:

1. $\mathbb{E}(X_{(n)}) = \frac{n}{n+1}h$
2. $\text{bias}(X_{(n)}) = -\frac{h}{n+1}$ so that $X_{(n)}$ is negatively biased but asymptotically unbiased.
3. $\text{var}(X_{(n)}) = \frac{n}{(n+2)(n+1)^2}h^2$
4. $\text{mse}(X_{(n)}) = \frac{2}{(n+1)(n+2)}h^2$ so that $X_{(n)}$ is consistent.

Proof

The likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is $L_{\mathbf{x}}(h) = 1/h^n$ for $h \geq x_i$ for each $i \in \{1, 2, \dots, n\}$. The domain is equivalent to $h \geq x_{(n)}$. The function $h \mapsto 1/h^n$ is decreasing, and so the maximum occurs at the smallest value, namely $x_{(n)}$. Parts (a) and (c) are restatements of results from the section on order statistics. Parts (b) and (d) follow from (a) and (c).

Since the expected value of $X_{(n)}$ is a known multiple of the parameter h , we can easily construct an unbiased estimator.

Let $V = \frac{n+1}{n}X_{(n)}$. The estimator V satisfies the following properties:

1. V is unbiased.
2. $\text{var}(V) = \frac{h^2}{n(n+2)}$ so that V is consistent.
3. The asymptotic relative efficiency of V to U is infinite.

Proof

Parts (a) and (b) follow from the [previous](#) result and basic properties of the expected value and variance. For part (c),

$$\frac{\text{var}(U)}{\text{var}(V)} = \frac{h^2/3n}{h^2/n(n+2)} = \frac{n+2}{3} \rightarrow \infty \text{ as } n \rightarrow \infty \quad (7.3.46)$$

The last part shows that the unbiased version V of the maximum likelihood estimator is a much better estimator than the method of moments estimator U . In fact, an estimator such as V , whose mean square error decreases on the order of $\frac{1}{n^2}$, is called *super efficient*. Now, having found a really good estimator, let's see if we can find a really bad one. A natural candidate is an estimator based on $X_{(1)} = \min\{X_1, X_2, \dots, X_n\}$, the *first order statistic*. The next result will make the computations very easy.

The sample $\mathbf{X} = (X_1, X_2, \dots, X_n)$ satisfies the following properties:

1. $h - X_i$ is uniformly distributed on $[0, h]$ for each i .
2. $(h - X_1, h - X_2, \dots, h - X_n)$ is also a random sample from the uniform distribution on $[0, h]$.
3. $X_{(1)}$ has the same distribution as $h - X_{(n)}$.

Proof

1. This is a simple consequence of the fact that uniform distributions are preserved under linear transformations on the random variable.
2. This follows from (a) and that the fact that if \mathbf{X} is a sequence of independent variables, then so is $(h - X_1, h - X_2, \dots, h - X_n)$.
3. From part (b), $X_{(1)} = \min\{X_1, X_2, \dots, X_n\}$ has the same distribution as $\min\{h - X_1, h - X_2, \dots, h - X_n\} = h - \max\{X_1, X_2, \dots, X_n\} = h - X_{(n)}$.

Now we can construct our really bad estimator.

Let $W = (n+1)X_{(1)}$. Then

1. W is an unbiased estimator of h .
2. $\text{var}(W) = \frac{n}{n+2}h^2$, so W is not even consistent.

Proof

These results follow from the ones above:

1. $\mathbb{E}(X_{(1)}) = h - \mathbb{E}(X_{(n)}) = h - \frac{n}{n+1}h = \frac{1}{n+1}h$ and hence $\mathbb{E}(W) = h$.
2. $\text{var}(W) = (n+1)^2 \text{var}(X_{(1)}) = (n+1)^2 \text{var}(h - X_{(n)}) = (n+1)^2 \frac{n}{(n+1)^2(n+2)}h^2 = \frac{n}{n+2}h^2$.

Run the uniform estimation experiment 1000 times for several values of the sample size n and the parameter a . In each case, compare the empirical bias and mean square error of the estimators with their theoretical values. Rank the estimators in terms of empirical mean square error.

Our next series of exercises will show that the maximum likelihood estimator is not necessarily unique. Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the uniform distribution on the interval $[a, a+1]$, where $a \in \mathbb{R}$ is an unknown parameter. Thus, the sampling distribution has probability density function

$$g(x) = 1, \quad a \leq x \leq a+1 \quad (7.3.47)$$

As usual, let's first review the method of moments estimator.

The method of moments estimator of a is $U = M - \frac{1}{2}$. The estimator U satisfies the following properties:

1. U is unbiased.

2. $\text{var}(U) = \frac{1}{12n}$ so U is consistent.

However, as promised, there is not a unique maximum likelihood estimator.

Any statistic $V \in [X_{(n)} - 1, X_{(1)}]$ is a maximum likelihood estimator of a .

Proof

The likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is $L_{\mathbf{x}}(a) = 1$ for $a \leq x_i \leq a+1$ and $i \in \{1, 2, \dots, n\}$. The domain is equivalent to $a \leq x_{(1)}$ and $a \geq x_{(n)} - 1$. Since the likelihood function is constant on this domain, the result follows.

For completeness, let's consider the full estimation problem. Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the uniform distribution on $[a, a+h]$ where $a \in \mathbb{R}$ and $h \in (0, \infty)$ are both unknown. Here's the result from the last section:

Let U and V denote the method of moments estimators of a and h , respectively. Then

$$U = 2M - \sqrt{3}T, \quad V = 2\sqrt{3}T \quad (7.3.48)$$

where $M = \frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean, and $T = \frac{1}{n} \sum_{i=1}^n (X_i - M)^2$ is the biased version of the sample variance.

It should come as no surprise at this point that the maximum likelihood estimators are functions of the largest and smallest order statistics.

The maximum likelihood estimators of a and h are $U = X_{(1)}$ and $V = X_{(n)} - X_{(1)}$, respectively.

1. $E(U) = a + \frac{h}{n+1}$ so U is positively biased and asymptotically unbiased.
2. $E(V) = h \frac{n-1}{n+1}$ so V is negatively biased and asymptotically unbiased.
3. $\text{var}(U) = h^2 \frac{n}{(n+1)^2(n+2)}$ so U is consistent.
4. $\text{var}(V) = h^2 \frac{2(n-1)}{(n+1)^2(n+2)}$ so V is consistent.

Proof

The likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is $L_{\mathbf{x}}(a, h) = \frac{1}{h^n}$ for $a \leq x_i \leq a+h$ and $i \in \{1, 2, \dots, n\}$. The domain is equivalent to $a \leq x_{(1)}$ and $a+h \geq x_{(n)}$. Since the likelihood function depends only on h in this domain and is decreasing, the maximum occurs when $a = x_{(1)}$ and $h = x_{(n)} - x_{(1)}$. Parts (a)–(d) follow from standard results for the order statistics from the uniform distribution.

The Hypergeometric Model

In all of our previous examples, the sequence of observed random variables $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from a distribution. However, maximum likelihood is a very general method that does not require the observation variables to be independent or identically distributed. In the *hypergeometric model*, we have a population of N objects with r of the objects *type 1* and the remaining $N - r$ objects *type 0*. The *population size* N , is a positive integer. The *type 1 size* r , is a nonnegative integer with $r \leq N$. These are the basic parameters, and typically one or both is unknown. Here are some typical examples:

1. The objects are devices, classified as *good* or *defective*.
2. The objects are persons, classified as *female* or *male*.
3. The objects are voters, classified as *for* or *against* a particular candidate.
4. The objects are wildlife or a particular type, either *tagged* or *untagged*.

We sample n objects from the population at random, without replacement. Let X_i be the type of the i th object selected, so that our sequence of observed variables is $\mathbf{X} = (X_1, X_2, \dots, X_n)$. The variables are identically distributed indicator variables, with $P(X_i = 1) = r/N$ for each $i \in \{1, 2, \dots, n\}$, but are dependent since the sampling is without replacement. The number of type 1 objects in the sample is $Y = \sum_{i=1}^n X_i$. This statistic has the *hypergeometric distribution* with parameter N , r , and n , and has probability density function given by

$$P(Y = y) = \frac{\binom{r}{y} \binom{N-r}{n-y}}{\binom{N}{n}} = \binom{n}{y} \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad y \in \{\max\{0, N-n+r\}, \dots, \min\{n, r\}\} \quad (7.3.49)$$

Recall the *falling power* notation: $x^{(k)} = x(x-1) \cdots (x-k+1)$ for $x \in \mathbb{R}$ and $k \in \mathbb{N}$. The hypergeometric model is studied in more detail in the chapter on Finite Sampling Models.

As above, let $\mathbf{X} = (X_1, X_2, \dots, X_n)$ be the observed variables in the hypergeometric model with parameters N and r . Then

1. The maximum likelihood estimator of r with N known is $U = \lfloor NM \rfloor = \lfloor NY/n \rfloor$.
2. The maximum likelihood estimator of N with r known is $V = \lfloor r/M \rfloor = \lfloor rn/Y \rfloor$ if $Y > 0$.

Proof

By a simple application of the multiplication rule, the PDF f of \mathbf{X} is

$$f(\mathbf{x}) = \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \quad (7.3.50)$$

where $y = \sum_{i=1}^n x_i$.

1. With N known, the likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ is

$$L_{\mathbf{x}}(r) = \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad r \in \{y, \dots, \min\{n, y + N - n\}\} \quad (7.3.51)$$

After some algebra, $L_{\mathbf{x}}(r-1) < L_{\mathbf{x}}(r)$ if and only if $(r-y)(N-r+1) < r(N-r-n+y+1)$ if and only if $r < Ny/n$. So the maximum of $L_{\mathbf{x}}(r)$ occurs when $r = \lfloor Ny/n \rfloor$.

2. Similarly, with r known, the likelihood function corresponding to the data $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ is

$$L_{\mathbf{x}}(N) = \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad N \in \{\max\{r, n\}, \dots\} \quad (7.3.52)$$

After some algebra, $L_{\mathbf{x}}(N-1) < L_{\mathbf{x}}(N)$ if and only if $(N-r-n+y)/(N-n) < (N-r)/N$ if and only if $N < rn/y$ (assuming $y > 0$). So the maximum of $L_{\mathbf{x}}(r)$ occurs when $N = \lfloor rn/y \rfloor$.

In the reliability example (1), we might typically know N and would be interested in estimating r . In the wildlife example (4), we would typically know r and would be interested in estimating N . This example is known as the *capture-recapture* model.

Clearly there is a close relationship between the hypergeometric model and the [Bernoulli trials model](#) above. In fact, if the sampling is *with* replacement, the Bernoulli trials model with $p = r/N$ would apply rather than the hypergeometric model. In addition, if the population size N is large compared to the sample size n , the hypergeometric model is well approximated by the Bernoulli trials model, again with $p = r/N$.

This page titled [7.3: Maximum Likelihood](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.4: Bayesian Estimation

Basic Theory

The General Method

Suppose again that we have an observable random variable \mathbf{X} for an experiment, that takes values in a set S . Suppose also that distribution of \mathbf{X} depends on a parameter θ taking values in a parameter space T . Of course, our data variable \mathbf{X} is almost always vector-valued, so that typically $S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$. Depending on the nature of the sample space S , the distribution of \mathbf{X} may be discrete or continuous. The parameter θ may also be vector-valued, so that typically $T \subseteq \mathbb{R}^k$ for some $k \in \mathbb{N}_+$.

In *Bayesian analysis*, named for the famous Thomas Bayes, we model the deterministic, but unknown parameter θ with a random variable Θ that has a specified distribution on the parameter space T . Depending on the nature of the parameter space, this distribution may also be either discrete or continuous. It is called the *prior distribution* of Θ and is intended to reflect our knowledge of the parameter θ , *before* we gather data. After observing $\mathbf{X} = \mathbf{x} \in S$, we then use Bayes' theorem, to compute the conditional distribution of Θ given $\mathbf{X} = \mathbf{x}$. This distribution is called the *posterior distribution* of Θ , and is an updated distribution, given the information in the data. Here is the mathematical description, stated in terms of probability density functions.

Suppose that the *prior distribution* of Θ on T has probability density function h , and that given $\Theta = \theta \in T$, the conditional probability density function of \mathbf{X} on S is $f(\cdot | \theta)$. Then the probability density function of the *posterior distribution* of Θ given $\mathbf{X} = \mathbf{x} \in S$ is

$$h(\theta | \mathbf{x}) = \frac{h(\theta)f(\mathbf{x} | \theta)}{f(\mathbf{x})}, \quad \theta \in T \quad (7.4.1)$$

where the function in the denominator is defined as follows, in the discrete and continuous cases, respectively:

$$f(\mathbf{x}) = \sum_{\theta \in T} h(\theta)f(\mathbf{x} | \theta), \quad \mathbf{x} \in S$$

$$f(\mathbf{x}) = \int_T h(\theta)f(\mathbf{x} | \theta) d\theta, \quad \mathbf{x} \in S$$

Proof

This is just Bayes' theorem with new terminology. Recall that the joint probability density function of (\mathbf{X}, Θ) is the mapping on $S \times T$ given by

$$(\mathbf{x}, \theta) \mapsto h(\theta)f(\mathbf{x} | \theta) \quad (7.4.2)$$

Then the function in the denominator is the marginal probability density function of \mathbf{X} . So by definition, $h(\theta | \mathbf{x}) = h(\theta)f(\mathbf{x} | \theta)/f(\mathbf{x})$ for $\theta \in T$ is the conditional probability density function of Θ given $\mathbf{X} = \mathbf{x}$.

For $\mathbf{x} \in S$, note that $f(\mathbf{x})$ is simply the *normalizing constant* for the function $\theta \mapsto h(\theta)f(\mathbf{x} | \theta)$. It may not be necessary to explicitly compute $f(\mathbf{x})$, if one can recognize the functional form of $\theta \mapsto h(\theta)f(\mathbf{x} | \theta)$ as that of a known distribution. This will indeed be the case in several of the examples explored below.

If the parameter space T has finite measure c (counting measure in the discrete case or Lebesgue measure in the continuous case), then one possible prior distribution is the uniform distribution on T , with probability density function $h(\theta) = 1/c$ for $\theta \in T$. This distribution reflects no prior knowledge about the parameter, and so is called the *non-informative* prior distribution.

Random Samples

Of course, an important and essential special case occurs when $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of a basic variable X . Specifically, suppose that X takes values in a set R and has probability density function $g(\cdot | \theta)$ for a given $\theta \in T$. In this case, $S = R^n$ and the probability density function $f(\cdot | \theta)$ of \mathbf{X} given θ is

$$f(x_1, x_2, \dots, x_n | \theta) = g(x_1 | \theta)g(x_2 | \theta) \cdots g(x_n | \theta), \quad (x_1, x_2, \dots, x_n) \in S \quad (7.4.3)$$

Real Parameters

Suppose that θ is a real-valued parameter, so that $T \subseteq \mathbb{R}$. Here is our main definition.

The conditional expected value $\mathbb{E}(\Theta | \mathbf{X})$ is the *Bayesian estimator* of θ .

1. If Θ has a discrete distribution on T then

$$\mathbb{E}(\Theta | \mathbf{X} = \mathbf{x}) = \sum_{\theta \in T} \theta h(\theta | \mathbf{x}), \quad \mathbf{x} \in S \quad (7.4.4)$$

2. If Θ has a continuous distribution on T then

$$\mathbb{E}(\Theta | \mathbf{X} = \mathbf{x}) = \int_T \theta h(\theta | \mathbf{x}) d\theta, \quad \mathbf{x} \in S \quad (7.4.5)$$

Recall that $\mathbb{E}(\Theta | \mathbf{X})$ is a function of \mathbf{X} and, among all functions of \mathbf{X} , is closest to Θ in the mean square sense. Of course, once we collect the data and observe $\mathbf{X} = \mathbf{x}$, the *Bayesian estimate* of θ is $\mathbb{E}(\Theta | \mathbf{X} = \mathbf{x})$. As always, the term *estimator* refers to a random variable, before the data are collected, and the term *estimate* refers to an observed value of the random variable after the data are collected. The definitions of bias and mean square error are as before, but now conditioned on $\Theta = \theta \in T$.

Suppose that U is the Bayes estimator of θ .

1. The *bias* of U is $\text{bias}(U | \theta) = \mathbb{E}(U - \theta | \Theta = \theta)$ for $\theta \in T$.
2. The *mean square error* of U is $\text{mse}(U | \theta) = \mathbb{E}[(U - \theta)^2 | \Theta = \theta]$ for $\theta \in T$.

As before, $\text{bias}(U | \theta) = \mathbb{E}(U | \theta) - \theta$ and $\text{mse}(U | \theta) = \text{var}(U | \theta) + \text{bias}^2(U | \theta)$. Suppose now that we observe the random variables (X_1, X_2, X_3, \dots) sequentially, and we compute the Bayes estimator U_n of θ based on (X_1, X_2, \dots, X_n) for each $n \in \mathbb{N}_+$. Again, the most common case is when we are sampling from a distribution, so that the sequence is independent and identically distributed (given θ). We have the natural asymptotic properties that we have seen before.

Let $\mathbf{U} = (U_n : n \in \mathbb{N}_+)$ be the sequence of Bayes estimators of θ as above.

1. \mathbf{U} is *asymptotically unbiased* if $\text{bias}(U_n | \theta) \rightarrow 0$ as $n \rightarrow \infty$ for each $\theta \in T$.
2. \mathbf{U} is *mean-square consistent* if $\text{mse}(U_n | \theta) \rightarrow 0$ as $n \rightarrow \infty$ for each $\theta \in T$.

Often we cannot construct unbiased Bayesian estimators, but we do hope that our estimators are at least asymptotically unbiased and consistent. It turns out that the sequence of Bayesian estimators \mathbf{U} is a martingale. The theory of martingales provides some powerful tools for studying these estimators.

From the Bayesian perspective, the posterior distribution of Θ given the data $\mathbf{X} = \mathbf{x}$ is of primary importance. Point estimates of θ derived from this distribution are of secondary importance. In particular, the mean square error function $u \mapsto \mathbb{E}[(\Theta - u)^2 | \mathbf{X} = \mathbf{x}]$, minimized as we have noted at $\mathbb{E}(\Theta | \mathbf{X} = \mathbf{x})$, is not the only *loss function* that can be used. (Although it's the only one that we consider.) Another possible loss function, among many, is the mean absolute error function $u \mapsto \mathbb{E}[|\Theta - u| | \mathbf{X} = \mathbf{x}]$, which we know is minimized at the median(s) of the posterior distribution.

Conjugate Families

Often, the prior distribution of Θ is itself a member of a parametric family, with the parameters specified to reflect our prior knowledge of θ . In many important special cases, the parametric family can be chosen so that the posterior distribution of Θ given $\mathbf{X} = \mathbf{x}$ belongs to the same family for each $\mathbf{x} \in S$. In such a case, the family of distributions of Θ is said to be *conjugate* to the family of distributions of \mathbf{X} . Conjugate families are nice from a computational point of view, since we can often compute the posterior distribution through a simple formula involving the parameters of the family, without having to use Bayes' theorem directly. Similarly, in the case that the parameter is real valued, we can often compute the Bayesian estimator through a simple formula involving the parameters of the conjugate family.

Special Distributions

The Bernoulli Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is sequence of independent variables, each having the Bernoulli distribution with unknown success parameter $p \in (0, 1)$. In short, \mathbf{X} is a sequence of Bernoulli trials, given p . In the usual language of reliability, $X_i = 1$ means success on trial i and $X_i = 0$ means failure on trial i . Recall that given p , the Bernoulli distribution has probability density function

$$g(x | p) = p^x (1 - p)^{1-x}, \quad x \in \{0, 1\} \quad (7.4.6)$$

Note that the number of successes in the first n trials is $Y_n = \sum_{i=1}^n X_i$. Given p , random variable Y_n has the binomial distribution with parameters n and p .

Suppose now that we model p with a random variable P that has a prior beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$, where a and b are chosen to reflect our initial information about p . So P has probability density function

$$h(p) = \frac{1}{B(a, b)} p^{a-1} (1 - p)^{b-1}, \quad p \in (0, 1) \quad (7.4.7)$$

and has mean $a/(a+b)$. For example, if we know nothing about p , we might let $a = b = 1$, so that the prior distribution is uniform on the parameter space $(0, 1)$ (the non-informative prior). On the other hand, if we believe that p is about $\frac{2}{3}$, we might let $a = 4$ and $b = 2$, so that the prior distribution is unimodal, with mean $\frac{2}{3}$. As a random process, the sequence \mathbf{X} with p randomized by P , is known as the beta-Bernoulli process, and is very interesting on its own, outside of the context of Bayesian estimation.

For $n \in \mathbb{N}_+$, the posterior distribution of P given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is beta with left parameter $a + Y_n$ and right parameter $b + (n - Y_n)$.

Proof

Fix $n \in \mathbb{N}_+$. Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, and let $y = \sum_{i=1}^n x_i$. Then

$$f(\mathbf{x} | p) = g(x_1 | p) g(x_2 | p) \cdots g(x_n | p) = p^y (1 - p)^{n-y} \quad (7.4.8)$$

Hence

$$h(p) f(\mathbf{x} | p) = \frac{1}{B(a, b)} p^{a-1} (1 - p)^{b-1} p^y (1 - p)^{n-y} = \frac{1}{B(a, b)} p^{a+y-1} (1 - p)^{b+n-y-1}, \quad p \in (0, 1) \quad (7.4.9)$$

As a function of p this expression is proportional to the beta PDF with parameters $a + y$, $b + n - y$. Note that it's not necessary to compute the normalizing factor $f(\mathbf{x})$.

Thus, the beta distribution is conjugate to the Bernoulli distribution. Note also that the posterior distribution depends on the data vector \mathbf{X}_n only through the number of successes Y_n . This is true because Y_n is a sufficient statistic for p . In particular, note that the left beta parameter is increased by the number of successes Y_n and the right beta parameter is increased by the number of failures $n - Y_n$.

The Bayesian estimator of p given \mathbf{X}_n is

$$U_n = \frac{a + Y_n}{a + b + n} \quad (7.4.10)$$

Proof

Recall that the mean of the beta distribution is the left parameter divided by the sum of the parameters, so this result follows from the previous result.

In the beta coin experiment, set $n = 20$ and $p = 0.3$, and set $a = 4$ and $b = 2$. Run the simulation 100 times and note the estimate of p and the shape and location of the posterior probability density function of p on each run.

Next let's compute the bias and mean-square error functions.

For $n \in \mathbb{N}_+$,

$$\text{bias}(U_n | p) = \frac{a(1-p) - bp}{a+b+n}, \quad p \in (0, 1) \quad (7.4.11)$$

The sequence $\mathbf{U} = (U_n : n \in \mathbb{N}_+)$ is asymptotically unbiased.

Proof

Given p , Y_n has the binomial distribution with parameters n and p so $E(Y_n | p) = np$. Hence

$$\text{bias}(U_n | p) = E(U_n | p) - p = \frac{a + np}{a + b + n} - p \quad (7.4.12)$$

Simplifying gives the formula above. Clearly $\text{bias}(U_n | p) \rightarrow 0$ as $n \rightarrow \infty$.

Note also that we cannot choose a and b to make U_n unbiased, since such a choice would involve the true value of p , which we do not know.

In the beta coin experiment, vary the parameters and note the change in the bias. Now set $n = 20$ and $p = 0.8$, and set $a = 2$ and $b = 6$. Run the simulation 1000 times. Note the estimate of p and the shape and location of the posterior probability density function of p on each update. Compare the empirical bias to the true bias.

For $n \in \mathbb{N}_+$,

$$\text{mse}(U_n | p) = \frac{p[n - 2a(a+b)] + p^2[(a+b)^2 - n] + a^2}{(a+b+n)^2}, \quad p \in (0, 1) \quad (7.4.13)$$

The sequence $(U_n : n \in \mathbb{N}_+)$ is mean-square consistent.

Proof

Once again, given p , Y_n has the binomial distribution with parameters n and p so

$$\text{var}(U_n | p) = \frac{np(1-p)}{(a+b+n)^2} \quad (7.4.14)$$

Hence

$$\text{mse}(U_n | p) = \frac{np(1-p)}{(a+b+n)^2} + \left[\frac{a(1-p) - bp}{a+b+n} \right]^2 \quad (7.4.15)$$

Simplifying gives the result. Clearly $\text{mse}(U_n | p) \rightarrow 0$ as $n \rightarrow \infty$.

In the beta coin experiment, vary the parameters and note the change in the mean square error. Now set $n = 10$ and $p = 0.7$, and set $a = b = 1$. Run the simulation 1000 times. Note the estimate of p and the shape and location of the posterior probability density function of p on each update. Compare the empirical mean square error to the true mean square error.

Interestingly, we can choose a and b so that U has mean square error that is independent of the unknown parameter p :

Let $n \in \mathbb{N}_+$ and let $a = b = \sqrt{n}/2$. Then

$$\text{mse}(U_n | p) = \frac{n}{4(n + \sqrt{n})^2}, \quad p \in (0, 1) \quad (7.4.16)$$

In the beta coin experiment, set $n = 36$ and $a = b = 3$. Vary p and note that the mean square error does not change. Now set $p = 0.8$ and run the simulation 1000 times. Note the estimate of p and the shape and location of the posterior probability density function on each update. Compare the empirical bias and mean square error to the true values.

Recall that the method of moments estimator and the maximum likelihood estimator of p (on the interval $(0, 1)$) is the sample mean (the proportion of successes):

$$M_n = \frac{Y}{n} = \frac{1}{n} \sum_{i=1}^n X_i \quad (7.4.17)$$

This estimator has mean square error $\text{mse}(M_n | p) = \frac{1}{n} p(1-p)$. To see the connection between the estimators, note from (6) that

$$U_n = \frac{a+b}{a+b+n} \frac{a}{a+b} + \frac{n}{a+b+n} M_n \quad (7.4.18)$$

So U_n is a weighted average of $a/(a+b)$ (the mean of the prior distribution) and M_n (the maximum likelihood estimator).

Another Bernoulli Distribution

Bayesian estimation, like other forms of parametric estimation, depends critically on the parameter space. Suppose again that (X_1, X_2, \dots) is a sequence of Bernoulli trials, given the unknown success parameter p , but suppose now that the parameter space is $\{\frac{1}{2}, 1\}$. This setup corresponds to the tossing of a coin that is either fair or two-headed, but we don't know which. We model p with a random variable P that has the prior probability density function h given by $h(1) = a$, $h(\frac{1}{2}) = 1-a$, where $a \in (0, 1)$ is chosen to reflect our prior knowledge of the probability that the coin is two-headed. If we are completely ignorant, we might let $a = \frac{1}{2}$ (the non-informative prior). If we think the coin is more likely to be two-headed, we might let $a = \frac{3}{4}$. Again let $Y_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}_+$.

The posterior distribution of P given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is

1. $h(1 | \mathbf{X}_n) = \frac{2^n a}{2^n a + (1-a)}$ if $Y_n = n$ and $h(1 | \mathbf{X}_n) = 0$ if $Y_n < n$
2. $h(\frac{1}{2} | \mathbf{X}_n) = \frac{1-a}{2^n a + (1-a)}$ if $Y_n = n$ and $h(\frac{1}{2} | \mathbf{X}_n) = 1$ if $Y_n < n$

Proof

Fix $n \in \mathbb{N}_+$. Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, and let $y = \sum_{i=1}^n x_i$. As before,

$$f(\mathbf{x} | p) = p^y (1-p)^{n-y} \quad (7.4.19)$$

We adopt the usual conventions (which gives the correct mathematics) that $0^k = 0$ if $k \in \mathbb{N}_+$ but $0^0 = 1$. So from Bayes' theorem,

$$h(1 | \mathbf{x}) = \frac{h(1)f(\mathbf{x} | 1)}{h(1/2)f(\mathbf{x} | 1/2) + h(1)f(\mathbf{x} | 1)} \quad (7.4.20)$$

$$= \frac{a 1^y 0^{n-y}}{(1-a)(1/2)^n + a 1^y 0^{n-y}} \quad (7.4.21)$$

So if $y < n$ then $h(1 | \mathbf{x}) = 0$ while if $y = n$

$$h(1 | \mathbf{x}) = \frac{a}{(1-a)(1/2)^n + a} \quad (7.4.22)$$

Of course, $h(\frac{1}{2} | \mathbf{x}) = 1 - h(1 | \mathbf{x})$. The results now follow after a bit of algebra.

Now let

$$p_n = \frac{2^{n+1}a + (1-a)}{2^{n+1}a + 2(1-a)} \quad (7.4.23)$$

The Bayes' estimator of p given \mathbf{X}_n the statistic U_n defined by

1. $U_n = p_n$ if $Y_n = n$
2. $U_n = \frac{1}{2}$ if $Y_n < n$

Proof

By definition, the Bayes' estimator is $U_n = E(P | \mathbf{X}_n)$. From the previous result, if $Y_n = n$ then

$$U_n = 1 \cdot \frac{2^n a}{2^n a + (1-a)} + \frac{1}{2} \cdot \frac{1-a}{2^n a + (1-a)} \quad (7.4.24)$$

which simplifies to p_n . If $Y_n < n$ then $U = 1 \cdot 0 + \frac{1}{2} \cdot 1 = \frac{1}{2}$.

If we observe $Y_n < n$ then U_n gives the correct answer $\frac{1}{2}$. This certainly makes sense since we know that we do not have the two-headed coin. On the other hand, if we observe $Y_n = n$ then we are not certain which coin we have, and the Bayesian estimate p_n is not even in the parameter space! But note that $p_n \rightarrow 1$ as $n \rightarrow \infty$ exponentially fast. Next let's compute the bias and mean-square error for a given $p \in \{\frac{1}{2}, 1\}$.

For $n \in \mathbb{N}_+$,

1. $\text{bias}(U_n | 1) = p_n - 1$
2. $\text{bias}(U_n | \frac{1}{2}) = (\frac{1}{2})^n (p_n - \frac{1}{2})$

The sequence of estimators $(U_n : n \in \mathbb{N}_+)$ is asymptotically unbiased.

Proof

By definition, $\text{bias}(U_n | p) = E(U - p | p)$. Hence from the previous result,

$$\text{bias}(U | p) = (p_n - p)\mathbb{P}(Y = n | p) + \left(\frac{1}{2} - p\right)\mathbb{P}(Y < n | p) \quad (7.4.25)$$

$$= (p_n - p)p^n + \left(\frac{1}{2} - p\right)(1 - p^n) \quad (7.4.26)$$

Substituting $p = 1$ and $p = \frac{1}{2}$ gives the results. In both cases, $\text{bias}(U_n | p) \rightarrow 0$ as $n \rightarrow \infty$ since $p_n \rightarrow 1$ and $(\frac{1}{2})^n \rightarrow 0$ as $n \rightarrow \infty$.

If $p = 1$, the estimator U_n is negatively biased; we noted this earlier. If $p = \frac{1}{2}$, then U_n is positively biased for sufficiently large n (depending on a).

For $n \in \mathbb{N}_+$,

1. $\text{mse}(U_n | 1) = (p_n - 1)^2$
2. $\text{mse}(U_n | \frac{1}{2}) = (\frac{1}{2})^n (p_n - \frac{1}{2})^2$

The sequence of estimators $U = (U_n : n \in \mathbb{N}_+)$ is mean-square consistent.

Proof

By definition, $\text{mse}(U_n | p) = \mathbb{E}[(U_n - p)^2 | p]$. Hence

$$\text{mse}(U_n | p) = (p_n - p)^2\mathbb{P}(Y_n = n | p) + \left(\frac{1}{2} - p\right)^2\mathbb{P}(Y_n < n | p) \quad (7.4.27)$$

$$= (p_n - p)^2p^n + \left(\frac{1}{2} - p\right)^2(1 - p^n) \quad (7.4.28)$$

Substituting $p = 1$ and $p = \frac{1}{2}$ gives the results. In both cases, $\text{mse}(U_n | p) \rightarrow 0$ as $n \rightarrow \infty$ since $p_n \rightarrow 1$ and $(\frac{1}{2})^n \rightarrow 0$ as $n \rightarrow \infty$.

The Geometric distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables, each having the geometric distribution on \mathbb{N}_+ with unknown success parameter $p \in (0, 1)$. Recall that these variables can be interpreted as the number of trials between successive successes in a sequence of Bernoulli trials. Given p , the geometric distribution has probability density function

$$g(x | p) = p(1 - p)^{x-1}, \quad x \in \mathbb{N}_+ \quad (7.4.29)$$

Once again for $n \in \mathbb{N}_+$, let $Y_n = \sum_{i=1}^n X_i$. In this setting, Y_n is the trial number of the n th success, and given p , has the negative binomial distribution with parameters n and p .

Suppose now that we model p with a random variable P having a prior beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. As usual, a and b are chosen to reflect our prior knowledge of p .

The posterior distribution of P given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is beta with left parameter $a+n$ and right parameter $b+(Y_n-n)$.

Proof

Fix $n \in \mathbb{N}_+$. Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}_+^n$ and let $y = \sum_{i=1}^n x_i$. Then

$$f(\mathbf{x} | p) = g(x_1 | p)g(x_2 | p) \cdots g(x_n | p) = p^n(1-p)^{y-n} \quad (7.4.30)$$

Hence

$$h(p)f(\mathbf{x} | p) = \frac{1}{B(a, b)} p^{a-1}(1-p)^{b-1} p^n(1-p)^{y-n} = \frac{1}{B(a, b)} p^{a+n-1}(1-p)^{b+y-n-1}, \quad p \in (0, 1) \quad (7.4.31)$$

As a function of $p \in (0, 1)$ this expression is proportional to the beta PDF with parameters $a+n$ and $b+y-n$. Note that it's not necessary to compute the normalizing constant $f(\mathbf{x})$.

Thus, the beta distribution is conjugate to the geometric distribution. Moreover, note that in the posterior beta distribution, the left parameter is increased by the number of successes n while the right parameter is increased by the number of failures $Y-n$, just as in the [Bernoulli model](#). In particular, the posterior left parameter is deterministic and depends on the data only through the sample size n .

The Bayesian estimator of p based on \mathbf{X}_n is

$$V_n = \frac{a+n}{a+b+Y_n} \quad (7.4.32)$$

Proof

By definition, the Bayesian estimator is the mean of the posterior distribution. Recall again that the mean of the beta distribution is the left parameter divided by the sum of the parameters, so the result follows from our previous theorem.

Recall that the method of moments estimator of p , and the maximum likelihood estimator of p on the interval $(0, 1)$ are both $W_n = 1/M_n = n/Y_n$. To see the connection between the estimators, note from (19) that

$$\frac{1}{V_n} = \frac{a}{a+n} \frac{a+b}{a} + \frac{n}{a+n} \frac{1}{W_n} \quad (7.4.33)$$

So $1/V_n$ (the reciprocal of the Bayesian estimator) is a weighted average of $(a+b)/a$ (the reciprocal of the mean of the prior distribution) and $1/W_n$ (the reciprocal of the maximum likelihood estimator).

The Poisson Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of random variable each having the *Poisson distribution* with unknown parameter $\lambda \in (0, \infty)$. Recall that the Poisson distribution is often used to model the number of “random points” in a region of time or space and is studied in more detail in the chapter on the Poisson Process. The distribution is named for the inimitable Simeon Poisson and given λ , has probability density function

$$g(x | \lambda) = e^{-\lambda} \frac{\lambda^x}{x!}, \quad x \in \mathbb{N} \quad (7.4.34)$$

Once again, for $n \in \mathbb{N}_+$, let $Y_n = \sum_{i=1}^n X_i$. Given λ , random variable Y_n also has a Poisson distribution, but with parameter $n\lambda$.

Suppose now that we model λ with a random variable Λ having a prior gamma distribution with shape parameter $k \in (0, \infty)$ and rate parameter $r \in (0, \infty)$. As usual k and r are chosen to reflect our prior knowledge of λ . Thus the prior probability density function of Λ is

$$h(\lambda) = \frac{r^k}{\Gamma(k)} \lambda^{k-1} e^{-r\lambda}, \quad \lambda \in (0, \infty) \quad (7.4.35)$$

and the mean is k/r . The scale parameter of the gamma distribution is $b = 1/r$, but the formulas will work out nicer if we use the rate parameter.

The posterior distribution of Λ given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is gamma with shape parameter $k + Y_n$ and rate parameter $r + n$.

Proof

Fix $n \in \mathbb{N}_+$. Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}^n$ and $y = \sum_{i=1}^n x_i$. Then

$$f(\mathbf{x} \mid \lambda) = g(x_1 \mid \lambda)g(x_2 \mid \lambda) \cdots g(x_n \mid \lambda) = e^{-n\lambda} \frac{\lambda^y}{x_1!x_2! \cdots x_n!} \quad (7.4.36)$$

Hence

$$h(\lambda)f(\mathbf{x} \mid \lambda) = \frac{r^k}{\Gamma(k)} \lambda^{k-1} e^{-r\lambda} e^{-n\lambda} \frac{\lambda^y}{x_1!x_2! \cdots x_n!} \quad (7.4.37)$$

$$= \frac{r^k}{\Gamma(k)x_1!x_2! \cdots x_n!} e^{-(r+n)\lambda} \lambda^{k+y-1}, \quad \lambda \in (0, \infty) \quad (7.4.38)$$

As a function of $\lambda \in (0, \infty)$ the last expression is proportional to the gamma PDF with shape parameter $k + y$ and rate parameter $r + n$. Note again that it's not necessary to compute the normalizing constant $f(\mathbf{x})$.

It follows that the gamma distribution is conjugate to the Poisson distribution. Note that the posterior rate parameter is deterministic and depends on the data only through the sample size n .

The Bayesian estimator of λ based on $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is

$$V_n = \frac{k + Y_n}{r + n} \quad (7.4.39)$$

Proof

By definition, the Bayes estimator is the mean of the posterior distribution. Recall that mean of the gamma distribution is the shape parameter divided by the rate parameter.

Since V_n is a linear function of Y_n , and we know the distribution of Y_n given $\lambda \in (0, \infty)$, we can compute the bias and mean-square error functions.

For $n \in \mathbb{N}_+$,

$$\text{bias}(V_n \mid \lambda) = \frac{k - r\lambda}{r + n}, \quad \lambda \in (0, \infty) \quad (7.4.40)$$

The sequence of estimators $\mathbf{V} = (V_n : n \in \mathbb{N}_+)$ is asymptotically unbiased.

Proof

The computation is simple, since the distribution of Y_n given λ is Poisson with parameter $n\lambda$.

$$\text{bias}(V_n \mid \lambda) = \mathbb{E}(V_n \mid \lambda) - \lambda = \frac{k + n\lambda}{r + n} - \lambda = \frac{k - r\lambda}{r + n} \quad (7.4.41)$$

Clearly $\text{bias}(V_n \mid \lambda) \rightarrow 0$ as $n \rightarrow \infty$.

Note that, as before, we cannot choose k and r to make V_n unbiased, without knowledge of λ .

For $n \in \mathbb{N}_+$,

$$\text{mse}(V_n \mid \lambda) = \frac{n\lambda + (k - r\lambda)^2}{(r + n)^2}, \quad \lambda \in (0, \infty) \quad (7.4.42)$$

The sequence of estimators $\mathbf{V} = (V_n : n \in \mathbb{N}_+)$ is mean-square consistent.

Proof

Again, the computation is easy since the distribution of Y_n given λ is Poisson with parameter $n\lambda$.

$$\text{mse}(V | \lambda) = \text{var}(V_n | \lambda) + \text{bias}^2(V_n | \lambda) = \frac{n\lambda}{(r+n)^2} + \left(\frac{k-r\lambda}{r+n} \right)^2 \quad (7.4.43)$$

Clearly $\text{mse}(V_n | \lambda) \rightarrow 0$ as $n \rightarrow \infty$.

Recall that the method of moments estimator of λ and the maximum likelihood estimator of λ on the interval $(0, \infty)$ are both $M_n = Y_n/n$, the sample mean. This estimator is unbiased and has mean square error λ/n . To see the connection between the estimators, note from (21) that

$$V_n = \frac{r}{r+n} \frac{k}{r} + \frac{n}{r+n} M_n \quad (7.4.44)$$

So V_n is a weighted average of k/r (the mean of the prior distribution) and M_n (the maximum likelihood estimator).

The Normal Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables, each having the normal distribution with unknown mean $\mu \in \mathbb{R}$ but known variance $\sigma^2 \in (0, \infty)$. Of course, the normal distribution plays an especially important role in statistics, in part because of the central limit theorem. The normal distribution is widely used to model physical quantities subject to numerous small, random errors. In many statistical applications, the variance of the normal distribution is more stable than the mean, so the assumption that the variance is known is not entirely artificial. Recall that the normal probability density function (given μ) is

$$g(x | \mu) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2} \left(\frac{x-\mu}{\sigma} \right)^2 \right], \quad x \in \mathbb{R} \quad (7.4.45)$$

Again, for $n \in \mathbb{N}_+$ let $Y_n = \sum_{i=1}^n X_i$. Recall that Y_n also has a normal distribution (given μ) but with mean $n\mu$ and variance $n\sigma^2$.

Suppose now that μ is modeled by a random variable Ψ that has a prior normal distribution with mean $a \in \mathbb{R}$ and variance $b^2 \in (0, \infty)$. As usual, a and b are chosen to reflect our prior knowledge of μ . An interesting special case is when we take $b = \sigma$, so the variance of the prior distribution of Ψ is the same as the variance of the underlying sampling distribution.

For $n \in \mathbb{N}_+$, the posterior distribution of Ψ given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is normal with mean and variance given by

$$\mathbb{E}(\Psi | \mathbf{X}_n) = \frac{Y_n b^2 + a \sigma^2}{n b^2 + \sigma^2} \quad (7.4.46)$$

$$\text{var}(\Psi | \mathbf{X}_n) = \frac{\sigma^2 b^2}{n b^2 + \sigma^2} \quad (7.4.47)$$

Proof

Fix $n \in \mathbb{N}_+$. Suppose $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}$ and let $y = \sum_{i=1}^n x_i$ and $w^2 = \sum_{i=1}^n x_i^2$. Then

$$f(\mathbf{x} | \mu) = g(x_1 | \mu) g(x_2 | \mu) \cdots g(x_n | \mu) = \frac{1}{(2\pi)^{n/2} \sigma^n} \exp \left[-\frac{1}{2} \sum_{i=1}^n \left(\frac{x_i - \mu}{\sigma} \right)^2 \right] \quad (7.4.48)$$

$$= \frac{1}{(2\pi)^{n/2} \sigma^n} \exp \left[-\frac{1}{2\sigma^2} (w^2 - 2\mu y + n\mu^2) \right] \quad (7.4.49)$$

On the other hand, of course

$$h(\mu) = \frac{1}{\sqrt{2\pi}b} \exp \left[-\frac{1}{2} \left(\frac{\mu - a}{b} \right)^2 \right] = \frac{1}{\sqrt{2\pi}b} \exp \left[-\frac{1}{2b^2} (\mu^2 - 2a\mu + a^2) \right] \quad (7.4.50)$$

Therefore,

$$h(\mu) f(\mathbf{x} | \mu) = C \exp \left\{ -\frac{1}{2} \left[\left(\frac{1}{b^2} + \frac{n}{\sigma^2} \right) \mu^2 - 2 \left(\frac{a}{b^2} + \frac{y}{\sigma^2} \right) \mu \right] \right\} \quad (7.4.51)$$

where C depends on $n, \sigma, a, b, \mathbf{x}$, but importantly *not* on μ . So we don't really care what C is. Completing the square in μ in the expression above gives

$$h(\mu)f(\mathbf{x} | \mu) = K \exp \left[-\frac{1}{2} \left(\frac{1}{b^2} + \frac{n}{\sigma^2} \right) \left(\mu - \frac{a/b^2 + y/\sigma^2}{1/b^2 + n/\sigma^2} \right)^2 \right] \quad (7.4.52)$$

where K is yet another factor that depends on lots of stuff, but not μ . As a function of μ , this expression is proportional to the normal distribution with mean and variance, respectively, given by

$$\frac{a/b^2 + y/\sigma^2}{1/b^2 + n/\sigma^2} = \frac{yb^2 + a\sigma^2}{nb^2 + \sigma^2} \quad (7.4.53)$$

$$\frac{1}{1/b^2 + n/\sigma^2} = \frac{\sigma^2 b^2}{\sigma^2 + nb^2} \quad (7.4.54)$$

Once again, it was not necessary to compute the normalizing constant $f(\mathbf{x})$, which would have been yet another factor that we do not care about.

Therefore, the normal distribution is conjugate to the normal distribution with unknown mean and known variance. Note that the posterior variance is deterministic, and depends on the data only through the sample size n . In the special case that $b = \sigma$, the posterior distribution of Ψ given \mathbf{X}_n is normal with mean $(Y_n + a)/(n + 1)$ and variance $\sigma^2/(n + 1)$.

The Bayesian estimator of μ is

$$U_n = \frac{Y_n b^2 + a\sigma^2}{nb^2 + \sigma^2} \quad (7.4.55)$$

Proof

This follows immediately from the previous result.

Note that $U_n = (Y_n + a)/(n + 1)$ in the special case that $b = \sigma$.

For $n \in \mathbb{N}_+$,

$$\text{bias}(U_n | \mu) = \frac{\sigma^2(a - \mu)}{\sigma^2 + nb^2}, \quad \mu \in \mathbb{R} \quad (7.4.56)$$

The sequence of estimators $\mathbf{U} = (U_n : n \in \mathbb{N}_+)$ is asymptotically unbiased.

Proof

Recall that Y_n has mean $n\mu$ given μ . Hence

$$\text{bias}(U_n | \mu) = \mathbb{E}(U_n | \mu) - \mu = \frac{nb^2\mu + a\sigma^2}{nb^2 + \sigma^2} - \mu = \frac{(a - \mu)\sigma^2}{nb^2 + \sigma^2} \quad (7.4.57)$$

Clearly $\text{bias}(U_n | \mu) \rightarrow 0$ as $n \rightarrow \infty$ for every $\mu \in \mathbb{R}$.

When $b = \sigma$, $\text{bias}(U_n | \mu) = (a - \mu)/(n + 1)$.

For $n \in \mathbb{N}_+$,

$$\text{mse}(U_n | \mu) = \frac{n\sigma^2 b^4 + \sigma^4(a - \mu)^2}{(\sigma^2 + nb^2)^2}, \quad \mu \in \mathbb{R} \quad (7.4.58)$$

The sequence of estimators $\mathbf{U} = (U_n : n \in \mathbb{N}_+)$ is mean-square consistent.

Proof

Recall that Y_n has variance $n\sigma^2$. Hence

$$\text{mse}(U_n | \mu) = \text{var}(U_n | \mu) + \text{bias}^2(U_n | \mu) = \left(\frac{b^2}{nb^2 + \sigma^2} \right)^2 n\sigma^2 + \left(\frac{(a - \mu)\sigma^2}{nb^2 + \sigma^2} \right)^2 \quad (7.4.59)$$

Clearly $\text{mse}(U_n | \mu) \rightarrow 0$ as $n \rightarrow \infty$ for every $\mu \in \mathbb{R}$.

When $b = \sigma$, $\text{mse}(U | \mu) = [n\sigma^2 + (a - \mu)^2]/(n + 1)^2$. Recall that the method of moments estimator of μ and the maximum likelihood estimator of μ on \mathbb{R} are both $M_n = Y_n/n$, the sample mean. This estimator is unbiased and has mean square error $\text{var}(M) = \sigma^2/n$. To see the connection between the estimators, note from (25) that

$$U_n = \frac{\sigma^2}{nb^2 + \sigma^2} a + \frac{nb^2}{nb^2 + \sigma^2} M_n \quad (7.4.60)$$

So U_n is a weighted average of a (the mean of the prior distribution) and M_n (the maximum likelihood estimator).

The Beta Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables each having the beta distribution with unknown left shape parameter $a \in (0, \infty)$ and right shape parameter $b = 1$. The beta distribution is widely used to model random proportions and probabilities and other variables that take values in bounded intervals (scaled to take values in $(0, 1)$). Recall that the probability density function (given a) is

$$g(x | a) = a x^{a-1}, \quad x \in (0, 1) \quad (7.4.61)$$

Suppose now that a is modeled by a random variable A that has a prior gamma distribution with shape parameter $k \in (0, \infty)$ and rate parameter $r \in (0, \infty)$. As usual, k and r are chosen to reflect our prior knowledge of a . Thus the prior probability density function of A is

$$h(a) = \frac{r^k}{\Gamma(k)} a^{k-1} e^{-ra}, \quad a \in (0, \infty) \quad (7.4.62)$$

The mean of the prior distribution is k/r .

The posterior distribution of A given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is gamma, with shape parameter $k + n$ and rate parameter $r - \ln(X_1 X_2 \cdots X_n)$.

Proof

Fix $n \in \mathbb{N}_+$. Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in (0, 1)^n$ and let $z = x_1 x_2 \cdots x_n$. Then

$$f(\mathbf{x} | a) = g(x_1 | a)g(x_2 | a) \cdots g(x_n | a) = a^n z^{a-1} = \frac{a^n}{z} e^{a \ln z} \quad (7.4.63)$$

Hence

$$h(a)f(\mathbf{x} | a) = \frac{r^k}{z\Gamma(k)} a^{n+k-1} e^{-a(r - \ln z)}, \quad a \in (0, \infty) \quad (7.4.64)$$

As a function of $a \in (0, \infty)$ this expression is proportional to the gamma PDF with shape parameter $n + k$ and scale parameter $r - \ln z$. Once again, it's not necessary to compute the normalizing constant $f(\mathbf{x})$.

Thus, the gamma distribution is conjugate to the beta distribution with unknown left parameter and right parameter 1. Note that the posterior shape parameter is deterministic and depends on the data only through the sample size n .

The Bayesian estimator of a based on \mathbf{X}_n is

$$U_n = \frac{k + n}{r - \ln(X_1 X_2 \cdots X_n)} \quad (7.4.65)$$

Proof

The mean of the gamma distribution is the shape parameter divided by the rate parameter, so this follows from the previous theorem.

Given the complicated structure, the bias and mean square error of U_n given $a \in (0, \infty)$ would be difficult to compute explicitly. Recall that the maximum likelihood estimator of a is $W_n = -n / \ln(X_1 X_2 \cdots X_n)$. To see the connection between the estimators, note from (29) that

$$\frac{1}{U_n} = \frac{k}{k+n} \frac{r}{k} + \frac{n}{k+n} \frac{1}{W_n} \quad (7.4.66)$$

So $1/U_n$ (the reciprocal of the Bayesian estimator) is a weighted average of r/k (the reciprocal of the mean of the prior distribution) and $1/W_n$ (the reciprocal of the maximum likelihood estimator).

The Pareto Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables each having the Pareto distribution with unknown shape parameter $a \in (0, \infty)$ and scale parameter $b = 1$. The Pareto distribution is used to model certain financial variables and other variables with heavy-tailed distributions, and is named for Vilfredo Pareto. Recall that the probability density function (given a) is

$$g(x | a) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (7.4.67)$$

Suppose now that a is modeled by a random variable A that has a prior gamma distribution with shape parameter $k \in (0, \infty)$ and rate parameter $r \in (0, \infty)$. As usual, k and r are chosen to reflect our prior knowledge of a . Thus the prior probability density function of A is

$$h(a) = \frac{r^k}{\Gamma(k)} a^{k-1} e^{-ra}, \quad a \in (0, \infty) \quad (7.4.68)$$

For $n \in \mathbb{N}_+$, the posterior distribution of A given $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is gamma, with shape parameter $k+n$ and rate parameter $r + \ln(X_1 X_2 \cdots X_n)$.

Proof

Fix $n \in \mathbb{N}_+$. Let $\mathbf{x} = (x_1, x_2, \dots, x_n) \in [1, \infty)^n$ and let $z = x_1 x_2 \cdots x_n$. Then

$$f(\mathbf{x} | a) = g(x_1 | a)g(x_2 | a) \cdots g(x_n | a) = \frac{a^n}{z^{a+1}} = \frac{a^n}{z} e^{-a \ln z} \quad (7.4.69)$$

Hence

$$h(a)f(\mathbf{x} | a) = \frac{r^k}{z\Gamma(k)} a^{n+k-1} e^{-a(r+\ln z)}, \quad a \in (0, \infty) \quad (7.4.70)$$

As a function of $a \in (0, \infty)$ this expression is proportional to the gamma PDF with shape parameter $n+k$ and scale parameter $r + \ln z$. Once again, it's not necessary to compute the normalizing constant $f(\mathbf{x})$.

Thus, the gamma distribution is conjugate to Pareto distribution with unknown shape parameter. Note that the posterior shape parameter is deterministic and depends on the data only through the sample size n .

The Bayesian estimator of a based on \mathbf{X}_n is

$$U_n = \frac{k+n}{r + \ln(X_1 X_2 \cdots X_n)} \quad (7.4.71)$$

Proof

Once again, the mean of the gamma distribution is the shape parameter divided by the rate parameter, so this follows from the previous theorem.

Given the complicated structure, the bias and mean square error of U given $a \in (0, \infty)$ would be difficult to compute explicitly. Recall that the maximum likelihood estimator of a is $W_n = n / \ln(X_1 X_2 \cdots X_n)$. To see the connection between the estimators, note from (31) that

$$\frac{1}{U_n} = \frac{k}{k+n} \frac{r}{k} + \frac{n}{k+n} \frac{1}{W_n} \quad (7.4.72)$$

So $1/U_n$ (the reciprocal of the Bayesian estimator) is a weighted average of r/k (the reciprocal of the mean of the prior distribution) and $1/W_n$ (the maximum likelihood estimator).

This page titled [7.4: Bayesian Estimation](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.5: Best Unbiased Estimators

Basic Theory

Consider again the basic statistical model, in which we have a random experiment that results in an observable random variable \mathbf{X} taking values in a set S . Once again, the experiment is typically to sample n objects from a population and record one or more measurements for each item. In this case, the observable random variable has the form

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (7.5.1)$$

where X_i is the vector of measurements for the i th item.

Suppose that θ is a real parameter of the distribution of \mathbf{X} , taking values in a parameter space Θ . Let f_θ denote the probability density function of \mathbf{X} for $\theta \in \Theta$. Note that the expected value, variance, and covariance operators also depend on θ , although we will sometimes suppress this to keep the notation from becoming too unwieldy.

Definitions

Suppose now that $\lambda = \lambda(\theta)$ is a parameter of interest that is derived from θ . (Of course, λ might be θ itself, but more generally might be a function of θ .) In this section we will consider the general problem of finding the best estimator of λ among a given class of unbiased estimators. Recall that if U is an unbiased estimator of λ , then $\text{var}_\theta(U)$ is the mean square error. Mean square error is our measure of the quality of unbiased estimators, so the following definitions are natural.

Suppose that U and V are unbiased estimators of λ .

1. If $\text{var}_\theta(U) \leq \text{var}_\theta(V)$ for all $\theta \in \Theta$ then U is a *uniformly better* estimator than V .
2. If U is uniformly better than every other unbiased estimator of λ , then U is a *Uniformly Minimum Variance Unbiased Estimator (UMVUE)* of λ .

Given unbiased estimators U and V of λ , it may be the case that U has smaller variance for some values of θ while V has smaller variance for other values of θ , so that neither estimator is uniformly better than the other. Of course, a minimum variance unbiased estimator is the best we can hope for.

The Cramér-Rao Lower Bound

We will show that under mild conditions, there is a lower bound on the variance of any unbiased estimator of the parameter λ . Thus, if we can find an estimator that achieves this lower bound for all θ , then the estimator must be an UMVUE of λ . The derivative of the log likelihood function, sometimes called the *score*, will play a critical role in our analysis. A lesser, but still important role, is played by the negative of the second derivative of the log-likelihood function. Life will be much easier if we give these functions names.

For $\mathbf{x} \in S$ and $\theta \in \Theta$, define

$$L_1(\mathbf{x}, \theta) = \frac{d}{d\theta} \ln(f_\theta(\mathbf{x})) \quad (7.5.2)$$

$$L_2(\mathbf{x}, \theta) = -\frac{d}{d\theta} L_1(\mathbf{x}, \theta) = -\frac{d^2}{d\theta^2} \ln(f_\theta(\mathbf{x})) \quad (7.5.3)$$

In the rest of this subsection, we consider statistics $h(\mathbf{X})$ where $h : S \rightarrow \mathbb{R}$ (and so in particular, h does not depend on θ). We need a fundamental assumption:

We will consider only statistics $h(\mathbf{X})$ with $\mathbb{E}_\theta(h^2(\mathbf{X})) < \infty$ for $\theta \in \Theta$. We also assume that

$$\frac{d}{d\theta} \mathbb{E}_\theta(h(\mathbf{X})) = \mathbb{E}_\theta(h(\mathbf{X})L_1(\mathbf{X}, \theta)) \quad (7.5.4)$$

This is equivalent to the assumption that the derivative operator $d/d\theta$ can be interchanged with the expected value operator \mathbb{E}_θ .

Proof

Note first that

$$\frac{d}{d\theta} \mathbb{E}(h(\mathbf{X})) = \frac{d}{d\theta} \int_S h(\mathbf{x}) f_{\theta}(\mathbf{x}) d\mathbf{x} \quad (7.5.5)$$

On the other hand,

$$\mathbb{E}_{\theta}(h(\mathbf{X})L_1(\mathbf{X}, \theta)) = \mathbb{E}_{\theta}\left(h(\mathbf{X})\frac{d}{d\theta}\ln(f_{\theta}(\mathbf{X}))\right) = \int_S h(\mathbf{x})\frac{d}{d\theta}\ln(f_{\theta}(\mathbf{x}))f_{\theta}(\mathbf{x}) d\mathbf{x} \quad (7.5.6)$$

$$= \int_S h(\mathbf{x})\frac{\frac{d}{d\theta}f_{\theta}(\mathbf{x})}{f_{\theta}(\mathbf{x})}f_{\theta}(\mathbf{x}) d\mathbf{x} = \int_S h(\mathbf{x})\frac{d}{d\theta}f_{\theta}(\mathbf{x}) d\mathbf{x} = \int_S \frac{d}{d\theta}h(\mathbf{x})f_{\theta}(\mathbf{x}) d\mathbf{x} \quad (7.5.7)$$

Thus the two expressions are the same if and only if we can interchange the derivative and integral operators.

Generally speaking, the fundamental assumption will be satisfied if $f_{\theta}(\mathbf{x})$ is differentiable as a function of θ , with a derivative that is jointly continuous in \mathbf{x} and θ , and if the support set $\{\mathbf{x} \in S : f_{\theta}(\mathbf{x}) > 0\}$ does not depend on θ .

$\mathbb{E}_{\theta}(L_1(\mathbf{X}, \theta)) = 0$ for $\theta \in \Theta$.

Proof

This follows from the [fundamental assumption](#) by letting $h(\mathbf{x}) = 1$ for $\mathbf{x} \in S$.

If $h(\mathbf{X})$ is a statistic then

$$\text{cov}_{\theta}(h(\mathbf{X}), L_1(\mathbf{X}, \theta)) = \frac{d}{d\theta} \mathbb{E}_{\theta}(h(\mathbf{X})) \quad (7.5.8)$$

Proof

First note that the covariance is simply the expected value of the product of the variables, since the second variable has mean 0 by the [previous theorem](#). The result then follows from the [basic condition](#).

$\text{var}_{\theta}(L_1(\mathbf{X}, \theta)) = \mathbb{E}_{\theta}(L_1^2(\mathbf{X}, \theta))$

Proof

This follows since $L_1(\mathbf{X}, \theta)$ has mean 0 by the [theorem above](#).

The following theorem gives the general *Cramér-Rao lower bound* on the variance of a statistic. The lower bound is named for Harold Cramér and CR Rao:

If $h(\mathbf{X})$ is a statistic then

$$\text{var}_{\theta}(h(\mathbf{X})) \geq \frac{\left(\frac{d}{d\theta} \mathbb{E}_{\theta}(h(\mathbf{X}))\right)^2}{\mathbb{E}_{\theta}(L_1^2(\mathbf{X}, \theta))} \quad (7.5.9)$$

Proof

From the Cauchy-Scharwtz (correlation) inequality,

$$\text{cov}_{\theta}^2(h(\mathbf{X}), L_1(\mathbf{X}, \theta)) \leq \text{var}_{\theta}(h(\mathbf{X})) \text{var}_{\theta}(L_1(\mathbf{X}, \theta)) \quad (7.5.10)$$

The result now follows from the previous two theorems.

We can now give the first version of the Cramér-Rao lower bound for unbiased estimators of a parameter.

Suppose now that $\lambda(\theta)$ is a parameter of interest and $h(\mathbf{X})$ is an unbiased estimator of λ . Then

$$\text{var}_{\theta}(h(\mathbf{X})) \geq \frac{(d\lambda/d\theta)^2}{\mathbb{E}_{\theta}(L_1^2(\mathbf{X}, \theta))} \quad (7.5.11)$$

Proof

This follows immediately from the Cramér-Rao lower bound, since $\mathbb{E}_\theta(h(\mathbf{X})) = \lambda$ for $\theta \in \Theta$.

An estimator of λ that achieves the Cramér-Rao lower bound must be a uniformly minimum variance unbiased estimator (UMVUE) of λ .

Equality holds in the previous theorem, and hence $h(\mathbf{X})$ is an UMVUE, if and only if there exists a function $u(\theta)$ such that (with probability 1)

$$h(\mathbf{X}) = \lambda(\theta) + u(\theta)L_1(\mathbf{X}, \theta) \quad (7.5.12)$$

Proof

Equality holds in the Cauchy-Schwartz inequality if and only if the random variables are linear transformations of each other. Recall also that $L_1(\mathbf{X}, \theta)$ has mean 0.

The quantity $\mathbb{E}_\theta(L^2(\mathbf{X}, \theta))$ that occurs in the denominator of the lower bounds in the previous two theorems is called the *Fisher information number* of \mathbf{X} , named after Sir Ronald Fisher. The following theorem gives an alternate version of the Fisher information number that is usually computationally better.

If the appropriate derivatives exist and if the appropriate interchanges are permissible then

$$\mathbb{E}_\theta(L_1^2(\mathbf{X}, \theta)) = \mathbb{E}_\theta(L_2(\mathbf{X}, \theta)) \quad (7.5.13)$$

The following theorem gives the second version of the Cramér-Rao lower bound for unbiased estimators of a parameter.

If $\lambda(\theta)$ is a parameter of interest and $h(\mathbf{X})$ is an unbiased estimator of λ then

$$\text{var}_\theta(h(\mathbf{X})) \geq \frac{(d\lambda/d\theta)^2}{\mathbb{E}_\theta(L_2(\mathbf{X}, \theta))} \quad (7.5.14)$$

Proof

This follows from the results above.

Random Samples

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of a random variable X having probability density function g_θ and taking values in a set R . Thus $S = R^n$. We will use lower-case letters for the derivative of the log likelihood function of X and the negative of the second derivative of the log likelihood function of X .

For $x \in R$ and $\theta \in \Theta$ define

$$l(x, \theta) = \frac{d}{d\theta} \ln(g_\theta(x)) \quad (7.5.15)$$

$$l_2(x, \theta) = -\frac{d^2}{d\theta^2} \ln(g_\theta(x)) \quad (7.5.16)$$

L^2 can be written in terms of l^2 and L_2 can be written in terms of l_2 :

1. $\mathbb{E}_\theta(L^2(\mathbf{X}, \theta)) = n\mathbb{E}_\theta(l^2(X, \theta))$
2. $\mathbb{E}_\theta(L_2(\mathbf{X}, \theta)) = n\mathbb{E}_\theta(l_2(X, \theta))$

The following theorem gives the second version of the general Cramér-Rao lower bound on the variance of a statistic, specialized for random samples.

If $h(\mathbf{X})$ is a statistic then

$$\text{var}_{\theta}(h(\mathbf{X})) \geq \frac{\left(\frac{d}{d\theta} \mathbb{E}_{\theta}(h(\mathbf{X}))\right)^2}{n \mathbb{E}_{\theta}(l^2(X, \theta))} \quad (7.5.17)$$

The following theorem give the third version of the Cramér-Rao lower bound for unbiased estimators of a parameter, specialized for random samples.

Suppose now that $\lambda(\theta)$ is a parameter of interest and $h(\mathbf{X})$ is an unbiased estimator of λ . Then

$$\text{var}_{\theta}(h(\mathbf{X})) \geq \frac{(d\lambda/d\theta)^2}{n \mathbb{E}_{\theta}(l^2(X, \theta))} \quad (7.5.18)$$

Note that the Cramér-Rao lower bound varies inversely with the sample size n . The following version gives the fourth version of the Cramér-Rao lower bound for unbiased estimators of a parameter, again specialized for random samples.

If the appropriate derivatives exist and the appropriate interchanges are permissible) then

$$\text{var}_{\theta}(h(\mathbf{X})) \geq \frac{(d\lambda/d\theta)^2}{n \mathbb{E}_{\theta}(l_2(X, \theta))} \quad (7.5.19)$$

To summarize, we have four versions of the Cramér-Rao lower bound for the variance of an unbiased estimate of λ : [version 1](#) and [version 2](#) in the general case, and [version 1](#) and [version 2](#) in the special case that \mathbf{X} is a random sample from the distribution of X . If an unbiased estimator of λ achieves the lower bound, then the estimator is an UMVUE.

Examples and Special Cases

We will apply the results above to several parametric families of distributions. First we need to recall some standard notation. Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of a real-valued random variable X with mean μ and variance σ^2 . The sample mean is

$$M = \frac{1}{n} \sum_{i=1}^n X_i \quad (7.5.20)$$

Recall that $\mathbb{E}(M) = \mu$ and $\text{var}(M) = \sigma^2/n$. The special version of the sample variance, when μ is known, and standard version of the sample variance are, respectively,

$$W^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 \quad (7.5.21)$$

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M)^2 \quad (7.5.22)$$

The Bernoulli Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Bernoulli distribution with unknown success parameter $p \in (0, 1)$. In the usual language of reliability, $X_i = 1$ means success on trial i and $X_i = 0$ means failure on trial i ; the distribution is named for Jacob Bernoulli. Recall that the Bernoulli distribution has probability density function

$$g_p(x) = p^x(1-p)^{1-x}, \quad x \in \{0, 1\} \quad (7.5.23)$$

The [basic assumption](#) is satisfied. Moreover, recall that the mean of the Bernoulli distribution is p , while the variance is $p(1-p)$.

$p(1-p)/n$ is the [Cramér-Rao lower bound](#) for the variance of unbiased estimators of p .

The sample mean M (which is the proportion of successes) attains the lower bound in the previous exercise and hence is an UMVUE of p .

The Poisson Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Poisson distribution with parameter $\theta \in (0, \infty)$. Recall that this distribution is often used to model the number of “random points” in a region of time or space and is studied in more detail in the chapter on the Poisson Process. The Poisson distribution is named for Simeon Poisson and has probability density function

$$g_\theta(x) = e^{-\theta} \frac{\theta^x}{x!}, \quad x \in \mathbb{N} \quad (7.5.24)$$

The [basic assumption](#) is satisfied. Recall also that the mean and variance of the distribution are both θ .

θ/n is the [Cramér-Rao lower bound](#) for the variance of unbiased estimators of θ .

The sample mean M attains the lower bound in the previous exercise and hence is an UMVUE of θ .

The Normal Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the normal distribution with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$. Recall that the normal distribution plays an especially important role in statistics, in part because of the central limit theorem. The normal distribution is widely used to model physical quantities subject to numerous small, random errors, and has probability density function

$$g_{\mu, \sigma^2}(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R} \quad (7.5.25)$$

The [basic assumption](#) is satisfied with respect to both of these parameters. Recall also that the fourth central moment is $\mathbb{E}((X - \mu)^4) = 3\sigma^4$.

σ^2/n is the [Cramér-Rao lower bound](#) for the variance of unbiased estimators of μ .

The sample mean M attains the lower bound in the previous exercise and hence is an UMVUE of μ .

$\frac{2\sigma^4}{n}$ is the [Cramér-Rao lower bound](#) for the variance of unbiased estimators of σ^2 .

The sample variance S^2 has variance $\frac{2\sigma^4}{n-1}$ and hence does not attain the lower bound in the previous exercise.

If μ is known, then the special sample variance W^2 attains the lower bound [above](#) and hence is an UMVUE of σ^2 .

If μ is unknown, no unbiased estimator of σ^2 attains the Cramér-Rao lower bound [above](#).

Proof

This follows from the result [above](#) on equality in the Cramér-Rao inequality.

The Gamma Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the gamma distribution with known shape parameter $k > 0$ and unknown scale parameter $b > 0$. The gamma distribution is often used to model random times and certain other types of positive random variables, and is studied in more detail in the chapter on Special Distributions. The probability density function is

$$g_b(x) = \frac{1}{\Gamma(k)b^k} x^{k-1} e^{-x/b}, \quad x \in (0, \infty) \quad (7.5.26)$$

The [basic assumption](#) is satisfied with respect to b . Moreover, the mean and variance of the gamma distribution are kb and kb^2 , respectively.

$\frac{b^2}{nk}$ is the [Cramér-Rao lower bound](#) for the variance of unbiased estimators of b .

$\frac{M}{k}$ attains the lower bound in the previous exercise and hence is an UMVUE of b .

The Beta Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the beta distribution with left parameter $a > 0$ and right parameter $b = 1$. Beta distributions are widely used to model random proportions and other random variables that take values in bounded intervals, and are studied in more detail in the chapter on Special Distributions. In our specialized case, the probability density function of the sampling distribution is

$$g_a(x) = a x^{a-1}, \quad x \in (0, 1) \quad (7.5.27)$$

The **basic assumption** is satisfied with respect to a .

The mean and variance of the distribution are

$$\begin{aligned} 1. \mu &= \frac{a}{a+1} \\ 2. \sigma^2 &= \frac{a}{(a+1)^2(a+2)} \end{aligned}$$

The **Cramér-Rao lower bound** for the variance of unbiased estimators of μ is $\frac{a^2}{n(a+1)^4}$.

The sample mean M does not achieve the Cramér-Rao lower bound in the previous exercise, and hence is not an UMVUE of μ .

The Uniform Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the uniform distribution on $[0, a]$ where $a > 0$ is the unknown parameter. Thus, the probability density function of the sampling distribution is

$$g_a(x) = \frac{1}{a}, \quad x \in [0, a] \quad (7.5.28)$$

The **basic assumption** is *not* satisfied.

The **Cramér-Rao lower bound** for the variance of unbiased estimators of a is $\frac{a^2}{n}$. Of course, the Cramér-Rao Theorem does not apply, by the previous exercise.

Recall that $V = \frac{n+1}{n} \max\{X_1, X_2, \dots, X_n\}$ is unbiased and has variance $\frac{a^2}{n(n+2)}$. This variance is smaller than the Cramér-Rao bound in the previous exercise.

The reason that the **basic assumption** is not satisfied is that the support set $\{x \in \mathbb{R} : g_a(x) > 0\}$ depends on the parameter a .

Best Linear Unbiased Estimators

We now consider a somewhat specialized problem, but one that fits the general theme of this section. Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of observable real-valued random variables that are uncorrelated and have the same unknown mean $\mu \in \mathbb{R}$, but possibly different standard deviations. Let $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_n)$ where $\sigma_i = \text{sd}(X_i)$ for $i \in \{1, 2, \dots, n\}$.

We will consider estimators of μ that are linear functions of the outcome variables. Specifically, we will consider estimators of the following form, where the vector of coefficients $\mathbf{c} = (c_1, c_2, \dots, c_n)$ is to be determined:

$$Y = \sum_{i=1}^n c_i X_i \quad (7.5.29)$$

Y is unbiased if and only if $\sum_{i=1}^n c_i = 1$.

The variance of Y is

$$\text{var}(Y) = \sum_{i=1}^n c_i^2 \sigma_i^2 \quad (7.5.30)$$

The variance is minimized, subject to the unbiased constraint, when

$$c_j = \frac{1/\sigma_j^2}{\sum_{i=1}^n 1/\sigma_i^2}, \quad j \in \{1, 2, \dots, n\} \quad (7.5.31)$$

Proof

Use the method of Lagrange multipliers (named after Joseph-Louis Lagrange).

This exercise shows how to construct the *Best Linear Unbiased Estimator (BLUE)* of μ , assuming that the vector of standard deviations σ is known.

Suppose now that $\sigma_i = \sigma$ for $i \in \{1, 2, \dots, n\}$ so that the outcome variables have the same standard deviation. In particular, this would be the case if the outcome variables form a random sample of size n from a distribution with mean μ and standard deviation σ .

In this case the variance is minimized when $c_i = 1/n$ for each i and hence $Y = M$, the sample mean.

This exercise shows that the sample mean M is the best linear unbiased estimator of μ when the standard deviations are the same, and that moreover, we do not need to know the value of the standard deviation.

This page titled [7.5: Best Unbiased Estimators](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

7.6: Sufficient, Complete and Ancillary Statistics

Basic Theory

The Basic Statistical Model

Consider again the basic statistical model, in which we have a random experiment with an observable random variable \mathbf{X} taking values in a set S . Once again, the experiment is typically to sample n objects from a population and record one or more measurements for each item. In this case, the outcome variable has the form

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (7.6.1)$$

where X_i is the vector of measurements for the i th item. In general, we suppose that the distribution of \mathbf{X} depends on a parameter θ taking values in a parameter space T . The parameter θ may also be vector-valued. We will sometimes use subscripts in probability density functions, expected values, etc. to denote the dependence on θ .

As usual, the most important special case is when \mathbf{X} is a sequence of independent, identically distributed random variables. In this case \mathbf{X} is a random sample from the common distribution.

Sufficient Statistics

Let $U = u(\mathbf{X})$ be a statistic taking values in a set R . Intuitively, U is sufficient for θ if U contains all of the information about θ that is available in the entire data variable \mathbf{X} . Here is the formal definition:

A statistic U is *sufficient* for θ if the conditional distribution of \mathbf{X} given U does not depend on $\theta \in T$.

Sufficiency is related to the concept of *data reduction*. Suppose that \mathbf{X} takes values in \mathbb{R}^n . If we can find a sufficient statistic U that takes values in \mathbb{R}^j , then we can reduce the original data vector \mathbf{X} (whose dimension n is usually large) to the vector of statistics U (whose dimension j is usually much smaller) with no loss of information about the parameter θ .

The following result gives a condition for sufficiency that is equivalent to this definition.

Let $U = u(\mathbf{X})$ be a statistic taking values in R , and let f_θ and h_θ denote the probability density functions of \mathbf{X} and U respectively. Then U is sufficient for θ if and only if the function on S given below does not depend on $\theta \in T$:

$$\mathbf{x} \mapsto \frac{f_\theta(\mathbf{x})}{h_\theta[u(\mathbf{x})]} \quad (7.6.2)$$

Proof

The joint distribution of (\mathbf{X}, U) is concentrated on the set $\{(\mathbf{x}, y) : \mathbf{x} \in S, y = u(\mathbf{x})\} \subseteq S \times R$. The conditional PDF of \mathbf{X} given $U = u(\mathbf{x})$ is $f_\theta(\mathbf{x})/h_\theta[u(\mathbf{x})]$ on this set, and is 0 otherwise.

The definition precisely captures the intuitive notion of sufficiency given above, but can be difficult to apply. We must know in advance a candidate statistic U , and then we must be able to compute the conditional distribution of \mathbf{X} given U . The *Fisher-Neyman factorization theorem* given next often allows the identification of a sufficient statistic from the form of the probability density function of \mathbf{X} . It is named for Ronald Fisher and Jerzy Neyman.

Fisher-Neyman Factorization Theorem. Let f_θ denote the probability density function of \mathbf{X} and suppose that $U = u(\mathbf{X})$ is a statistic taking values in R . Then U is sufficient for θ if and only if there exists $G : R \times T \rightarrow [0, \infty)$ and $r : S \rightarrow [0, \infty)$ such that

$$f_\theta(\mathbf{x}) = G[u(\mathbf{x}), \theta]r(\mathbf{x}); \quad \mathbf{x} \in S, \theta \in T \quad (7.6.3)$$

Proof

Let h_θ denote the PDF of U for $\theta \in T$. If U is sufficient for θ , then from the previous theorem, the function $r(\mathbf{x}) = f_\theta(\mathbf{x})/h_\theta[u(\mathbf{x})]$ for $\mathbf{x} \in S$ does not depend on $\theta \in T$. Hence $f_\theta(\mathbf{x}) = h_\theta[u(\mathbf{x})]r(\mathbf{x})$ for $(\mathbf{x}, \theta) \in S \times T$ and so $(\mathbf{x}, \theta) \mapsto f_\theta(\mathbf{x})$ has the form given in the theorem. Conversely, suppose that $(\mathbf{x}, \theta) \mapsto f_\theta(\mathbf{x})$ has the form given in the theorem. Then there exists a positive constant C such that $h_\theta(y) = CG(y, \theta)$ for $\theta \in T$ and $y \in R$. Hence $f_\theta(\mathbf{x})/h_\theta[u(\mathbf{x})] = r(\mathbf{x})/C$ for $\mathbf{x} \in S$, independent of $\theta \in T$.

Note that r depends only on the data \mathbf{x} but not on the parameter θ . Less technically, $u(\mathbf{X})$ is sufficient for θ if the probability density function $f_\theta(\mathbf{x})$ depends on the data vector \mathbf{x} and the parameter θ only through $u(\mathbf{x})$.

If U and V are equivalent statistics and U is sufficient for θ then V is sufficient for θ .

Minimal Sufficient Statistics

The entire data variable \mathbf{X} is trivially sufficient for θ . However, as noted above, there usually exists a statistic U that is sufficient for θ and has smaller dimension, so that we can achieve real data reduction. Naturally, we would like to find the statistic U that has the smallest dimension possible. In many cases, this smallest dimension j will be the same as the dimension k of the parameter vector θ . However, as we will see, this is not necessarily the case; j can be smaller or larger than k . An example based on the uniform distribution is given in (38).

Suppose that a statistic U is sufficient for θ . Then U is *minimally sufficient* if U is a function of any other statistic V that is sufficient for θ .

Once again, the definition precisely captures the notion of minimal sufficiency, but is hard to apply. The following result gives an equivalent condition.

Let f_θ denote the probability density function of \mathbf{X} corresponding to the parameter value $\theta \in T$ and suppose that $U = u(\mathbf{X})$ is a statistic taking values in R . Then U is minimally sufficient for θ if the following condition holds: for $\mathbf{x} \in S$ and $\mathbf{y} \in S$

$$\frac{f_\theta(\mathbf{x})}{f_\theta(\mathbf{y})} \text{ is independent of } \theta \text{ if and only if } u(\mathbf{x}) = u(\mathbf{y}) \quad (7.6.4)$$

Proof

Suppose that the condition in the theorem is satisfied. Then the PDF f_θ of \mathbf{X} must have the form given in the factorization theorem (3) so U is sufficient for θ . Next, suppose that $V = v(\mathbf{X})$ is another sufficient statistic for θ , taking values in R . From the factorization theorem, there exists $G : R \times T \rightarrow [0, \infty)$ and $r : S \rightarrow [0, \infty)$ such that $f_\theta(\mathbf{x}) = G[v(\mathbf{x}), \theta]r(\mathbf{x})$ for $(\mathbf{x}, \theta) \in S \times T$. Hence if $\mathbf{x}, \mathbf{y} \in S$ and $v(\mathbf{x}) = v(\mathbf{y})$ then

$$\frac{f_\theta(\mathbf{x})}{f_\theta(\mathbf{y})} = \frac{G[v(\mathbf{x}), \theta]r(\mathbf{x})}{G[v(\mathbf{y}), \theta]r(\mathbf{y})} = \frac{r(\mathbf{x})}{r(\mathbf{y})} \quad (7.6.5)$$

does not depend on $\theta \in \Theta$. Hence from the condition in the theorem, $u(\mathbf{x}) = u(\mathbf{y})$ and it follows that U is a function of V .

If U and V are equivalent statistics and U is minimally sufficient for θ then V is minimally sufficient for θ .

Properties of Sufficient Statistics

Sufficiency is related to several of the methods of constructing estimators that we have studied.

Suppose that U is sufficient for θ and that there exists a maximum likelihood estimator of θ . Then there exists a maximum likelihood estimator V that is a function of U .

Proof

From the factorization theorem (3), the log likelihood function for $\mathbf{x} \in S$ is

$$\theta \mapsto \ln G[u(\mathbf{x}), \theta] + \ln r(\mathbf{x}) \quad (7.6.6)$$

Hence a value of θ that maximizes this function, if it exists, must be a function of $u(\mathbf{x})$.

In particular, suppose that V is the unique maximum likelihood estimator of θ and that V is sufficient for θ . If U is sufficient for θ then V is a function of U by the previous theorem. Hence it follows that V is minimally sufficient for θ . Our next result applies to Bayesian analysis.

Suppose that the statistic $U = u(\mathbf{X})$ is sufficient for the parameter θ and that θ is modeled by a random variable Θ with values in T . Then the posterior distribution of Θ given $\mathbf{X} = \mathbf{x} \in S$ is a function of $u(\mathbf{x})$.

Proof

Let h denote the prior PDF of Θ and $f(\cdot | \theta)$ the conditional PDF of \mathbf{X} given $\Theta = \theta \in T$. By the factorization theorem (3), this conditional PDF has the form $f(\mathbf{x} | \theta) = G[u(\mathbf{x}), \theta]r(\mathbf{x})$ for $\mathbf{x} \in S$ and $\theta \in T$. The posterior PDF of Θ given $\mathbf{X} = \mathbf{x} \in S$ is

$$h(\theta | \mathbf{x}) = \frac{h(\theta)f(\mathbf{x} | \theta)}{f(\mathbf{x})}, \quad \theta \in T \quad (7.6.7)$$

where the function in the denominator is the marginal PDF of \mathbf{X} , or simply the normalizing constant for the function of θ in the numerator. Let's suppose that Θ has a continuous distribution on T , so that $f(\mathbf{x}) = \int_T h(t)G[u(\mathbf{x}), t]r(\mathbf{x})dt$ for $\mathbf{x} \in S$. Then the posterior PDF simplifies to

$$h(\theta | \mathbf{x}) = \frac{h(\theta)G[u(\mathbf{x}), \theta]}{\int_T h(t)G[u(\mathbf{x}), t]dt} \quad (7.6.8)$$

which depends on $\mathbf{x} \in S$ only through $u(\mathbf{x})$.

Continuing with the setting of Bayesian analysis, suppose that θ is a real-valued parameter. If we use the usual mean-square loss function, then the Bayesian estimator is $V = \mathbb{E}(\Theta | \mathbf{X})$. By the previous result, V is a function of the sufficient statistics U . That is, $\mathbb{E}(\Theta | \mathbf{X}) = \mathbb{E}(\Theta | U)$.

The next result is the *Rao-Blackwell theorem*, named for CR Rao and David Blackwell. The theorem shows how a sufficient statistic can be used to improve an unbiased estimator.

Rao-Blackwell Theorem. Suppose that U is sufficient for θ and that V is an unbiased estimator of a real parameter $\lambda = \lambda(\theta)$. Then $\mathbb{E}_\theta(V | U)$ is also an unbiased estimator of λ and is uniformly better than V .

Proof

This follows from basic properties of conditional expected value and conditional variance. First, since V is a function of \mathbf{X} and U is sufficient for θ , $\mathbb{E}_\theta(V | U)$ is a valid statistic; that is, it does not depend on θ , in spite of the formal dependence on θ in the expected value. Next, $\mathbb{E}_\theta(V | U)$ is a

function of U and $\mathbb{E}_\theta[\mathbb{E}_\theta(V | U)] = \mathbb{E}_\theta(V) = \lambda$ for $\theta \in \Theta$. Thus $\mathbb{E}_\theta(V | U)$ is an unbiased estimator of λ . Finally $\text{var}_\theta[\mathbb{E}_\theta(V | U)] = \text{var}_\theta(V) - \mathbb{E}_\theta[\text{var}_\theta(V | U)] \leq \text{var}_\theta(V)$ for any $\theta \in T$.

Complete Statistics

Suppose that $U = u(\mathbf{X})$ is a statistic taking values in a set R . Then U is a *complete* statistic for θ if for any function $r : R \rightarrow \mathbb{R}$

$$\mathbb{E}_\theta[r(U)] = 0 \text{ for all } \theta \in T \implies \mathbb{P}_\theta[r(U) = 0] = 1 \text{ for all } \theta \in T \quad (7.6.9)$$

To understand this rather strange looking condition, suppose that $r(U)$ is a statistic constructed from U that is being used as an estimator of 0 (thought of as a function of θ). The completeness condition means that the only such unbiased estimator is the statistic that is 0 with probability 1.

If U and V are equivalent statistics and U is complete for θ then V is complete for θ .

The next result shows the importance of statistics that are both complete and sufficient; it is known as the *Lehmann-Scheffé theorem*, named for Erich Lehmann and Henry Scheffé.

Lehmann-Scheffé Theorem. Suppose that U is sufficient and complete for θ and that $V = r(U)$ is an unbiased estimator of a real parameter $\lambda = \lambda(\theta)$. Then V is a uniformly minimum variance unbiased estimator (UMVUE) of λ .

Proof

Suppose that W is an unbiased estimator of λ . By the Rao-Blackwell theorem (10), $\mathbb{E}(W | U)$ is also an unbiased estimator of λ and is uniformly better than W . Since $\mathbb{E}(W | U)$ is a function of U , it follows from completeness that $V = \mathbb{E}(W | U)$ with probability 1.

Ancillary Statistics

Suppose that $V = v(\mathbf{X})$ is a statistic taking values in a set R . If the distribution of V does not depend on θ , then V is called an *ancillary* statistic for θ .

Thus, the notion of an ancillary statistic is complementary to the notion of a sufficient statistic. A sufficient statistic contains all available information about the parameter; an ancillary statistic contains no information about the parameter. The following result, known as *Basu's Theorem* and named for Debabrata Basu, makes this point more precisely.

Basu's Theorem. Suppose that U is complete and sufficient for a parameter θ and that V is an ancillary statistic for θ . Then U and V are independent.

Proof

Let g denote the probability density function of V and let $v \mapsto g(v | U)$ denote the conditional probability density function of V given U . From properties of conditional expected value, $\mathbb{E}[g(v | U)] = g(v)$ for $v \in R$. But then from completeness, $g(v | U) = g(v)$ with probability 1.

If U and V are equivalent statistics and U is ancillary for θ then V is ancillary for θ .

Applications and Special Distributions

In this subsection, we will explore sufficient, complete, and ancillary statistics for a number of special distributions. As always, be sure to try the problems yourself before looking at the solutions.

The Bernoulli Distribution

Recall that the *Bernoulli distribution* with parameter $p \in (0, 1)$ is a discrete distribution on $\{0, 1\}$ with probability density function g defined by

$$g(x) = p^x(1-p)^{1-x}, \quad x \in \{0, 1\} \quad (7.6.10)$$

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Bernoulli distribution with parameter p . Equivalently, \mathbf{X} is a sequence of *Bernoulli trials*, so that in the usual language of reliability, $X_i = 1$ if trial i is a success, and $X_i = 0$ if trial i is a failure. The Bernoulli distribution is named for Jacob Bernoulli and is studied in more detail in the chapter on Bernoulli Trials

Let $Y = \sum_{i=1}^n X_i$ denote the number of successes. Recall that Y has the binomial distribution with parameters n and p , and has probability density function h defined by

$$h(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (7.6.11)$$

Y is sufficient for p . Specifically, for $y \in \{0, 1, \dots, n\}$, the conditional distribution of \mathbf{X} given $Y = y$ is uniform on the set of points

$$D_y = \{(x_1, x_2, \dots, x_n) \in \{0, 1\}^n : x_1 + x_2 + \dots + x_n = y\} \quad (7.6.12)$$

Proof

The joint PDF f of \mathbf{X} is defined by

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = p^y(1-p)^{n-y}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \quad (7.6.13)$$

where $y = \sum_{i=1}^n x_i$. Now let $y \in \{0, 1, \dots, n\}$. Given $Y = y$, \mathbf{X} is concentrated on D_y and

$$\mathbb{P}(\mathbf{X} = \mathbf{x} \mid Y = y) = \frac{\mathbb{P}(\mathbf{X} = \mathbf{x})}{\mathbb{P}(Y = y)} = \frac{p^y(1-p)^{n-y}}{\binom{n}{y} p^y(1-p)^{n-y}} = \frac{1}{\binom{n}{y}}, \quad \mathbf{x} \in D_y \quad (7.6.14)$$

Of course, $\binom{n}{y}$ is the cardinality of D_y .

This result is intuitively appealing: in a sequence of Bernoulli trials, all of the information about the probability of success p is contained in the number of successes Y . The particular *order* of the successes and failures provides no additional information. Of course, the sufficiency of Y follows more easily from the factorization theorem (3), but the conditional distribution provides additional insight.

Y is complete for p on the parameter space $(0, 1)$.

Proof

If $r : \{0, 1, \dots, n\} \rightarrow \mathbb{R}$, then

$$\mathbb{E}[r(Y)] = \sum_{y=0}^n r(y) \binom{n}{y} p^y(1-p)^{n-y} = (1-p)^n \sum_{y=0}^n r(y) \binom{n}{y} \left(\frac{p}{1-p}\right)^y \quad (7.6.15)$$

The last sum is a polynomial in the variable $t = \frac{p}{1-p} \in (0, \infty)$. If this polynomial is 0 for all $t \in (0, \infty)$, then all of the coefficients must be 0. Hence we must have $r(y) = 0$ for $y \in \{0, 1, \dots, n\}$.

The proof of the last result actually shows that if the parameter space is any subset of $(0, 1)$ containing an interval of positive length, then Y is complete for p . But the notion of completeness depends very much on the parameter space. The following result considers the case where p has a finite set of values.

Suppose that the parameter space $T \subset (0, 1)$ is a finite set with $k \in \mathbb{N}_+$ elements. If the sample size n is at least k , then Y is not complete for p .

Proof

Suppose that $r : \{0, 1, \dots, n\} \rightarrow \mathbb{R}$ and that $\mathbb{E}[r(Y)] = 0$ for $p \in T$. Then we have

$$\sum_{y=0}^n \binom{n}{y} p^y(1-p)^{n-y} r(y) = 0, \quad p \in T \quad (7.6.16)$$

This is a set of k linear, homogenous equations in the variables $(r(0), r(1), \dots, r(n))$. Since $n \geq k$, we have at least $k+1$ variables, so there are infinitely many nontrivial solutions.

The sample mean $M = Y/n$ (the sample proportion of successes) is clearly equivalent to Y (the number of successes), and hence is also sufficient for p and is complete for $p \in (0, 1)$. Recall that the sample mean M is the method of moments estimator of p , and is the maximum likelihood estimator of p on the parameter space $(0, 1)$.

In Bayesian analysis, the usual approach is to model p with a random variable P that has a prior beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. Then the posterior distribution of P given \mathbf{X} is beta with left parameter $a + Y$ and right parameter $b + (n - Y)$. The posterior distribution depends on the data only through the sufficient statistic Y , as guaranteed by theorem (9).

The sample variance S^2 is an UMVUE of the distribution variance $p(1-p)$ for $p \in (0, 1)$, and can be written as

$$S^2 = \frac{Y}{n-1} \left(1 - \frac{Y}{n}\right) \quad (7.6.17)$$

Proof

Recall that the sample variance can be written as

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n X_i^2 - \frac{n}{n-1} M^2 \quad (7.6.18)$$

But $X_i^2 = X_i$ since X_i is an indicator variable, and $M = Y/n$. Substituting gives the representation above. In general, S^2 is an unbiased estimator of the distribution variance σ^2 . But in this case, S^2 is a function of the complete, sufficient statistic Y , and hence by the Lehmann-Scheffé theorem (13), S^2 is an UMVUE of $\sigma^2 = p(1-p)$.

The Poisson Distribution

Recall that the *Poisson distribution* with parameter $\theta \in (0, \infty)$ is a discrete distribution on \mathbb{N} with probability density function g defined by

$$g(x) = e^{-\theta} \frac{\theta^x}{x!}, \quad x \in \mathbb{N} \quad (7.6.19)$$

The Poisson distribution is named for Simeon Poisson and is used to model the number of “random points” in region of time or space, under certain ideal conditions. The parameter θ is proportional to the size of the region, and is both the mean and the variance of the distribution. The Poisson distribution is studied in more detail in the chapter on Poisson process.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Poisson distribution with parameter θ . Recall that the sum of the scores $Y = \sum_{i=1}^n X_i$ also has the Poisson distribution, but with parameter $n\theta$.

The statistic Y is sufficient for θ . Specifically, for $y \in \mathbb{N}$, the conditional distribution of \mathbf{X} given $Y = y$ is the multinomial distribution with y trials, n trial values, and uniform trial probabilities.

Proof

The joint PDF f of \mathbf{X} is defined by

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = \frac{e^{-n\theta}\theta^y}{x_1!x_2! \cdots x_n!}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}^n \quad (7.6.20)$$

where $y = \sum_{i=1}^n x_i$. Given $Y = y \in \mathbb{N}$, random vector \mathbf{X} takes values in the set $D_y = \{\mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{N}^n : \sum_{i=1}^n x_i = y\}$. Moreover,

$$\mathbb{P}(\mathbf{X} = \mathbf{x} \mid Y = y) = \frac{\mathbb{P}(\mathbf{X} = \mathbf{x})}{\mathbb{P}(Y = y)} = \frac{e^{-n\theta}\theta^y / (x_1!x_2! \cdots x_n!)}{e^{-n\theta}(n\theta)^y / y!} = \frac{y!}{x_1!x_2! \cdots x_n!} \frac{1}{n^y}, \quad \mathbf{x} \in D_y \quad (7.6.21)$$

The last expression is the PDF of the multinomial distribution stated in the theorem. Of course, the important point is that the conditional distribution does not depend on θ .

As before, it's easier to use the factorization theorem to prove the sufficiency of Y , but the conditional distribution gives some additional insight.

Y is complete for $\theta \in (0, \infty)$.

Proof

If $r : \mathbb{N} \rightarrow \mathbb{R}$ then

$$\mathbb{E}[r(Y)] = \sum_{y=0}^{\infty} e^{-n\theta} \frac{(n\theta)^y}{y!} r(y) = e^{-n\theta} \sum_{y=0}^{\infty} \frac{n^y}{y!} r(y) \theta^y \quad (7.6.22)$$

The last sum is a power series in θ with coefficients $n^y r(y) / y!$ for $y \in \mathbb{N}$. If this series is 0 for all θ in an open interval, then the coefficients must be 0 and hence $r(y) = 0$ for $y \in \mathbb{N}$.

As with our discussion of Bernoulli trials, the sample mean $M = Y/n$ is clearly equivalent to Y and hence is also sufficient for θ and complete for $\theta \in (0, \infty)$. Recall that M is the method of moments estimator of θ and is the maximum likelihood estimator on the parameter space $(0, \infty)$.

An UMVUE of the parameter $\mathbb{P}(X = 0) = e^{-\theta}$ for $\theta \in (0, \infty)$ is

$$U = \left(\frac{n-1}{n} \right)^Y \quad (7.6.23)$$

Proof

The probability generating function of Y is

$$P(t) = \mathbb{E}(t^Y) = e^{n\theta(t-1)}, \quad t \in \mathbb{R} \quad (7.6.24)$$

Hence

$$\mathbb{E} \left[\left(\frac{n-1}{n} \right)^Y \right] = \exp \left[n\theta \left(\frac{n-1}{n} - 1 \right) \right] = e^{-\theta}, \quad \theta \in (0, \infty) \quad (7.6.25)$$

So $U = [(n-1)/n]^Y$ is an unbiased estimator of $e^{-\theta}$. Since U is a function of the complete, sufficient statistic Y , it follows from the Lehmann-Scheffé theorem (13) that U is an UMVUE of $e^{-\theta}$.

The Normal Distribution

Recall that the *normal distribution* with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$ is a continuous distribution on \mathbb{R} with probability density function g defined by

$$g(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\frac{1}{2} \left(\frac{x-\mu}{\sigma} \right)^2 \right], \quad x \in \mathbb{R} \quad (7.6.26)$$

The normal distribution is often used to model physical quantities subject to small, random errors, and is studied in more detail in the chapter on Special Distributions. Because of the central limit theorem, the normal distribution is perhaps the most important distribution in statistics.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with mean μ and variance σ^2 . Then each of the following pairs of statistics is minimally sufficient for (μ, σ^2)

1. (Y, V) where $Y = \sum_{i=1}^n X_i$ and $V = \sum_{i=1}^n X_i^2$.
2. (M, S^2) where $M = \frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean and $S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M)^2$ is the sample variance.
3. (M, T^2) where $T^2 = \frac{1}{n} \sum_{i=1}^n (X_i - M)^2$ is the biased sample variance.

Proof

1. The joint PDF f of \mathbf{X} is given by

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = \frac{1}{(2\pi)^{n/2}\sigma^n} \exp\left[-\frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2\right], \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (7.6.27)$$

After some algebra, this can be written as

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{n/2}\sigma^n} e^{-n\mu^2/\sigma^2} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n x_i^2 + \frac{2\mu}{\sigma^2} \sum_{i=1}^n x_i\right), \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (7.6.28)$$

It follows from the factorization theorem (3) that (Y, V) is sufficient for (μ, σ^2) . Minimal sufficiency follows from the condition in theorem (6).

2. Note that $M = \frac{1}{n}Y$, $S^2 = \frac{1}{n-1}V - \frac{n}{n-1}M^2$. Hence (M, S^2) is equivalent to (Y, V) and so (M, S^2) is also minimally sufficient for (μ, σ^2) .
3. Similarly, $M = \frac{1}{n}Y$ and $T^2 = \frac{1}{n}V - M^2$. Hence (M, T^2) is equivalent to (Y, V) and so (M, T^2) is also minimally sufficient for (μ, σ^2) .

Recall that M and T^2 are the method of moments estimators of μ and σ^2 , respectively, and are also the maximum likelihood estimators on the parameter space $\mathbb{R} \times (0, \infty)$.

Run the normal estimation experiment 1000 times with various values of the parameters. Compare the estimates of the parameters in terms of bias and mean square error.

Sometimes the variance σ^2 of the normal distribution is known, but not the mean μ . It's rarely the case that μ is known but not σ^2 . Nonetheless we can give sufficient statistics in both cases.

Suppose again that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in (0, \infty)$. If

1. If σ^2 is known then $Y = \sum_{i=1}^n X_i$ is minimally sufficient for μ .
2. If μ is known then $U = \sum_{i=1}^n (X_i - \mu)^2$ is sufficient for σ^2 .

Proof

1. This results follow from the second displayed equation for the PDF $f(\mathbf{x})$ of \mathbf{X} in the proof of the previous theorem.
2. This result follows from the first displayed equation for the PDF $f(\mathbf{x})$ of \mathbf{X} in the proof of the previous theorem.

Of course by equivalence, in part (a) the sample mean $M = Y/n$ is minimally sufficient for μ , and in part (b) the special sample variance $W = U/n$ is minimally sufficient for σ^2 . Moreover, in part (a), M is complete for μ on the parameter space \mathbb{R} and the sample variance S^2 is ancillary for μ (Recall that $(n-1)S^2/\sigma^2$ has the chi-square distribution with $n-1$ degrees of freedom.) It follows from Basu's theorem (15) that the sample mean M and the sample variance S^2 are independent. We proved this by more direct means in the section on special properties of normal samples, but the formulation in terms of sufficient and ancillary statistics gives additional insight.

The Gamma Distribution

Recall that the *gamma distribution* with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ is a continuous distribution on $(0, \infty)$ with probability density function g given by

$$g(x) = \frac{1}{\Gamma(k)b^k} x^{k-1} e^{-x/b}, \quad x \in (0, \infty) \quad (7.6.29)$$

The gamma distribution is often used to model random times and certain other types of positive random variables, and is studied in more detail in the chapter on Special Distributions.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the gamma distribution with shape parameter k and scale parameter b . Each of the following pairs of statistics is minimally sufficient for (k, b)

1. (Y, V) where $Y = \sum_{i=1}^n X_i$ is the sum of the scores and $V = \prod_{i=1}^n X_i$ is the product of the scores.
2. (M, U) where $M = Y/n$ is the sample (arithmetic) mean of \mathbf{X} and $U = V^{1/n}$ is the sample geometric mean of \mathbf{X} .

Proof

1. The joint PDF f of \mathbf{X} is given by

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = \frac{1}{\Gamma^n(k)b^{nk}} (x_1 x_2 \cdots x_n)^{k-1} e^{-(x_1 + x_2 + \cdots + x_n)/b}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in (0, \infty)^n \quad (7.6.30)$$

From the factorization theorem (3), (Y, V) is sufficient for (k, b) . Minimal sufficiency follows from condition (6).

2. Clearly $M = Y/n$ is equivalent to Y and $U = V^{1/n}$ is equivalent to V . Hence (M, U) is also minimally sufficient for (k, b) .

Recall that the method of moments estimators of k and b are M^2/T^2 and T^2/M , respectively, where $M = \frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean and $T^2 = \frac{1}{n} \sum_{i=1}^n (X_i - M)^2$ is the biased sample variance. If the shape parameter k is known, $\frac{1}{k}M$ is both the method of moments estimator of b and the maximum likelihood estimator on the parameter space $(0, \infty)$. Note that T^2 is not a function of the sufficient statistics (Y, V) , and hence estimators based on T^2 suffer from a loss of information.

Run the gamma estimation experiment 1000 times with various values of the parameters and the sample size n . Compare the estimates of the parameters in terms of bias and mean square error.

The proof of the last theorem actually shows that Y is sufficient for b if k is known, and that V is sufficient for k if b is known.

Suppose again that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the gamma distribution with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. Then $Y = \sum_{i=1}^n X_i$ is complete for b .

Proof

Y has the gamma distribution with shape parameter nk and scale parameter b . Hence, if $r : [0, \infty) \rightarrow \mathbb{R}$, then

$$\mathbb{E}[r(Y)] = \int_0^\infty \frac{1}{\Gamma(nk)b^{nk}} y^{nk-1} e^{-y/b} r(y) dy = \frac{1}{\Gamma(nk)b^{nk}} \int_0^\infty y^{nk-1} r(y) e^{-y/b} dy \quad (7.6.31)$$

The last integral can be interpreted as the Laplace transform of the function $y \mapsto y^{nk-1} r(y)$ evaluated at $1/b$. If this transform is 0 for all b in an open interval, then $r(y) = 0$ almost everywhere in $(0, \infty)$.

Suppose again that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the gamma distribution on $(0, \infty)$ with shape parameter $k \in (0, \infty)$ and scale parameter $b \in (0, \infty)$. Let $M = \frac{1}{n} \sum_{i=1}^n X_i$ denote the sample mean and $U = (X_1 X_2 \dots X_n)^{1/n}$ the sample geometric mean, as before. Then

1. M/U is ancillary for b .
2. M and M/U are independent.

Proof

1. We can take $X_i = bZ_i$ for $i \in \{1, 2, \dots, n\}$ where $\mathbf{Z} = (Z_1, Z_2, \dots, Z_n)$ is a random sample of size n from the gamma distribution with shape parameter k and scale parameter 1 (the standard gamma distribution with shape parameter k). Then

$$\frac{M}{U} = \frac{1}{n} \sum_{i=1}^n \frac{X_i}{(X_1 X_2 \dots X_n)^{1/n}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{X_i^n}{X_1 X_2 \dots X_n} \right)^{1/n} = \frac{1}{n} \sum_{i=1}^n \left(\prod_{j \neq i} \frac{X_i}{X_j} \right)^{1/n} \quad (7.6.32)$$

But $X_i/X_j = Z_i/Z_j$ for $i \neq j$, and the distribution of $\{Z_i/Z_j : i, j \in \{1, 2, \dots, n\}, i \neq j\}$ does not depend on b . Hence the distribution of M/U does not depend on b .

2. This follows from Basu's theorem (15), since M is complete and sufficient for b and M/U is ancillary for b .

The Beta Distribution

Recall that the *beta distribution* with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$ is a continuous distribution on $(0, 1)$ with probability density function g given by

$$g(x) = \frac{1}{B(a, b)} x^{a-1} (1-x)^{b-1}, \quad x \in (0, 1) \quad (7.6.33)$$

where B is the beta function. The beta distribution is often used to model random proportions and other random variables that take values in bounded intervals. It is studied in more detail in the chapter on Special Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the beta distribution with left parameter a and right parameter b . Then (P, Q) is minimally sufficient for (a, b) where $P = \prod_{i=1}^n X_i$ and $Q = \prod_{i=1}^n (1 - X_i)$.

Proof

The joint PDF f of \mathbf{X} is given by

$$f(\mathbf{x}) = g(x_1)g(x_2) \dots g(x_n) = \frac{1}{B^n(a, b)} (x_1 x_2 \dots x_n)^{a-1} [(1-x_1)(1-x_2) \dots (1-x_n)]^{b-1}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in (0, 1)^n \quad (7.6.34)$$

From the factorization theorem (3), it follows that (U, V) is sufficient for (a, b) . Minimal sufficiency follows from condition (6).

The proof also shows that P is sufficient for a if b is known, and that Q is sufficient for b if a is known. Recall that the method of moments estimators of a and b are

$$U = \frac{M(M - M^{(2)})}{M^{(2)} - M^2}, \quad V = \frac{(1 - M)(M - M^{(2)})}{M^{(2)} - M^2} \quad (7.6.35)$$

respectively, where $M = \frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean and $M^{(2)} = \frac{1}{n} \sum_{i=1}^n X_i^2$ is the second order sample mean. If b is known, the method of moments estimator of a is $U_b = bM/(1-M)$, while if a is known, the method of moments estimator of b is $V_a = a(1-M)/M$. None of these estimators is a function of the sufficient statistics (P, Q) and so all suffer from a loss of information. On the other hand, if $b = 1$, the maximum likelihood estimator of a on the interval $(0, \infty)$ is $W = -n / \sum_{i=1}^n \ln X_i$, which is a function of P (as it must be).

Run the beta estimation experiment 1000 times with various values of the parameters. Compare the estimates of the parameters.

The Pareto Distribution

Recall that the *Pareto distribution* with shape parameter $a \in (0, \infty)$ and scale parameter $b \in (0, \infty)$ is a continuous distribution on $[b, \infty)$ with probability density function g given by

$$g(x) = \frac{ab^a}{x^{a+1}}, \quad b \leq x < \infty \quad (7.6.36)$$

The Pareto distribution, named for Vilfredo Pareto, is a heavy-tailed distribution often used to model income and certain other types of random variables. It is studied in more detail in the chapter on Special Distribution.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Pareto distribution with shape parameter a and scale parameter b . Then $(P, X_{(1)})$ is minimally sufficient for (a, b) where $P = \prod_{i=1}^n X_i$ is the product of the sample variables and where $X_{(1)} = \min\{X_1, X_2, \dots, X_n\}$ is the first order statistic.

Proof

The joint PDF f of \mathbf{X} at $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is given by

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = \frac{a^n b^{na}}{(x_1 x_2 \cdots x_n)^{a+1}}, \quad x_1 \geq b, x_2 \geq b, \dots, x_n \geq b \quad (7.6.37)$$

which can be rewritten as

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = \frac{a^n b^{na}}{(x_1 x_2 \cdots x_n)^{a+1}} \mathbf{1}(x_{(1)} \geq b), \quad (x_1, x_2, \dots, x_n) \in (0, \infty)^n \quad (7.6.38)$$

So the result follows from the factorization theorem (3). Minimal sufficiency follows from condition (6).

The proof also shows that P is sufficient for a if b is known (which is often the case), and that $X_{(1)}$ is sufficient for b if a is known (much less likely). Recall that the method of moments estimators of a and b are

$$U = 1 + \sqrt{\frac{M^{(2)}}{M^{(2)} - M^2}}, \quad V = \frac{M^{(2)}}{M} \left(1 - \sqrt{\frac{M^{(2)} - M^2}{M^{(2)}}} \right) \quad (7.6.39)$$

respectively, where as before $M = \frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean and $M^{(2)} = \frac{1}{n} \sum_{i=1}^n X_i^2$ the second order sample mean. These estimators are not functions of the sufficient statistics and hence suffers from loss of information. On the other hand, the maximum likelihood estimators of a and b on the interval $(0, \infty)$ are

$$W = \frac{n}{\sum_{i=1}^n \ln X_i - n \ln X_{(1)}}, \quad X_{(1)} \quad (7.6.40)$$

respectively. These are functions of the sufficient statistics, as they must be.

Run the Pareto estimation experiment 1000 times with various values of the parameters a and b and the sample size n . Compare the method of moments estimates of the parameters with the maximum likelihood estimates in terms of the empirical bias and mean square error.

The Uniform Distribution

Recall that the *continuous uniform distribution* on the interval $[a, a+h]$, where $a \in \mathbb{R}$ is the location parameter and $h \in (0, \infty)$ is the scale parameter, has probability density function g given by

$$g(x) = \frac{1}{h}, \quad x \in [a, a+h] \quad (7.6.41)$$

Continuous uniform distributions are widely used in applications to model a number chosen “at random” from an interval. Continuous uniform distributions are studied in more detail in the chapter on Special Distributions. Let's first consider the case where both parameters are unknown.

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the uniform distribution on the interval $[a, a+h]$. Then $(X_{(1)}, X_{(n)})$ is minimally sufficient for (a, h) , where $X_{(1)} = \min\{X_1, X_2, \dots, X_n\}$ is the first order statistic and $X_{(n)} = \max\{X_1, X_2, \dots, X_n\}$ is the last order statistic.

Proof

The PDF f of \mathbf{X} is given by

$$f(\mathbf{x}) = g(x_1)g(x_2) \cdots g(x_n) = \frac{1}{h^n}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in [a, a+h]^n \quad (7.6.42)$$

We can rewrite the PDF as

$$f(\mathbf{x}) = \frac{1}{h^n} \mathbf{1}_{[x_{(1)} \geq a]} \mathbf{1}_{[x_{(n)} \leq a+h]}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (7.6.43)$$

It then follows from the factorization theorem (3) that $(X_{(1)}, X_{(n)})$ is sufficient for (a, h) . Next, suppose that $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and that $x_{(1)} \neq y_{(1)}$ or $x_{(n)} \neq y_{(n)}$. For a given $h \in (0, \infty)$, we can easily find values of $a \in \mathbb{R}$ such that $f(\mathbf{x}) = 0$ and $f(\mathbf{y}) = 1/h^n$, and other values of $a \in \mathbb{R}$ such that $f(\mathbf{x}) = f(\mathbf{y}) = 1/h^n$. By condition (6), $(X_{(1)}, X_{(n)})$ is minimally sufficient.

If the location parameter a is known, then the largest order statistic is sufficient for the scale parameter h . But if the scale parameter h is known, we still need both order statistics for the location parameter a . So in this case, we have a single real-valued parameter, but the minimally sufficient statistic is a pair of real-valued random variables.

Suppose again that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the uniform distribution on the interval $[a, a+h]$.

1. If $a \in \mathbb{R}$ is known, then $X_{(n)}$ is sufficient for h .
2. If $h \in (0, \infty)$ is known, then $(X_{(1)}, X_{(n)})$ is minimally sufficient for a .

Proof

Both parts follow easily from the analysis given in the proof of the last theorem.

Run the uniform estimation experiment 1000 times with various values of the parameter. Compare the estimates of the parameter.

Recall that if both parameters are unknown, the method of moments estimators of a and h are $U = 2M - \sqrt{3}T$ and $V = 2\sqrt{3}T$, respectively, where $M = \frac{1}{n} \sum_{i=1}^n X_i$ is the sample mean and $T^2 = \frac{1}{n} \sum_{i=1}^n (X_i - M)^2$ is the biased sample variance. If a is known, the method of moments estimator of h is $V_a = 2(M - a)$, while if h is known, the method of moments estimator of a is $U_h = M - \frac{1}{2}h$. None of these estimators are functions of the minimally sufficient statistics, and hence result in loss of information.

The Hypergeometric Model

So far, in all of our examples, the basic variables have formed a random sample from a distribution. In this subsection, our basic variables will be dependent.

Recall that in the *hypergeometric model*, we have a population of N objects, and that r of the objects are *type 1* and the remaining $N - r$ are *type 0*. The population size N is a positive integer and the type 1 size r is a nonnegative integer with $r \leq N$. Typically one or both parameters are unknown. We select a random sample of n objects, without replacement from the population, and let X_i be the type of the i th object chosen. So our basic sequence of random variables is $\mathbf{X} = (X_1, X_2, \dots, X_n)$. The variables are identically distributed indicator variables with $\mathbb{P}(X_i = 1) = r/N$ for $i \in \{1, 2, \dots, n\}$, but are dependent. Of course, the sample size n is a positive integer with $n \leq N$.

The variable $Y = \sum_{i=1}^n X_i$ is the number of type 1 objects in the sample. This variable has the *hypergeometric distribution* with parameters N, r , and n , and has probability density function h given by

$$h(y) = \frac{\binom{r}{y} \binom{N-r}{n-y}}{\binom{N}{n}} = \binom{n}{y} \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad y \in \{\max\{0, N-n+r\}, \dots, \min\{n, r\}\} \quad (7.6.44)$$

(Recall the *falling power* notation $x^{(k)} = x(x-1) \cdots (x-k+1)$). The hypergeometric distribution is studied in more detail in the chapter on Finite Sampling Models.

Y is sufficient for (N, r) . Specifically, for $y \in \{\max\{0, N-n+r\}, \dots, \min\{n, r\}\}$, the conditional distribution of \mathbf{X} given $Y = y$ is uniform on the set of points

$$D_y = \{(x_1, x_2, \dots, x_n) \in \{0, 1\}^n : x_1 + x_2 + \cdots + x_n = y\} \quad (7.6.45)$$

Proof

By a simple application of the multiplication rule of combinatorics, the PDF f of \mathbf{X} is given by

$$f(\mathbf{x}) = \frac{r^{(y)} (N-r)^{(n-y)}}{N^{(n)}}, \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \quad (7.6.46)$$

where $y = \sum_{i=1}^n x_i$. If $y \in \{\max\{0, N-n+r\}, \dots, \min\{n, r\}\}$, the conditional distribution of \mathbf{X} given $Y = y$ is concentrated on D_y and

$$\mathbb{P}(\mathbf{X} = \mathbf{x} \mid Y = y) = \frac{\mathbb{P}(\mathbf{X} = \mathbf{x})}{\mathbb{P}(Y = y)} = \frac{r^{(y)} (N-r)^{(n-y)} / N^{(n)}}{\binom{n}{y} r^{(y)} (N-r)^{(n-y)} / N^{(n)}} = \frac{1}{\binom{n}{y}}, \quad \mathbf{x} \in D_y \quad (7.6.47)$$

Of course, $\binom{n}{y}$ is the cardinality of D_y .

There are clearly strong similarities between the hypergeometric model and the Bernoulli trials model above. Indeed if the sampling were *with* replacement, the Bernoulli trials model with $p = r/N$ would apply rather than the hypergeometric model. It's also interesting to note that we have a single real-valued statistic that is sufficient for two real-valued parameters.

Once again, the sample mean $M = Y/n$ is equivalent to Y and hence is also sufficient for (N, r) . Recall that the method of moments estimator of r with N known is NM and the method of moment estimator of N with r known is r/M . The estimator of r is the one that is used in the capture-recapture experiment.

Exponential Families

Suppose now that our data vector \mathbf{X} takes values in a set S , and that the distribution of \mathbf{X} depends on a parameter vector $\boldsymbol{\theta}$ taking values in a parameter space Θ . The distribution of \mathbf{X} is a k -parameter *exponential family* if S does not depend on $\boldsymbol{\theta}$ and if the probability density function of \mathbf{X} can be written as

$$f_{\boldsymbol{\theta}}(\mathbf{x}) = \alpha(\boldsymbol{\theta})r(\mathbf{x}) \exp\left(\sum_{i=1}^k \beta_i(\boldsymbol{\theta})u_i(\mathbf{x})\right); \quad \mathbf{x} \in S, \boldsymbol{\theta} \in \Theta \quad (7.6.48)$$

where α and $(\beta_1, \beta_2, \dots, \beta_k)$ are real-valued functions on Θ , and where r and (u_1, u_2, \dots, u_k) are real-valued functions on S . Moreover, k is assumed to be the smallest such integer. The parameter vector $\boldsymbol{\beta} = (\beta_1(\boldsymbol{\theta}), \beta_2(\boldsymbol{\theta}), \dots, \beta_k(\boldsymbol{\theta}))$ is sometimes called the *natural parameter* of the distribution, and the random vector $\mathbf{U} = (u_1(\mathbf{X}), u_2(\mathbf{X}), \dots, u_k(\mathbf{X}))$ is sometimes called the *natural statistic* of the distribution. Although the definition may look intimidating, exponential families are useful because they have many nice mathematical properties, and because many special parametric families are exponential families. In particular, the sampling distributions from the [Bernoulli](#), [Poisson](#), [gamma](#), [normal](#), [beta](#), and [Pareto](#) considered above are exponential families. Exponential families of distributions are studied in more detail in the chapter on special distributions.

\mathbf{U} is minimally sufficient for $\boldsymbol{\theta}$.

Proof

That \mathbf{U} is sufficient for $\boldsymbol{\theta}$ follows immediately from the factorization theorem. That \mathbf{U} is minimally sufficient follows since k is the smallest integer in the exponential formulation.

It turns out that \mathbf{U} is complete for $\boldsymbol{\theta}$ as well, although the proof is more difficult.

This page titled [7.6: Sufficient, Complete and Ancillary Statistics](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

8: Set Estimation

Set estimation refers to the process of constructing a subset of the parameter space, based on observed data from a probability distribution. The subset will contain the true value of the parameter with a specified *confidence level*. In this chapter, we explore the basic method of set estimation using pivot variables. We study set estimation in some of the most important models: the single variable normal model, the two-variable normal model, and the Bernoulli model.

[8.1: Introduction to Set Estimation](#)

[8.2: Estimation the Normal Model](#)

[8.3: Estimation in the Bernoulli Model](#)

[8.4: Estimation in the Two-Sample Normal Model](#)

[8.5: Bayesian Set Estimation](#)

This page titled [8: Set Estimation](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.1: Introduction to Set Estimation

Basic Theory

The Basic Statistical Model

As usual, our starting point is a random experiment with an underlying sample space and a probability measure \mathbb{P} . In the basic statistical model, we have an observable random variable \mathbf{X} taking values in a set S . In general, \mathbf{X} can have quite a complicated structure. For example, if the experiment is to sample n objects from a population and record various measurements of interest, then

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (8.1.1)$$

where X_i is the vector of measurements for the i th object. The most important special case occurs when (X_1, X_2, \dots, X_n) are independent and identically distributed. In this case, we have a random sample of size n from the common distribution.

Suppose also that the distribution of \mathbf{X} depends on a parameter θ taking values in a parameter space Θ . The parameter may also be vector-valued, in which case $\Theta \subseteq \mathbb{R}^k$ for some $k \in \mathbb{N}_+$ and the parameter vector has the form $\theta = (\theta_1, \theta_2, \dots, \theta_k)$.

Confidence Sets

A *confidence set* is a subset $C(\mathbf{X})$ of the parameter space Θ that depends only on the data variable \mathbf{X} , and no unknown parameters. the *confidence level* is the smallest probability that $\theta \in C(\mathbf{X})$:

$$\min \{\mathbb{P}[\theta \in C(\mathbf{X})] : \theta \in \Theta\} \quad (8.1.2)$$

Thus, in a sense, a confidence set is a *set-valued statistic*. A confidence set is an estimator of θ in the sense that we hope that $\theta \in C(\mathbf{X})$ with high probability, so that the confidence level is high. Note that since the distribution of \mathbf{X} depends on θ , there is a dependence on θ in the probability measure \mathbb{P} in the definition of confidence level. However, we usually suppress this, just to keep the notation simple. Usually, we try to construct a confidence set for θ with a prescribed confidence level $1 - \alpha$ where $0 < \alpha < 1$. Typical confidence levels are 0.9, 0.95, and 0.99. Sometimes the best we can do is to construct a confidence set whose confidence level is *at least* $1 - \alpha$; this is called a *conservative* $1 - \alpha$ confidence set for θ .




Figure 8.1.1: A set estimate that successfully captured the parameter

Suppose that $C(\mathbf{X})$ is $1 - \alpha$ level confidence set for a parameter θ . Note that when we run the experiment and observe the data \mathbf{x} , the *computed* confidence set is $C(\mathbf{x})$. The true value of θ is either in this set, or is not, and we will usually never know. However, by the law of large numbers, if we were to repeat the confidence experiment over and over, the proportion of sets that contain θ would converge to $\mathbb{P}[\theta \in C(\mathbf{X})] = 1 - \alpha$. This is the precise meaning of the term *confidence*. In the usual terminology of statistics, the random set $C(\mathbf{X})$ is the *estimator*; the deterministic set $C(\mathbf{x})$ based on an observed value \mathbf{x} is the *estimate*.

Next, note that the quality of a confidence set, as an estimator of θ , is based on two factors: the confidence level and the *precision* as measured by the “size” of the set. A good estimator has small size (and hence gives a precise estimate of θ) and large confidence. However, for a given \mathbf{X} , there is usually a tradeoff between confidence level and precision—increasing the confidence level comes only at the expense of increasing the size of the set, and decreasing the size of the set comes only at the expense of decreasing the confidence level. How we measure the “size” of the confidence set depends on the dimension of the parameter space and the nature of the confidence set. Moreover, the size of the set is usually random, although in some special cases it may be deterministic.

Considering the extreme cases may give us some insight. First, suppose that $C(\mathbf{X}) = \Theta$. This set estimator has maximum confidence 1, but no precision and hence it is worthless (we already *knew* that $\theta \in \Theta$). At the other extreme, suppose that $C(\mathbf{X})$ is a singleton set. This set estimator has the best possible precision, but typically for continuous distributions, would have confidence 0. In between these extremes, hopefully, are set estimators that have high confidence and high precision.

Suppose that $C_i(\mathbf{X})$ is a $1 - \alpha_i$ level confidence set for θ for $i \in \{1, 2, \dots, k\}$. If $\alpha = \alpha_1 + \alpha_2 + \dots + \alpha_k < 1$ then $C_1(\mathbf{X}) \cap C_2(\mathbf{X}) \cap \dots \cap C_k(\mathbf{X})$ is a conservative $1 - \alpha$ level confidence set for θ .

Proof

This follows from Bonferroni's inequality.

Real-Valued Parameters

In many cases, we are interested in estimating a real-valued parameter $\lambda = \lambda(\theta)$ taking values in an interval parameter space (a, b) , where $a, b \in \mathbb{R}$ with $a < b$. Of course, it's possible that $a = -\infty$ or $b = \infty$. In this context our confidence set frequently has the form

$$C(\mathbf{X}) = \{\theta \in \Theta : L(\mathbf{X}) < \lambda(\theta) < U(\mathbf{X})\} \quad (8.1.3)$$

where $L(\mathbf{X})$ and $U(\mathbf{X})$ are real-valued statistics. In this case $(L(\mathbf{X}), U(\mathbf{X}))$ is called a *confidence interval* for λ . If $L(\mathbf{X})$ and $U(\mathbf{X})$ are both random, then the confidence interval is often said to be *two-sided*. In the special case that $U(\mathbf{X}) = b$, $L(\mathbf{X})$ is called a *confidence lower bound* for λ . In the special case that $L(\mathbf{X}) = a$, $U(\mathbf{X})$ is called a *confidence upper bound* for λ .

Suppose that $L(\mathbf{X})$ is a $1 - \alpha$ level confidence lower bound for λ and that $U(\mathbf{X})$ is a $1 - \beta$ level confidence upper bound for λ . If $\alpha + \beta < 1$ then $(L(\mathbf{X}), U(\mathbf{X}))$ is a conservative $1 - (\alpha + \beta)$ level confidence interval for λ .

Proof

This follows immediately from (2).

Pivot Variables

You might think that it should be very difficult to construct confidence sets for a parameter θ . However, in many important special cases, confidence sets can be constructed easily from certain random variables known as *pivot variables*.

Suppose that V is a function from $S \times \Theta$ into a set T . The random variable $V(\mathbf{X}, \theta)$ is a *pivot variable* for θ if its distribution does not depend on θ . Specifically, $\mathbb{P}[V(\mathbf{X}, \theta) \in B]$ is constant in $\theta \in \Theta$ for each $B \subseteq T$.

The basic idea is that we try to combine \mathbf{X} and θ algebraically in such a way that we *factor out* the dependence on θ in the distribution of the resulting random variable $V(\mathbf{X}, \theta)$. If we know the distribution of the pivot variable, then for a given α , we can try to find $B \subseteq T$ (that does not depend on θ) such that $\mathbb{P}_\theta[V(\mathbf{X}, \theta) \in B] = 1 - \alpha$. It then follows that a $1 - \alpha$ confidence set for the parameter is given by $C(\mathbf{X}) = \{\theta \in \Theta : V(\mathbf{X}, \theta) \in B\}$.




Figure 8.1.2: A confidence set constructed from a pivot variable

Suppose now that our pivot variable $V(\mathbf{X}, \theta)$ is real-valued, which for simplicity, we will assume has a continuous distribution. For $p \in (0, 1)$, let $v(p)$ denote the quantile of order p for the pivot variable $V(\mathbf{X}, \theta)$. By the very meaning of pivot variable, $v(p)$ does not depend on θ .

For any $p \in (0, 1)$, a $1 - \alpha$ level confidence set for θ is

$$\{\theta \in \Theta : v(\alpha - p\alpha) < V(\mathbf{X}, \theta) < v(1 - p\alpha)\} \quad (8.1.4)$$

Proof

By definition, the probability of the event is $(1 - p\alpha) - (\alpha - p\alpha) = 1 - \alpha$.

The confidence set above corresponds to $(1-p)\alpha$ in the left tail and $p\alpha$ in the right tail, in terms of the distribution of the pivot variable $V(\mathbf{X}, \lambda)$. The special case $p = \frac{1}{2}$ is the *equal-tailed case*, the most common case.




Figure 8.1.3: Distribution of the pivot variable showing $(1-p)\alpha$ is the left tail and $p\alpha$ is the right tail.

The confidence set (5) is decreasing in α and hence increasing in $1-\alpha$ (in the sense of the subset relation) for fixed p .

For the confidence set (5), we would naturally like to choose p that minimizes the size of the set in some sense. However this is often a difficult problem. The equal-tailed interval, corresponding to $p = \frac{1}{2}$, is the most commonly used case, and is sometimes (but not always) an optimal choice. Pivot variables are far from unique; the challenge is to find a pivot variable whose distribution is known and which gives tight bounds on the parameter (high precision).

Suppose that $V(\mathbf{X}, \theta)$ is a pivot variable for θ . If g is a function defined on the range of V and g involves no unknown parameters, then $U = g[V(\mathbf{X}, \theta)]$ is also a pivot variable for θ .

Examples and Special Cases

Location-Scale Families

In the case of location-scale families of distributions, we can easily find pivot variables. Suppose that Z is a real-valued random variable with a continuous distribution that has probability density function g , and no unknown parameters. Let $X = \mu + \sigma Z$ where $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$ are parameters. Recall that the probability density function of X is given by

$$f_{\mu, \sigma}(x) = \frac{1}{\sigma} g\left(\frac{x - \mu}{\sigma}\right), \quad x \in \mathbb{R} \quad (8.1.5)$$

and the corresponding family of distributions is called the *location-scale family* associated with the distribution of Z ; μ is the *location parameter* and σ is the *scale parameter*. Generally, we are assuming that these parameters are unknown.

Now suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the distribution of X ; this is our observable outcome vector. For each i , let

$$Z_i = \frac{X_i - \mu}{\sigma} \quad (8.1.6)$$

The random vector $\mathbf{Z} = (Z_1, Z_2, \dots, Z_n)$ is a random sample of size n from the distribution of Z .

In particular, note that \mathbf{Z} is a pivot variable for (μ, σ) , since \mathbf{Z} is a function of \mathbf{X} , μ , and σ , but the distribution of \mathbf{Z} does not depend on μ or σ . Hence, any function of \mathbf{Z} will also be a pivot variable for (μ, σ) , (if the function does not involve the parameters). Of course, some of these pivot variables will be much more useful than others in estimating μ and σ . In the following exercises, we will explore two common and important pivot variables.

Let $M(\mathbf{X})$ and $M(\mathbf{Z})$ denote the sample means of \mathbf{X} and \mathbf{Z} , respectively. Then $M(\mathbf{Z})$ is a pivot variable for (μ, σ) since

$$M(\mathbf{Z}) = \frac{M(\mathbf{X}) - \mu}{\sigma} \quad (8.1.7)$$

Let m denote the quantile function of the pivot variable $M(\mathbf{Z})$. For any $p \in (0, 1)$, a $1-\alpha$ confidence set for (μ, σ) is

$$Z_{\alpha, p}(\mathbf{X}) = \{(\mu, \sigma) : M(\mathbf{X}) - m(1-p\alpha)\sigma < \mu < M(\mathbf{X}) - m(\alpha-p\alpha)\sigma\} \quad (8.1.8)$$

The confidence set constructed above is a “cone” in the (μ, σ) parameter space, with vertex at $(M(\mathbf{X}), 0)$ and boundary lines of slopes $-1/m(1-p\alpha)$ and $-1/m(\alpha-p\alpha)$, as shown in the graph below. (Note, however, that both slopes might be negative or both positive.)




Figure 8.1.4: The confidence set for (μ, σ) constructed from M

The fact that the confidence set is unbounded is clearly not good, but is perhaps not surprising; we are estimating two real parameters with a single real-valued pivot variable. However, if σ is known, the confidence set defines a *confidence interval* for μ . Geometrically, the confidence interval simply corresponds to the horizontal cross section at σ .

$1 - \alpha$ confidence sets for (μ, σ) are

1. $Z_{\alpha,1}(\mathbf{X}) = \{(\mu, \sigma) : M(\mathbf{X}) - m(1 - \alpha)\sigma < \mu < \infty\}$
2. $Z_{\alpha,0}(\mathbf{X}) = \{(\mu, \sigma) : -\infty < \mu < M(\mathbf{X}) - m(\alpha)\sigma\}$

Proof

In the confidence set constructed above, let $p \uparrow 1$ and $p \downarrow 0$, respectively.

If σ is known, then (a) gives a $1 - \alpha$ confidence lower bound for μ and (b) gives a $1 - \alpha$ confidence upper bound for μ .

Let $S(\mathbf{X})$ and $S(\mathbf{Z})$ denote the sample standard deviations of \mathbf{X} and \mathbf{Z} , respectively. Then $S(\mathbf{Z})$ is a pivot variable for (μ, σ) and a pivot variable for σ since

$$S(\mathbf{Z}) = \frac{S(\mathbf{X})}{\sigma} \quad (8.1.9)$$

Let s denote the quantile function of $S(\mathbf{Z})$. For any $\alpha \in (0, 1)$ and $p \in (0, 1)$, a $1 - \alpha$ confidence set for (μ, σ) is

$$V_{\alpha,p}(\mathbf{X}) = \left\{ (\mu, \sigma) : \frac{S(\mathbf{X})}{s(1-p\alpha)} < \sigma < \frac{S(\mathbf{X})}{s(\alpha-p\alpha)} \right\} \quad (8.1.10)$$

Note that the confidence set gives no information about μ since the random variable above is a pivot variable for σ alone. The confidence set can also be viewed as a bounded confidence *interval* for σ .




Figure 8.1.5: The confidence set for (μ, σ) constructed from S

$1 - \alpha$ confidence sets for (μ, σ) are

1. $V_{\alpha,1}(\mathbf{X}) = \{(\mu, \sigma) : S(\mathbf{X})/s(1 - \alpha) < \sigma < \infty\}$
2. $V_{\alpha,0}(\mathbf{X}) = \{(\mu, \sigma) : 0 < \sigma < S(\mathbf{X})/s(\alpha)\}$

Proof

In the confidence set constructed above, let $p \uparrow 1$ and $p \downarrow 0$, respectively.

The set in part (a) gives a $1 - \alpha$ confidence lower bound for σ and the set in part (b) gives a $1 - \alpha$ confidence upper bound for σ . We can intersect the confidence sets corresponding to the two pivot variables to produce conservative, bounded confidence sets.

If $\alpha, \beta, p, q \in (0, 1)$ with $\alpha + \beta < 1$ then $Z_{\alpha,p} \cap V_{\beta,q}$ is a conservative $1 - (\alpha + \beta)$ confidence set for (μ, σ) .

Proof




Figure 8.1.6: The bounded confidence set for (μ, σ) constructed from (M, S)

The most important location-scale family is the family of normal distributions. The problem of estimation in the normal model is considered in the next section. In the remainder of this section, we will explore another important scale family.

The Exponential Distribution

Recall that the exponential distribution with scale parameter $\sigma \in (0, \infty)$ has probability density function $f(x) = \frac{1}{\sigma} e^{-x/\sigma}$, $x \in [0, \infty)$. It is the scale family associated with the standard exponential distribution, which has probability density function $g(x) = e^{-x}$, $x \in [0, \infty)$. The exponential distribution is widely used to model random times (such as lifetimes and “arrival” times), particularly in the context of the Poisson model. Now suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the exponential distribution with unknown scale parameter σ . Let

$$Y = \sum_{i=1}^n X_i \quad (8.1.11)$$

The random variable $\frac{2}{\sigma} Y$ has the chi-square distribution with $2n$ degrees of freedom, and hence is a pivot variable for σ .

Note that this pivot variable is a multiple of the variable M constructed above for general location-scale families (with $\mu = 0$). For $p \in (0, 1)$ and $k \in (0, \infty)$, let $\chi_k^2(p)$ denote the quantile of order p for the chi-square distribution with k degrees of freedom. For selected values of k and p , $\chi_k^2(p)$ can be obtained from the special distribution calculator or from most statistical software packages.

Recall that

1. $\chi_k^2(p) \rightarrow 0$ as $p \downarrow 0$
2. $\chi_k^2(p) \rightarrow \infty$ as $p \uparrow 1$

For any $\alpha \in (0, 1)$ and any $p \in (0, 1)$, a $1 - \alpha$ confidence interval for σ is

$$\left(\frac{2Y}{\chi_{2n}^2(1-p\alpha)}, \frac{2Y}{\chi_{2n}^2(\alpha-p\alpha)} \right) \quad (8.1.12)$$

Note that

1. $2Y/\chi_{2n}^2(1-\alpha)$ is a $1 - \alpha$ confidence lower bound for σ .
2. $2Y/\chi_{2n}^2(\alpha)$ is a $1 - \alpha$ confidence lower bound for σ .

Of the two-sided confidence intervals constructed above, we would naturally prefer the one with the smallest length, because this interval gives the most information about the parameter σ . However, minimizing the length as a function of p is computationally difficult. The two-sided confidence interval that is typically used is the *equal tailed* interval obtained by letting $p = \frac{1}{2}$:

$$\left(\frac{2Y}{\chi_{2n}^2(1-\alpha/2)}, \frac{2Y}{\chi_{2n}^2(\alpha/2)} \right) \quad (8.1.13)$$

The lifetime of a certain type of component (in hours) has an exponential distribution with unknown scale parameter σ . Ten devices are operated until failure; the lifetimes are 592, 861, 1470, 2412, 335, 3485, 736, 758, 530, 1961.

1. Construct the 95% two-sided confidence interval for σ .
2. Construct the 95% confidence lower bound for σ .
3. Construct the 95% confidence upper bound for σ .

Answer

1. (769.1, 2740.1)
2. 836.7
3. 2421.9

This page titled [8.1: Introduction to Set Estimation](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.2: Estimation the Normal Model

Basic Theory

The Normal Model

The normal distribution is perhaps the most important distribution in the study of mathematical statistics, in part because of the central limit theorem. As a consequence of this theorem, a measured quantity that is subject to numerous small, random errors will have, at least approximately, a normal distribution. Such variables are ubiquitous in statistical experiments, in subjects varying from the physical and biological sciences to the social sciences.

So in this section, we assume that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with mean μ and standard deviation σ . Our goal is to construct confidence intervals for μ and σ individually, and then more generally, confidence sets for (μ, σ) . These are among of the most important special cases of set estimation. A parallel section on Tests in the Normal Model is in the chapter on Hypothesis Testing. First we need to review some basic facts that will be critical for our analysis.

Recall that the sample mean M and sample variance S^2 are

$$M = \frac{1}{n} \sum_{i=1}^n X_i, \quad S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M)^2 \quad (8.2.1)$$

From our study of point estimation, recall that M is an unbiased and consistent estimator of μ while S^2 is an unbiased and consistent estimator of σ^2 . From these basic statistics we can construct the pivot variables that will be used to construct our interval estimates. The following results were established in the section on Special Properties of the Normal Distribution.

Define

$$Z = \frac{M - \mu}{\sigma/\sqrt{n}}, \quad T = \frac{M - \mu}{S/\sqrt{n}}, \quad V = \frac{n-1}{\sigma^2} S^2 \quad (8.2.2)$$

1. Z has the standard normal distribution.
2. T has the student t distribution with $n-1$ degrees of freedom.
3. V has the chi-square distribution with $n-1$ degrees of freedom.
4. Z and V are independent.

It follows that each of these random variables is a pivot variable for (μ, σ) since the distributions do not depend on the parameters, but the variables themselves functionally depend on one or both parameters. Pivot variables Z and T will be used to construct interval estimates of μ while V will be used to construct interval estimates of σ^2 . To construct our estimates, we will need quantiles of these standard distributions. The quantiles can be computed using the special distribution calculator or from most mathematical and statistical software packages. Here is the notation we will use:

Let $p \in (0, 1)$ and $k \in \mathbb{N}_+$.

1. $z(p)$ denotes the quantile of order p for the standard normal distribution.
2. $t_k(p)$ denotes the quantile of order p for the student t distribution with k degrees of freedom.
3. $\chi_k^2(p)$ denotes the quantile of order p for the chi-square distribution with k degrees of freedom

Since the standard normal and student t distributions are symmetric about 0, it follows that $z(1-p) = -z(p)$ and $t_k(1-p) = -t_k(p)$ for $p \in (0, 1)$ and $k \in \mathbb{N}_+$. On the other hand, the chi-square distribution is not symmetric.

Confidence Intervals for μ with σ Known

For our first discussion, we assume that the distribution mean μ is unknown but the standard deviation σ is known. This is not always an artificial assumption. There are often situations where σ is stable over time, and hence is at least approximately known, while μ changes because of different “treatments”. Examples are given in the computational exercises below. The pivot variable Z leads to confidence intervals for μ .

For $\alpha \in (0, 1)$,

1. $\left[M - z\left(1 - \frac{\alpha}{2}\right) \frac{\sigma}{\sqrt{n}}, M + z\left(1 - \frac{\alpha}{2}\right) \frac{\sigma}{\sqrt{n}} \right]$ is a $1 - \alpha$ confidence interval for μ .
2. $M - z(1 - \alpha) \frac{\sigma}{\sqrt{n}}$ is a $1 - \alpha$ confidence lower bound for μ
3. $M + z(1 - \alpha) \frac{\sigma}{\sqrt{n}}$ is a $1 - \alpha$ confidence upper bound for μ

Proof

Since $Z = \frac{M - \mu}{\sigma/\sqrt{n}}$ has the standard normal distribution, each of the following events has probability $1 - \alpha$ by definition of the quantiles:

1. $\left\{ -z\left(1 - \frac{\alpha}{2}\right) \leq \frac{M - \mu}{\sigma/\sqrt{n}} \leq z\left(1 - \frac{\alpha}{2}\right) \right\}$
2. $\left\{ \frac{M - \mu}{\sigma/\sqrt{n}} \geq z(1 - \alpha) \right\}$
3. $\left\{ \frac{M - \mu}{\sigma/\sqrt{n}} \leq -z(1 - \alpha) \right\}$

In each case, solving the inequality for μ gives the result.

These are the standard interval estimates for μ when σ is known. The two-sided confidence interval in (a) is symmetric about the sample mean M , and as the proof shows, corresponds to equal probability $\frac{\alpha}{2}$ in each tail of the distribution of the pivot variable Z . But of course, this is not the only two-sided $1 - \alpha$ confidence interval; we can divide the probability α anyway we want between the left and right tails of the distribution of Z .

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ confidence interval for μ is

$$\left[M - z(1 - p\alpha) \frac{\sigma}{\sqrt{n}}, M - z(\alpha - p\alpha) \frac{\sigma}{\sqrt{n}} \right] \quad (8.2.3)$$

1. $p = \frac{1}{2}$ gives the symmetric, equal-tail confidence interval.
2. $p \rightarrow 0$ gives the interval with the confidence upper bound.
3. $p \rightarrow 1$ gives the interval with the confidence lower bound.

Proof

From the normal distribution of M and the definition of the quantile function,

$$\mathbb{P} \left[z(\alpha - p\alpha) < \frac{M - \mu}{\sigma/\sqrt{n}} < z(1 - p\alpha) \right] = 1 - \alpha \quad (8.2.4)$$

The result then follows by solving for μ in the inequality.

In terms of the distribution of the pivot variable Z , as the proof shows, the two-sided confidence interval above corresponds to $p\alpha$ in the right tail and $(1 - p)\alpha$ in the left tail. Next, let's study the length of this confidence interval.

For $\alpha, p \in (0, 1)$, the (deterministic) length of the two-sided $1 - \alpha$ confidence interval above is

$$L = \frac{[z(1 - p\alpha) - z(\alpha - p\alpha)] \sigma}{\sqrt{n}} \quad (8.2.5)$$

1. L is a decreasing function of α , and $L \downarrow 0$ as $\alpha \uparrow 1$ and $L \uparrow \infty$ as $\alpha \downarrow 0$.
2. L is a decreasing function of n , and $L \downarrow 0$ as $n \uparrow \infty$.
3. L is an increasing function of σ , and $L \downarrow 0$ as $\sigma \downarrow 0$ and $L \uparrow \infty$ as $\sigma \uparrow \infty$.
4. As a function of p , L decreases and then increases, with minimum at the point of symmetry $p = \frac{1}{2}$.

The last result shows again that there is a tradeoff between the confidence level and the length of the confidence interval. If n and p are fixed, we can decrease L , and hence tighten our estimate, only at the expense of decreasing our confidence in the estimate. Conversely, we can increase our confidence in the estimate only at the expense of increasing the length of the interval. In terms of

p , the best of the two-sided $1 - \alpha$ confidence intervals (and the one that is almost always used) is symmetric, equal-tail interval with $p = \frac{1}{2}$:

Use the mean estimation experiment to explore the procedure. Select the normal distribution and select normal pivot. Use various parameter values, confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 times. As the simulation runs, note that the confidence interval successfully captures the mean if and only if the value of the pivot variable is between the quantiles. Note the size and location of the confidence intervals and compare the proportion of successful intervals to the theoretical confidence level.

For the standard confidence intervals, let d denote the distance between the sample mean M and an endpoint. That is,

$$d = z_{\alpha} \frac{\sigma}{\sqrt{n}} \quad (8.2.6)$$

where $z_{\alpha} = z(1 - \alpha/2)$ for the two-sided interval and $z_{\alpha} = z(1 - \alpha)$ for the upper or lower confidence interval. The number d is the *margin of error* of the estimate.

Note that d is deterministic, and the length of the standard two-sided interval is $L = 2d$. In many cases, the first step in the *design of the experiment* is to determine the sample size needed to estimate μ with a given margin of error and a given confidence level.

The sample size needed to estimate μ with confidence $1 - \alpha$ and margin of error d is

$$n = \left\lceil \frac{z_{\alpha}^2 \sigma^2}{d^2} \right\rceil \quad (8.2.7)$$

Proof

This follows by solving for n in the definition of d above, and then rounding up to the next integer.

Note that n varies directly with z_{α}^2 and with σ^2 and inversely with d^2 . This last fact implies a *law of diminishing return* in reducing the margin of error. For example, if we want to reduce a given margin of error by a factor of $\frac{1}{2}$, we must increase the sample size by a factor of 4.

Confidence Intervals for μ with σ Unknown

For our next discussion, we assume that the distribution mean μ and standard deviation σ are unknown, the usual situation. In this case, we can use the T pivot variable, rather than the Z pivot variable, to construct confidence intervals for μ .

For $\alpha \in (0, 1)$,

1. $\left[M - t_{n-1} \left(1 - \frac{\alpha}{2} \right) \frac{S}{\sqrt{n}}, M + t_{n-1} \left(1 - \frac{\alpha}{2} \right) \frac{S}{\sqrt{n}} \right]$ is a $1 - \alpha$ confidence interval for μ .
2. $M - t_{n-1} \left(1 - \alpha \right) \frac{S}{\sqrt{n}}$ is a $1 - \alpha$ lower bound for μ
3. $M + t_{n-1} \left(1 - \alpha \right) \frac{S}{\sqrt{n}}$ is a $1 - \alpha$ upper bound for μ

Proof

Since $T = \frac{M - \mu}{S/\sqrt{n}}$ has the t distribution with $n - 1$ degrees of freedom, each of the following events has probability $1 - \alpha$, by definition of the quantiles:

1. $\left\{ -t_{n-1} \left(1 - \frac{\alpha}{2} \right) \leq \frac{M - \mu}{S/\sqrt{n}} \leq t_{n-1} \left(1 - \frac{\alpha}{2} \right) \right\}$
2. $\left\{ \frac{M - \mu}{S/\sqrt{n}} \geq t_{n-1} (1 - \alpha) \right\}$
3. $\left\{ \frac{M - \mu}{S/\sqrt{n}} \leq -t_{n-1} (1 - \alpha) \right\}$

In each case, solving for μ in the inequality gives the result.

These are the standard interval estimates of μ with σ unknown. The two-sided confidence interval in (a) is symmetric about the sample mean M and corresponds to equal probability $\frac{\alpha}{2}$ in each tail of the distribution of the pivot variable T . As before, this is

not the only confidence interval; we can divide α between the left and right tails any way that we want.

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ confidence interval for μ is

$$\left[M - t_{n-1}(1-p\alpha) \frac{S}{\sqrt{n}}, M - t_{n-1}(\alpha-p\alpha) \frac{S}{\sqrt{n}} \right] \quad (8.2.8)$$

1. $p = \frac{1}{2}$ gives the symmetric, equal-tail confidence interval.
2. $p \rightarrow 0$ gives the interval with the confidence upper bound.
3. $p \rightarrow 1$ gives the interval with the confidence lower bound.

Proof

Since T has the student t distribution with $n - 1$ degrees of freedom, it follows from the definition of the quantiles that

$$\mathbb{P} \left[t_{n-1}(\alpha - p\alpha) < \frac{M - \mu}{S/\sqrt{n}} < t_{n-1}(1 - p\alpha) \right] = 1 - \alpha \quad (8.2.9)$$

The result then follows by solving for μ in the inequality.

The two-sided confidence interval above corresponds to $p\alpha$ in the right tail and $(1 - p)\alpha$ in the left tail of the distribution of the pivot variable T . Next, let's study the length of this confidence interval.

For $\alpha, p \in (0, 1)$, the (random) length of the two-sided $1 - \alpha$ confidence interval above is

$$L = \frac{t_{n-1}(1-p\alpha) - t_{n-1}(\alpha-p\alpha)}{\sqrt{n}} S \quad (8.2.10)$$

1. L is a decreasing function of α , and $L \downarrow 0$ as $\alpha \uparrow 1$ and $L \uparrow \infty$ as $\alpha \downarrow 0$.
2. As a function of p , L decreases and then increases, with minimum at the point of symmetry $p = \frac{1}{2}$.

$$\mathbb{E}(L) = \frac{[t_{n-1}(1-p\alpha) - t_{n-1}(\alpha-p\alpha)] \sqrt{2} \sigma \Gamma(n/2)}{\sqrt{n(n-1)} \Gamma[(n-1)/2]} \quad (8.2.11)$$

$$\text{var}(L) = \frac{1}{n} [t_{n-1}(1-p\alpha) - t_{n-1}(\alpha-p\alpha)]^2 \sigma^2 \left[1 - \frac{2\Gamma^2(n/2)}{(n-1)\Gamma^2[(n-1)/2]} \right] \quad (8.2.12)$$

Proof

Parts (a) and (b) follow from properties of the student quantile function t_{n-1} . Parts (c) and (d) follow from the fact that $\frac{\sqrt{n-1}}{\sigma} S$ has a chi distribution with $n - 1$ degrees of freedom.

Once again, there is a tradeoff between the confidence level and the length of the confidence interval. If n and p are fixed, we can decrease L , and hence tighten our estimate, only at the expense of decreasing our confidence in the estimate. Conversely, we can increase our confidence in the estimate only at the expense of increasing the length of the interval. In terms of p , the best of the two-sided $1 - \alpha$ confidence intervals (and the one that is almost always used) is symmetric, equal-tail interval with $p = \frac{1}{2}$. Finally, note that it does not really make sense to consider L as a function of S , since S is a statistic rather than an algebraic variable. Similarly, it does not make sense to consider L as a function of n , since changing n means new data and hence a new value of S .

Use the mean estimation experiment to explore the procedure. Select the normal distribution and the T pivot. Use various parameter values, confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 times. As the simulation runs, note that the confidence interval successfully captures the mean if and only if the value of the pivot variable is between the quantiles. Note the size and location of the confidence intervals and compare the proportion of successful intervals to the theoretical confidence level.

Confidence Intervals for σ^2

Next we will construct confidence intervals for σ^2 using the pivot variable V given [above](#)

For $\alpha \in (0, 1)$,

1. $\left[\frac{n-1}{\chi_{n-1}^2(1-\alpha/2)} S^2, \frac{n-1}{\chi_{n-1}^2(\alpha/2)} S^2 \right]$ is a $1 - \alpha$ confidence interval for σ^2
2. $\frac{n-1}{\chi_{n-1}^2(1-\alpha)} S^2$ is a $1 - \alpha$ confidence lower bound for σ^2
3. $\frac{n-1}{\chi_{n-1}^2(\alpha)} S^2$ is a $1 - \alpha$ confidence upper bound for σ^2 .

Proof

Since $V = \frac{n-1}{\sigma^2} S^2$ has the chi-square distribution with $n - 1$ degrees of freedom, each of the following events has probability $1 - \alpha$ by definition of the quantiles:

1. $\left\{ \chi_{n-1}^2(\alpha/2) \leq \frac{n-1}{\sigma^2} S^2 \leq \chi_{n-1}^2(1-\alpha/2) \right\}$
2. $\left\{ \frac{n-1}{\sigma^2} S^2 \leq \chi_{n-1}^2(1-\alpha) \right\}$
3. $\left\{ \frac{n-1}{\sigma^2} S^2 \geq \chi_{n-1}^2(\alpha) \right\}$

In each case, solving for σ^2 in the inequality give the result.

These are the standard interval estimates for σ^2 . The two-sided interval in (a) is the *equal-tail* interval, corresponding to probability $\alpha/2$ in each tail of the distribution of the pivot variable V . Note however that this interval is not symmetric about the sample variance S^2 . Once again, we can partition the probability α between the left and right tails of the distribution of V any way that we like.

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ confidence interval for σ^2 is

$$\left[\frac{n-1}{\chi_{n-1}^2(1-p\alpha)} S^2, \frac{n-1}{\chi_{n-1}^2(\alpha-p\alpha)} S^2 \right] \quad (8.2.13)$$

1. $p = \frac{1}{2}$ gives the equal-tail $1 - \alpha$ confidence interval.
2. $p \rightarrow 0$ gives the interval with the $1 - \alpha$ upper bound
3. $p \rightarrow 1$ gives the interval with the $1 - \alpha$ lower bound.

In terms of the distribution of the pivot variable V , the confidence interval above corresponds to $p\alpha$ in the right tail and $(1-p)\alpha$ in the left tail. Once again, let's look at the length of the general two-sided confidence interval. The length is random, but is a multiple of the sample variance S^2 . Hence we can compute the expected value and variance of the length.

For $\alpha, p \in (0, 1)$, the (random) length of the two-sided confidence interval in the last theorem is

$$L = \left[\frac{1}{\chi_{n-1}^2(\alpha-p\alpha)} - \frac{1}{\chi_{n-1}^2(1-p\alpha)} \right] (n-1) S^2 \quad (8.2.14)$$

1. $\mathbb{E}(L) = \left[\frac{1}{\chi_{n-1}^2(\alpha-p\alpha)} - \frac{1}{\chi_{n-1}^2(1-p\alpha)} \right] (n-1) \sigma^2$
2. $\text{var}(L) = 2 \left[\frac{1}{\chi_{n-1}^2(\alpha-p\alpha)} - \frac{1}{\chi_{n-1}^2(1-p\alpha)} \right]^2 (n-1) \sigma^4$

To construct an optimal two-sided confidence interval, it would be natural to find p that minimizes the expected length. This is a complicated problem, but it turns out that for large n , the equal-tail interval with $p = \frac{1}{2}$ is close to optimal. Of course, taking square roots of the endpoints of any of the confidence intervals for σ^2 gives $1 - \alpha$ confidence intervals for the distribution standard deviation σ .

Use variance estimation experiment to explore the procedure. Select the normal distribution. Use various parameter values, confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 time. As the simulation runs, note that the confidence interval successfully captures the standard deviation if and only if the value of the pivot variable is between the quantiles. Note the size and location of the confidence intervals and compare the proportion of successful intervals to the theoretical confidence level.

Confidence Sets for (μ, σ)

In the discussion above, we constructed confidence intervals for μ and for σ separately (again, usually both parameters are unknown). In our next discussion, we will consider confidence *sets* for the parameter point (μ, σ) . These sets will be subsets of the underlying parameter space $\mathbb{R} \times (0, \infty)$.

Confidence Sets Constructed from the Pivot Variables

Each of the pivot variables Z , T , and V can be used to construct confidence sets for (μ, σ) . In isolation, each will produce an unbounded confidence set, not surprising since, we are using a single pivot variable to estimate two parameters. We consider the normal pivot variable Z first.

For any $\alpha, p \in (0, 1)$, a $1 - \alpha$ level confidence set for (μ, σ) is

$$Z_{\alpha,p} = \left\{ (\mu, \sigma) : M - z(1-p\alpha) \frac{\sigma}{\sqrt{n}} < \mu < M - z(\alpha-p\alpha) \frac{\sigma}{\sqrt{n}} \right\} \quad (8.2.15)$$

The confidence set is a “cone” in the (μ, σ) parameter space, with vertex at $(M, 0)$ and boundary lines of slopes $-\sqrt{n}/z(1-p\alpha)$ and $-\sqrt{n}/z(\alpha-p\alpha)$

Proof

From the normal distribution of M and the definition of the quantile function,

$$\mathbb{P} \left[z(\alpha-p\alpha) < \frac{M-\mu}{\sigma/\sqrt{n}} < z(1-p\alpha) \right] = 1 - \alpha \quad (8.2.16)$$

The result then follows by solving for μ in the inequality.

The confidence cone is shown in the graph below. (Note, however, that both slopes might be negative or both positive.)




Figure 8.2.1: The confidence set based on the normal pivot variable

The pivot variable T leads to the following result:

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ level confidence set for (μ, σ) is

$$T_{\alpha,p} = \left\{ (\mu, \sigma) : M - t_{n-1}(1-p\alpha) \frac{S}{\sqrt{n}} < \mu < M - t_{n-1}(\alpha-p\alpha) \frac{S}{\sqrt{n}} \right\} \quad (8.2.17)$$

Proof




Figure 8.2.2: The confidence set based on the T pivot variable

By design, this confidence set gives no information about σ . Finally, the pivot variable V leads to the following result:

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ level confidence set for (μ, σ) is

$$V_{\alpha,p} = \left\{ (\mu, \sigma) : \frac{(n-1)S^2}{\chi_{n-1}^2(1-p\alpha)} < \sigma^2 < \frac{(n-1)S^2}{\chi_{n-1}^2(\alpha-p\alpha)} \right\} \quad (8.2.18)$$

Proof




Figure 8.2.3: The confidence set based on the pivot variable V

By design, this confidence set gives no information about μ .

Intersections

We can now form intersections of some of the confidence sets constructed above to obtain *bounded* confidence sets for (μ, σ) . We will use the fact that the sample mean M and the sample variance S^2 are independent, one of the most important special properties of a normal sample. We will also need the result from the Introduction on the intersection of confidence intervals. In the following theorems, suppose that $\alpha, \beta, p, q \in (0, 1)$ with $\alpha + \beta < 1$.

The set $T_{\alpha,p} \cap V_{\beta,q}$ is a conservative $1 - (\alpha + \beta)$ confidence sets for (μ, σ) .




Figure 8.2.4: The confidence set $T_{\alpha,p} \cap V_{\beta,q}$

The set $Z_{\alpha,p} \cap V_{\beta,q}$ is a $(1 - \alpha)(1 - \beta)$ confidence set for (μ, σ) .




Figure 8.2.5: The confidence set $Z_{\alpha,p} \cap V_{\beta,q}$

It is interesting to note that the confidence set $T_{\alpha,p} \cap V_{\beta,q}$ is a product set as a subset of the parameter space, but is not a product set as a subset of the sample space. By contrast, the confidence set $Z_{\alpha,p} \cap V_{\beta,q}$ is not a product set as a subset of the parameter space, but is a product set as a subset of the sample space.

Exercises

Robustness

The main assumption that we made was that the underlying sampling distribution is normal. Of course, in real statistical problems, we are unlikely to know much about the sampling distribution, let alone whether or not it is normal. When a statistical procedure works reasonably well, even when the underlying assumptions are violated, the procedure is said to be *robust*. In this subsection, we will explore the robustness of the estimation procedures for μ and σ .

Suppose in fact that the underlying distribution is not normal. When the sample size n is relatively large, the distribution of the sample mean will still be approximately normal by the central limit theorem. Thus, our interval estimates of μ may still be approximately valid.

Use the simulation of the mean estimation experiment to explore the procedure. Select the gamma distribution and select student pivot. Use various parameter values, confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 times. Note the size and location of the confidence intervals and compare the proportion of successful intervals to the theoretical confidence level.

In the mean estimation experiment, repeat the previous exercise with the uniform distribution.

How large n needs to be for the interval estimation procedures of μ to work well depends, of course, on the underlying distribution; the more this distribution deviates from normality, the larger n must be. Fortunately, convergence to normality in the central limit theorem is rapid and hence, as you observed in the exercises, we can get away with relatively small sample sizes (30 or more) in most cases.

In general, the interval estimation procedures for σ are not robust; there is no analog of the central limit theorem to save us from deviations from normality.

In variance estimation experiment, select the gamma distribution. Use various parameter values, confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 times. Note the size and location of the confidence intervals and compare the proportion of successful intervals to the theoretical confidence level.

In variance estimation experiment, select the uniform distribution. Use various parameter values, confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 times. Note the size and location of the confidence intervals and compare the proportion of successful intervals to the theoretical confidence level.

Computational Exercises

In the following exercises, use the equal-tailed construction for two-sided confidence intervals, unless otherwise instructed.

The length of a certain machined part is supposed to be 10 centimeters but due to imperfections in the manufacturing process, the actual length is a normally distributed with mean μ and variance σ^2 . The variance is due to inherent factors in the process, which remain fairly stable over time. From historical data, it is known that $\sigma = 0.3$. On the other hand, μ may be set by adjusting various parameters in the process and hence may change to an unknown value fairly frequently. A sample of 100 parts has mean 10.2.

1. Construct the 95% confidence interval for μ .
2. Construct the 95% confidence upper bound for μ .
3. Construct the 95% confidence lower bound for μ .

Answer

1. (10.1, 10.26)
2. 10.25
3. 10.15

Suppose that the weight of a bag of potato chips (in grams) is a normally distributed random variable with mean μ and standard deviation σ , both unknown. A sample of 75 bags has mean 250 and standard deviation 10.

1. Construct the 90% confidence interval for μ .
2. Construct the 90% confidence interval for σ .
3. Construct a conservative 90% confidence rectangle for (μ, σ) .

Answer

1. (248.1, 251.9)
2. (8.8, 11.6)
3. (247.70, 252.30) \times (8.62, 11.92)

At a telemarketing firm, the length of a telephone solicitation (in seconds) is a normally distributed random variable with mean μ and standard deviation σ , both unknown. A sample of 50 calls has mean length 300 and standard deviation 60.

1. Construct the 95% confidence upper bound for μ .
2. Construct the 95% confidence lower bound for σ .

Answer

1. 314.3.
2. 51.6.

At a certain farm the weight of a peach (in ounces) at harvest time is a normally distributed random variable with standard deviation 0.5. How many peaches must be sampled to estimate the mean weight with a margin of error ± 2 and with 95% confidence.

Answer

25

The hourly salary for a certain type of construction work is a normally distributed random variable with standard deviation \$1.25 and unknown mean μ . How many workers must be sampled to construct a 95% confidence lower bound for μ with margin of error \$0.25?

Answer

68

Data Analysis Exercises

In Michelson's data, assume that the measured speed of light has a normal distribution with mean μ and standard deviation σ , both unknown.

1. Construct the 95% confidence interval for μ . Is the "true" value of the speed of light in this interval?
2. Construct the 95% confidence interval for σ .
3. Explore, in an informal graphical way, the assumption that the underlying distribution is normal.

Answer

1. (836.8, 868.0) No, the true value is not in the interval.
2. (69.4, 91.8)

In Cavendish's data, assume that the measured density of the earth has a normal distribution with mean μ and standard deviation σ , both unknown.

1. Construct the 95% confidence interval for μ . Is the "true" value of the density of the earth in this interval?
2. Construct the 95% confidence interval for σ .
3. Explore, in an informal graphical way, the assumption that the underlying distribution is normal.

Answer

1. (5.364, 5.532) Yes, the true value is in the interval.
2. (0.1725, 0.3074)

In Short's data, assume that the measured parallax of the sun has a normal distribution with mean μ and standard deviation σ , both unknown.

1. Construct the 95% confidence interval for μ . Is the "true" value of the parallax of the sun in this interval?
2. Construct the 95% confidence interval for σ .
3. Explore, in an informal graphical way, the assumption that the underlying distribution is normal.

Answer

1. (8.410, 8.822) Yes, the true value is in the interval.
2. (0.629, 0.927)

Suppose that the length of an iris petal of a given type (Setosa, Verginica, or Versicolor) is normally distributed. Use Fisher's iris data to construct 90% two-sided confidence intervals for each of the following parameters.

1. The mean length of a Sertosa iris petal.
2. The mean length of a Vergnica iris petal.
3. The mean length of a Versicolor iris petal.

Answer

1. (14.21, 15.03)
2. (54.21, 56.83)
3. (41.95, 44.49)

This page titled [8.2: Estimation the Normal Model](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.3: Estimation in the Bernoulli Model

Introduction

Recall that an *indicator variable* is a random variable that just takes the values 0 and 1. In applications, an indicator variable indicates which of two complementary events in a random experiment has occurred. Typical examples include

- A manufactured item subject to unavoidable random factors is either defective or acceptable.
- A voter selected from a population either supports a particular candidate or does not.
- A person selected from a population either does or does not have a particular medical condition.
- A student in a class either passes or fails a standardized test.
- A sample of radioactive material either does or does not emit an alpha particle in a specified ten-second period.

Recall also that the distribution of an indicator variable is known as the *Bernoulli distribution*, named for Jacob Bernoulli, and has probability density function given by $\mathbb{P}(X = 1) = p$, $\mathbb{P}(X = 0) = 1 - p$, where $p \in (0, 1)$ is the basic parameter. In the context of the examples above,

- p is the probability that the manufactured item is defective.
- p is the proportion of voters in the population who favor the candidate.
- p is the proportion of persons in the population that have the medical condition.
- p is the probability that a student in the class will pass the exam.
- p is the probability that the material will emit an alpha particle in the specified period.

Recall that the mean and variance of the Bernoulli distribution are $\mathbb{E}(X) = p$ and $\text{var}(X) = p(1 - p)$. Often in statistical applications, p is unknown and must be estimated from sample data. In this section, we will see how to construct interval estimates for the parameter from sample data. A parallel section on Tests in the Bernoulli Model is in the chapter on Hypothesis Testing.

The One-Sample Model

Preliminaries

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Bernoulli distribution with unknown parameter $p \in [0, 1]$. That is, \mathbf{X} is a sequence of Bernoulli trials. From the examples in the introduction above, note that often the underlying experiment is to sample at random from a dichotomous population. When the sampling is *with* replacement, \mathbf{X} really is a sequence of Bernoulli trials. When the sampling is *without* replacement, the variables are dependent, but the Bernoulli model is still approximately valid if the population size is large compared to the sample size n . For more on these points, see the discussion of sampling with and without replacement in the chapter on Finite Sampling Models.

Note that the sample mean of our data vector \mathbf{X} , namely

$$M = \frac{1}{n} \sum_{i=1}^n X_i \quad (8.3.1)$$

is the sample proportion of objects of the type of interest. By the central limit theorem, the standard score

$$Z = \frac{M - p}{\sqrt{p(1 - p)/n}} \quad (8.3.2)$$

has approximately a standard normal distribution and hence is (approximately) a pivot variable for p . For a given sample size n , the distribution of Z is closest to normal when p is near $\frac{1}{2}$ and farthest from normal when p is near 0 or 1 (extreme). Because the pivot variable is (approximately) normally distributed, the construction of confidence intervals for p in this model is similar to the construction of confidence intervals for the distribution mean μ in the normal model. But of course all of the confidence intervals so constructed are approximate.

As usual, for $r \in (0, 1)$, let $z(r)$ denote the quantile of order r for the standard normal distribution. Values of $z(r)$ can be obtained from the special distribution calculator, or from most statistical software packages.

Basic Confidence Intervals

For $\alpha \in (0, 1)$, the following are approximate $1 - \alpha$ confidence sets for p :

1. $\{p \in [0, 1] : M - z(1 - \alpha/2)\sqrt{p(1-p)/n} \leq p \leq M + z(1 - \alpha/2)\sqrt{p(1-p)/n}\}$
2. $\{p \in [0, 1] : p \leq M + z(1 - \alpha)\sqrt{p(1-p)/n}\}$
3. $\{p \in [0, 1] : M - z(1 - \alpha)\sqrt{p(1-p)/n} \leq p\}$

Proof

From our discussion above, $(M - p)/\sqrt{p(1-p)/n}$ has approximately a standard normal distribution. Hence by definition of the quantiles,

1. $\mathbb{P}[-z(1 - \alpha/2) \leq (M - p)/\sqrt{p(1-p)/n} \leq z(1 - \alpha/2)] \approx 1 - \alpha$
2. $\mathbb{P}[-z(1 - \alpha) \leq (M - p)/\sqrt{p(1-p)/n}] \approx 1 - \alpha$
3. $\mathbb{P}[(M - p)/\sqrt{p(1-p)/n} \leq z(1 - \alpha)] \approx 1 - \alpha$

Solving the inequalities for p in the numerator of $(M - p)/\sqrt{p(1-p)/n}$ for each event gives the corresponding confidence set.

These confidence sets are actually intervals, known as the *Wilson intervals*, in honor of Edwin Wilson.

The confidence sets for p in (1) are intervals. Let

$$U(z) = \frac{n}{n + z^2} \left(M + \frac{z^2}{2n} + z \sqrt{\frac{M(1-M)}{n} + \frac{z^2}{4n^2}} \right) \quad (8.3.3)$$

Then the following have approximate confidence level $1 - \alpha$ for p .

1. The two-sided interval $[U[-z(1 - \alpha/2)], U[z(1 - \alpha/2)]]$.
2. The upper bound $U[z(1 - \alpha)]$.
3. The lower bound $U[-z(1 - \alpha)]$.

Proof

This follows by solving the inequalities in (1) for p . For each inequality, we can isolate the square root term, and then square both sides. This gives quadratic inequalities, which can be solved using the quadratic formula.

As usual, the *equal-tailed* confidence interval in (a) is not the only two-sided $1 - \alpha$ confidence interval for p . We can divide the α probability between the left and right tails of the standard normal distribution in any way that we please.

For $\alpha, r \in (0, 1)$, an approximate two-sided $1 - \alpha$ confidence interval for p is $[U[z(\alpha - r\alpha)], U[z(1 - r\alpha)]]$ where U is the function in (2).

Proof

As in the proof of (1),

$$\mathbb{P} \left[z(\alpha - r\alpha) \leq \frac{M - p}{\sqrt{p(1-p)/n}} \leq z(1 - r\alpha) \right] \approx 1 - \alpha \quad (8.3.4)$$

Solving for p with the help of the quadratic formula gives the result.

In practice, the equal-tailed $1 - \alpha$ confidence interval in part (a) of (2), obtained by setting $r = \frac{1}{2}$, is the one that is always used. As $r \uparrow 1$, the right endpoint converges to the $1 - \alpha$ confidence upper bound in part (b), and as $r \downarrow 0$ the left endpoint converges to the $1 - \alpha$ confidence lower bound in part (c).

Simplified Confidence Intervals

Simplified approximate $1 - \alpha$ confidence intervals for p can be obtained by replacing the distribution mean p by the sample mean M in the extreme parts of the inequalities in (1).

For $\alpha \in (0, 1)$, the following have approximate confidence level $1 - \alpha$ for p :

1. The two-sided interval with endpoints $M \pm z(1 - \alpha/2)\sqrt{M(1 - M)/n}$.
2. The upper bound $M + z(1 - \alpha)\sqrt{M(1 - M)/n}$.
3. The lower bound $M - z(1 - \alpha)\sqrt{M(1 - M)/n}$.

Proof

As noted, these results follows from the confidence set in (1) by replacing p with M in the expression $\sqrt{p(1 - p)/n}$.

These confidence intervals are known as *Wald intervals*, in honor of Abraham Wald.. Note that the Wald interval can also be obtained from the Wilson intervals in (2) by assuming that n is large compared to z , so that $n/(n + z^2) \approx 1$, $z^2/2n \approx 0$, and $z^2/4n^2 \approx 0$. Note that this interval in (c) is symmetric about the sample proportion M but that the length of the interval, as well as the center is random. This is the two-sided interval that is normally used.

Use the simulation of the proportion estimation experiment to explore the procedure. Use various values of p and various confidence levels, sample sizes, and interval types. For each configuration, run the experiment 1000 times and compare the proportion of successful intervals to the theoretical confidence level.

As always, the equal-tailed interval in (4) is not the only two-sided, $1 - \alpha$ confidence interval.

For $\alpha, r \in (0, 1)$, an approximate two-sided $1 - \alpha$ confidence interval for p is

$$\left[M - z(1 - r\alpha)\sqrt{\frac{M(1 - M)}{n}}, M - z(\alpha - r\alpha)\sqrt{\frac{M(1 - M)}{n}} \right] \quad (8.3.5)$$

The interval with smallest length is the equal-tail interval with $r = \frac{1}{2}$.

Conservative Confidence Intervals

Note that the function $p \mapsto p(1 - p)$ on the interval $[0, 1]$ is maximized when $p = \frac{1}{2}$ and thus the maximum value is $\frac{1}{4}$. We can obtain conservative confidence intervals for p from the basic confidence intervals by using this fact.

For $\alpha \in (0, 1)$, the following have approximate confidence level at least $1 - \alpha$ for p :

1. The two-sided interval with endpoints $M \pm z(1 - \alpha/2)\frac{1}{2\sqrt{n}}$.
2. The upper bound $M + z(1 - \alpha)\frac{1}{2\sqrt{n}}$.
3. The lower bound $M - z(1 - \alpha)\frac{1}{2\sqrt{n}}$.

Proof

As noted, these results follows from the confidence sets in (1) by replacing p with $\frac{1}{2}$ in the expression $\sqrt{p(1 - p)/n}$.

Note that the confidence interval in (a) is symmetric about the sample proportion M and that the length of the interval is deterministic. Of course, the conservative confidence intervals will be larger than the approximate simplified confidence intervals in (4). The conservative estimate can be used to design the experiment. Recall that the *margin of error* is the distance between the sample proportion M and an endpoint of the confidence interval.

A conservative estimate of the sample size n needed to estimate p with confidence $1 - \alpha$ and margin of error d is

$$n = \left\lceil \frac{z_\alpha^2}{4d^2} \right\rceil \quad (8.3.6)$$

where $z_\alpha = z(1 - \alpha/2)$ for the two-sided interval and $z_\alpha = z(1 - \alpha)$ for the confidence upper or lower bound.

Proof

With confidence level $1 - \alpha$, the margin of error is $z_\alpha \frac{1}{2\sqrt{n}}$. Setting this equal to the prescribed value d and solving gives the result.

As always, the equal-tailed interval in (7) is not the only two-sided, conservative, $1 - \alpha$ confidence interval.

For $\alpha, r \in (0, 1)$, an approximate two-sided, conservative $1 - \alpha$ confidence interval for p is

$$\left[M - z(1 - r\alpha) \frac{1}{2\sqrt{n}}, M - z(\alpha - r\alpha) \frac{1}{2\sqrt{n}} \right] \quad (8.3.7)$$

The interval with smallest length is the equal-tail interval with $r = \frac{1}{2}$.

The Two-Sample Model

Preliminaries

Often we have two underlying Bernoulli distributions, with parameters $p_1, p_2 \in [0, 1]$ and we would like to estimate the difference $p_1 - p_2$. This problem could arise in the following typical examples:

- In a quality control setting, suppose that p_1 is the proportion of defective items produced under one set of manufacturing conditions while p_2 is the proportion of defectives under a different set of conditions.
- In an election, suppose that p_1 is the proportion of voters who favor a particular candidate at one point in the campaign, while p_2 is the proportion of voters who favor the candidate at a later point (perhaps after a scandal has erupted).
- Suppose that p_1 is the proportion of students who pass a certain standardized test with the usual test preparation methods while p_2 is the proportion of students who pass the test with a new set of preparation methods.
- Suppose that p_1 is the proportion of unvaccinated persons in a certain population who contract a certain disease, while p_2 is the proportion of vaccinated person who contract the disease.

Note that several of these examples can be thought of as *treatment-control* problems. Of course, we could construct interval estimates I_1 for p_1 and I_2 for p_2 separately, as in the subsections above. But as we noted in the Introduction, if these two intervals have confidence level $1 - \alpha$, then the product set $I_1 \times I_2$ has confidence level $(1 - \alpha)^2$ for (p_1, p_2) . So if $p_1 - p_2$ is our parameter of interest, we will use a different approach.

Simplified Confidence Intervals

Suppose now that $\mathbf{X} = (X_1, X_2, \dots, X_{n_1})$ is a random sample of size n_1 from the Bernoulli distribution with parameter p_1 , and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_{n_2})$ is a random sample of size n_2 from the Bernoulli distribution with parameter p_2 . We assume that the samples \mathbf{X} and \mathbf{Y} are independent. Let

$$M_1 = \frac{1}{n_1} \sum_{i=1}^{n_1} X_i, \quad M_2 = \frac{1}{n_2} \sum_{i=1}^{n_2} Y_i \quad (8.3.8)$$

denote the sample means (sample proportions) for the samples \mathbf{X} and \mathbf{Y} . A natural point estimate for $p_1 - p_2$, and the building block for our interval estimate, is $M_1 - M_2$. As noted in the one-sample model, if n_i is large, M_i has an approximate normal distribution with mean p_i and variance $p_i(1 - p_i)/n_i$ for $i \in \{1, 2\}$. Since the samples are independent, so are the sample means. Hence $M_1 - M_2$ has an approximate normal distribution with mean $p_1 - p_2$ and variance $p_1(1 - p_1)/n_1 + p_2(1 - p_2)/n_2$. We now have all the tools we need for a simplified, approximate confidence interval for $p_1 - p_2$.

For $\alpha \in (0, 1)$, the following have approximate confidence level $1 - \alpha$ for $p_1 - p_2$:

1. The two-sided interval with endpoints $(M_1 - M_2) \pm z(1 - \alpha/2) \sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2}$.
2. The lower bound $(M_1 - M_2) - z(1 - \alpha) \sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2}$.
3. The upper bound $(M_1 - M_2) + z(1 - \alpha) \sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2}$.

Proof

As noted above, if n_1 and n_2 are large,

$$\frac{(M_1 - M_2) - (p_1 - p_2)}{\sqrt{p_1(1 - p_1)/n_1 + p_2(1 - p_2)/n_2}} \quad (8.3.9)$$

has approximately a standard normal distribution, and hence so does

$$Z = \frac{(M_1 - M_2) - (p_1 - p_2)}{\sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2}} \quad (8.3.10)$$

1. $\mathbb{P}[-z(1 - \alpha/2) \leq Z \leq z(1 - \alpha/2)] \approx 1 - \alpha$. Solving for $p_1 - p_2$ gives the two-sided confidence interval.
2. $\mathbb{P}[-z(1 - \alpha) \leq Z] \approx 1 - \alpha$. Solving for $p_1 - p_2$ gives the confidence upper bound.
3. $\mathbb{P}[Z \leq z(1 - \alpha/2)] \approx 1 - \alpha$. Solving for $p_1 - p_2$ gives the confidence lower bound.

As always, the equal-tailed interval in (a) is not the only approximate two-sided $1 - \alpha$ confidence interval.

For $\alpha, r \in (0, 1)$, an approximate $1 - \alpha$ confidence set for $p_1 - p_2$ is

$$\begin{aligned} & \left[(M_1 - M_2) - z(1 - r\alpha) \sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2}, (M_1 - M_2) \right. \\ & \quad \left. - z(\alpha - r\alpha) \sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2} \right] \end{aligned} \quad (8.3.11)$$

Proof

As noted in the proof of the previous theorem,

$$Z = \frac{(M_1 - M_2) - (p_1 - p_2)}{\sqrt{M_1(1 - M_1)/n_1 + M_2(1 - M_2)/n_2}} \quad (8.3.12)$$

has approximately a standard normal distribution if n_1 and n_2 are large. Hence $\mathbb{P}[-z(\alpha - r\alpha) \leq Z \leq z(1 - r\alpha)] \approx 1 - \alpha$. Solving for $p_1 - p_2$ gives the two-sided confidence interval.

Conservative Confidence Intervals

Once again, $p \mapsto p(1 - p)$ is maximized when $p = \frac{1}{2}$ with maximum value $\frac{1}{4}$. We can use this to construct approximate conservative confidence intervals for $p_1 - p_2$.

For $\alpha \in (0, 1)$, the following have approximate confidence level at least $1 - \alpha$ for $p_1 - p_2$:

1. The two-sided interval with endpoints $(M_1 - M_2) \pm \frac{1}{2} z(1 - \alpha/2) \sqrt{1/n_1 + 1/n_2}$.
2. The lower bound $(M_1 - M_2) - \frac{1}{2} z(1 - \alpha) \sqrt{1/n_1 + 1/n_2}$.
3. The upper bound $(M_1 - M_2) + \frac{1}{2} z(1 - \alpha) \sqrt{1/n_1 + 1/n_2}$.

Proof

These results follow from the previous theorem by replacing $M_1(1 - M_1)$ and $M_2(1 - M_2)$ each with $\frac{1}{4}$.

Computational Exercises

In a poll of 1000 registered voters in a certain district, 427 prefer candidate X. Construct the 95% two-sided confidence interval for the proportion of all registered voters in the district that prefer X.

Answer

(0.396, 0.458)

A coin is tossed 500 times and results in 302 heads. Construct the 95% confidence lower bound for the probability of heads. Do you believe that the coin is fair?

Answer

0.579. No, the coin is almost certainly not fair.

A sample of 400 memory chips from a production line are tested, and 30 are defective. Construct the conservative 90% two-sided confidence interval for the proportion of defective chips.

Answer

(0.034, 0.116)

A drug company wants to estimate the proportion of persons who will experience an adverse reaction to a certain new drug. The company wants a two-sided interval with margin of error 0.03 with 95% confidence. How large should the sample be?

Answer

1068

An advertising agency wants to construct a 99% confidence lower bound for the proportion of dentists who recommend a certain brand of toothpaste. The margin of error is to be 0.02. How large should the sample be?

Answer

3382

The Buffon trial data set gives the results of 104 repetitions of Buffon's needle experiment. Theoretically, the data should correspond to Bernoulli trials with $p = 2/\pi$, but because real students dropped the needle, the true value of p is unknown. Construct a 95% confidence interval for p . Do you believe that p is the theoretical value?

Answer

(0.433, 0.634) The theoretical value is approximately 0.637, which is not in the confidence interval.

A manufacturing facility has two production lines for a certain item. In a sample of 150 items from line 1, 12 are defective. From a sample of 130 items from line 2, 10 are defective. Construct the two-sided 95% confidence interval for $p_1 - p_2$, where p_i is the proportion of defective items from line i , for $i \in \{1, 2\}$

Answer

$[-0.050, 0.056]$

The vaccine for influenza is tailored each year to match the predicted dominant strain of influenza. Suppose that of 500 unvaccinated persons, 45 contracted the flu in a certain time period. Of 300 vaccinated persons, 20 contracted the flu in the same time period. Construct the two-sided 99% confidence interval for $p_1 - p_2$, where p_1 is the incidence of flu in the unvaccinated population and p_2 the incidence of flu in the vaccinated population.

This page titled [8.3: Estimation in the Bernoulli Model](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.4: Estimation in the Two-Sample Normal Model

As we have noted before, the normal distribution is perhaps the most important distribution in the study of mathematical statistics, in part because of the central limit theorem. As a consequence of this theorem, measured quantities that are subject to numerous small, random errors will have, at least approximately, normal distributions. Such variables are ubiquitous in statistical experiments, in subjects varying from the physical and biological sciences to the social sciences.

In this section, we will study estimation problems in the two-sample normal model and in the bivariate normal model. This section parallels the section on Tests in the Two-Sample Normal Model in the Chapter on Hypothesis Testing.

The Two-Sample Normal Model

Preliminaries

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_m)$ is a random sample of size m from the normal distribution with mean μ and standard deviation σ , and that $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ is a random sample of size n from the normal distribution with mean ν and standard deviation τ . Moreover, suppose that the samples \mathbf{X} and \mathbf{Y} are independent. Usually, the parameters are unknown, so the parameter space for our vector of parameters (μ, ν, σ, τ) is $\mathbb{R}^2 \times (0, \infty)^2$.

This type of situation arises frequently when the random variables represent a measurement of interest for the objects of the population, and the samples correspond to two different treatments. For example, we might be interested in the blood pressure of a certain population of patients. The \mathbf{X} vector records the blood pressures of a control sample, while the \mathbf{Y} vector records the blood pressures of the sample receiving a new drug. Similarly, we might be interested in the yield of an acre of corn. The \mathbf{X} vector records the yields of a sample receiving one type of fertilizer, while the \mathbf{Y} vector records the yields of a sample receiving a different type of fertilizer.

Usually our interest is in a comparison of the parameters (either the means or standard deviations) for the two sampling distributions. In this section we will construct confidence intervals for the difference of the distribution means $\nu - \mu$ and for the ratio of the distribution variances τ^2 / σ^2 . As with previous estimation problems, the construction depends on finding appropriate pivot variables.

For a generic sample $\mathbf{U} = (U_1, U_2, \dots, U_k)$ from a distribution with mean a , we will use our standard notation for the sample mean and for the sample variance.

$$M(\mathbf{U}) = \frac{1}{k} \sum_{i=1}^k U_i \quad (8.4.1)$$

$$S^2(\mathbf{U}) = \frac{1}{k-1} \sum_{i=1}^k [U_i - M(\mathbf{U})]^2 \quad (8.4.2)$$

We will need to also recall the special properties of these statistics when the sampling distribution is normal. The special pivot distributions that will play a fundamental role in this section are the standard normal, the student t , and the Fisher F distributions. To construct our interval estimates we will need the quantiles of these distributions. The quantiles can be computed using the special distribution calculator or from most mathematical and statistical software packages. Here is the notation we will use:

Let $p \in (0, 1)$ and let $j, k \in \mathbb{N}_+$.

1. $z(p)$ denotes the quantile of order p for the standard normal distribution.
2. $t_k(p)$ denotes the quantile of order p for the student t distribution with k degrees of freedom.
3. $f_{j,k}(p)$ denotes the quantile of order p for the student f distribution with j degrees of freedom in the numerator and k degrees of freedom in the denominator.

Recall that by symmetry, $z(p) = -z(1-p)$ and $t_k(p) = -t_k(1-p)$ for $p \in (0, 1)$ and $k \in \mathbb{N}_+$. On the other hand, there is no simple relationship between the left and right tail probabilities of the F distribution.

Confidence Intervals for the Difference of the Means with Known Variances

First we will construct confidence intervals for $\nu - \mu$ under the assumption that the distribution variances σ^2 and τ^2 are known. This is not always an artificial assumption. As in the one sample normal model, the variances are sometime stable, and hence are at least approximately known, while the means change under different treatments. First recall the following basic facts:

The difference of the sample means $M(\mathbf{Y}) - M(\mathbf{X})$ has the normal distribution with mean $\nu - \mu$ and variance $\sigma^2/m + \tau^2/n$. Hence the standard score of the difference of the sample means

$$Z = \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - (\nu - \mu)}{\sqrt{\sigma^2/m + \tau^2/n}} \quad (8.4.3)$$

has the standard normal distribution. Thus, this variable is a pivotal variable for $\nu - \mu$ when σ, τ are known.

The basic confidence interval and upper and lower bound are now easy to construct.

For $\alpha \in (0, 1)$,

1. $\left[M(\mathbf{Y}) - M(\mathbf{X}) - z\left(1 - \frac{\alpha}{2}\right) \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}}, M(\mathbf{Y}) - M(\mathbf{X}) + z\left(1 - \frac{\alpha}{2}\right) \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}} \right]$ is a $1 - \alpha$ confidence interval for $\nu - \mu$.
2. $M(\mathbf{Y}) - M(\mathbf{X}) - z(1 - \alpha) \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}}$ is a $1 - \alpha$ confidence lower bound for $\nu - \mu$.
3. $M(\mathbf{Y}) - M(\mathbf{X}) + z(1 - \alpha) \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}}$ is a $1 - \alpha$ confidence upper bound for $\nu - \mu$.

Proof

The variable T given above has the standard normal distribution. Hence each of the following events has probability $1 - \alpha$ by definition of the quantiles:

1. $\{-z(1 - \frac{\alpha}{2}) \leq Z \leq z(1 - \frac{\alpha}{2})\}$
2. $\{Z \geq z(1 - \alpha)\}$
3. $\{Z \leq -z(1 - \alpha)\}$

In each case, solving the inequality for $\nu - \mu$ gives the result.

The two-sided interval in part (a) is the *symmetric interval* corresponding to $\alpha/2$ in both tails of the standard normal distribution. As usual, we can construct more general two-sided intervals by partitioning α between the left and right tails in anyway that we please.

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ confidence interval for $\nu - \mu$ is

$$\left[M(\mathbf{Y}) - M(\mathbf{X}) - z(1 - \alpha p) \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}}, M(\mathbf{Y}) - M(\mathbf{X}) - z(\alpha - p\alpha) \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}} \right] \quad (8.4.4)$$

1. $p = \frac{1}{2}$ gives the symmetric two-sided interval.
2. $p \rightarrow 1$ gives the interval with the confidence lower bound.
3. $p \rightarrow 0$ gives the interval with confidence upper bound.

Proof

From the distribution of the pivot variable and the definition of the quantile function,

$$\mathbb{P} \left[z(\alpha - p\alpha) < \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - (\nu - \mu)}{\sqrt{\sigma^2/m + \tau^2/n}} < z(1 - p\alpha) \right] = 1 - \alpha \quad (8.4.5)$$

Solving for $\nu - \mu$ in the inequality gives the confidence interval.

The following theorem gives some basic properties of the length of this interval.

The (deterministic) length of the general two-sided confidence interval is

$$L = [z(1 - \alpha p) - z(\alpha - p\alpha)] \sqrt{\frac{\sigma^2}{m} + \frac{\tau^2}{n}} \quad (8.4.6)$$

1. L is a decreasing function of m and a decreasing function of n .
2. L is an increasing function of σ and an increasing function of τ .
3. L is an decreasing function of α and hence an increasing function of the confidence level.
4. As a function of p , L decreases and then increases, with minimum value at $p = \frac{1}{2}$.

Part (a) means that we can make the estimate more precise by increasing either or both sample sizes. Part (b) means that the estimate becomes less precise as the variance in either distribution increases. Part (c) we have seen before. All other things being equal, we can increase the confidence level only at the expense of making the estimate less precise. Part (d) means that the symmetric, equal-tail confidence interval is the best of the two-sided intervals.

Confidence Intervals for the Difference of the Means with Unknown Variances

Our next method is a construction of confidence intervals for the difference of the means $\nu - \mu$ without needing to know the standard deviations σ and τ . However, there is a cost; we will assume that the standard deviations are the same, $\sigma = \tau$, but the common value is unknown. This assumption is reasonable if there is an inherent variability in the measurement variables that does not change even when different treatments are applied to the objects in the population. We need to recall some basic facts from our study of special properties of normal samples.

The *pooled estimate* of the common variance $\sigma^2 = \tau^2$ is

$$S^2(\mathbf{X}, \mathbf{Y}) = \frac{(m-1)S^2(\mathbf{X}) + (n-1)S^2(\mathbf{Y})}{m+n-2} \quad (8.4.7)$$

The random variable

$$T = \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - (\nu - \mu)}{S(\mathbf{X}, \mathbf{Y})\sqrt{1/m + 1/n}} \quad (8.4.8)$$

has the student t distribution with $m+n-2$ degrees of freedom

Note that $S^2(\mathbf{X}, \mathbf{Y})$ is a weighted average of the sample variances, with the degrees of freedom as the weight factors. Note also that T is a pivot variable for $\nu - \mu$ and so we can construct confidence intervals for $\nu - \mu$ in the usual way.

For $\alpha \in (0, 1)$,

1. $\left[M(\mathbf{Y}) - M(\mathbf{X}) - t_{m+n-2} \left(1 - \frac{\alpha}{2}\right) S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}}, M(\mathbf{Y}) - M(\mathbf{X}) + t_{m+n-2} \left(1 - \frac{\alpha}{2}\right) S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}} \right]$ is a $1 - \alpha$ confidence interval for $\nu - \mu$.
2. $M(\mathbf{Y}) - M(\mathbf{X}) - t_{m+n-2}(1 - \alpha)S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}}$ is a $1 - \alpha$ confidence lower bound for $\nu - \mu$.
3. $M(\mathbf{Y}) - M(\mathbf{X}) + t_{m+n-2}(1 - \alpha)S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}}$ is a $1 - \alpha$ confidence upper bound for $\nu - \mu$.

Proof

The variable T given above has the standard normal distribution. Hence each of the following events has probability $1 - \alpha$ by definition of the quantiles:

1. $\{-t_{m+n-2} \left(1 - \frac{\alpha}{2}\right) \leq T \leq t_{m+n-2} \left(1 - \frac{\alpha}{2}\right)\}$
2. $\{T \geq t_{m+n-2}(1 - \alpha)\}$
3. $\{T \leq -t_{m+n-2}(1 - \alpha)\}$

In each case, solving the inequality for $\nu - \mu$ gives the result.

The two-sided interval in part (a) is the *symmetric interval* corresponding to $\alpha/2$ in both tails of the student t distribution. As usual, we can construct more general two-sided intervals by partitioning α between the left and right tails in anyway that we please.

For every $\alpha, p \in (0, 1)$, a $1 - \alpha$ confidence interval for $\nu - \mu$ is

$$\left[M(\mathbf{Y}) - M(\mathbf{X}) - t_{m+n-2}(1 - \alpha p)S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}}, M(\mathbf{Y}) - M(\mathbf{X}) - t_{m+n-2}(\alpha - p\alpha)S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}} \right] \quad (8.4.9)$$

1. $p = \frac{1}{2}$ gives the symmetric two-sided interval.
2. $p \rightarrow 1$ gives the interval with the confidence lower bound.
3. $p \rightarrow 0$ gives the interval with confidence upper bound.

Proof

From the distribution of the pivot variable and the definition of the quantile function,

$$\mathbb{P} \left[t_{m+n-2}(\alpha - p\alpha) < \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - (\nu - \mu)}{S(\mathbf{X}, \mathbf{Y})\sqrt{1/m + 1/n}} < t_{m+n-2}(1 - p\alpha) \right] = 1 - \alpha \quad (8.4.10)$$

Solving for $\nu - \mu$ in the inequality gives the confidence interval.

The next result considers the length of the general two-sided interval.

The (random) length of the two-sided interval above is

$$L = [t_{m+n-2}(1-p\alpha) - t_{m+n-2}(\alpha-p\alpha)]S(\mathbf{X}, \mathbf{Y})\sqrt{\frac{1}{m} + \frac{1}{n}} \quad (8.4.11)$$

1. L is an decreasing function of α and hence an increasing function of the confidence level.
2. As a function of p , L decreases and then increases, with minimum value at $p = \frac{1}{2}$.

As in the case of known variances, part (c) means that all other things being equal, we can increase the confidence level only at the expense of making the estimate less precise. Part (b) means that the symmetric, equal-tail confidence interval is the best of the two-sided intervals.

Confidence Intervals for the Ratio of the Variances

Our next construction will produce interval estimates for the ratio of the variances τ^2/σ^2 (or by taking square roots, for the ratio of the standard deviations τ/σ). Once again, we need to recall some basic facts from our study of special properties of random samples from the normal distribution.

The ratio

$$U = \frac{S^2(\mathbf{X})\tau^2}{S^2(\mathbf{Y})\sigma^2} \quad (8.4.12)$$

has the F distribution with $m-1$ degrees of freedom in the numerator and $n-1$ degrees of freedom in the denominator, and hence this variable is a pivot variable for τ^2/σ^2 .

The pivot variable U can be used to construct confidence intervals for τ^2/σ^2 in the usual way.

For $\alpha \in (0, 1)$,

1. $\left[f_{m-1, n-1}\left(\frac{\alpha}{2}\right) \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})}, f_{m-1, n-1}\left(1 - \frac{\alpha}{2}\right) \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} \right]$ is a $1 - \alpha$ confidence interval for τ^2/σ^2 .
2. $f_{m-1, n-1}(1 - \alpha) \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})}$ is a $1 - \alpha$ confidence lower bound for τ^2/σ^2 .
3. $f_{m-1, n-1}(\alpha) \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})}$ is a $1 - \alpha$ confidence upper bound for $\nu - \mu$.

Proof

The variable U given above has the F distribution with $m-1$ degrees of freedom in the numerator and $n-1$ degrees of freedom in the denominator. Hence each of the following events has probability $1 - \alpha$ by definition of the quantiles:

1. $\{f_{m-1, n-1}\left(\frac{\alpha}{2}\right) \leq U \leq f_{m-1, n-1}\left(1 - \frac{\alpha}{2}\right)\}$
2. $\{U \geq f_{m-1, n-1}(1 - \alpha)\}$
3. $\{U \leq f_{m-1, n-1}(\alpha)\}$

In each case, solving the inequality for τ^2/σ^2 gives the result.

The two-sided confidence interval in part (a) is the *equal-tail* confidence interval, and is the one commonly used. But as usual, we can partition α between the left and right tails of the distribution of the pivot variable in any way that we please.

For every α , $p \in (0, 1)$, a $1 - \alpha$ confidence set for τ^2/σ^2 is

$$\left[f_{m-1, n-1}(\alpha - p\alpha) \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})}, f_{m-1, n-1}(1 - p\alpha) \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} \right] \quad (8.4.13)$$

1. $p = \frac{1}{2}$ gives the equal-tail, two-sided interval.
2. $p \rightarrow 1$ gives the interval with the confidence lower bound.
3. $p \rightarrow 0$ gives the interval with confidence upper bound.

Proof

From the F pivot variable and the definition of the quantile function,

$$\mathbb{P} \left[f_{m-1, n-1}(\alpha - p\alpha) < \frac{S^2(\mathbf{X}, \mu)\tau^2}{S^2(\mathbf{Y}, \nu)\sigma^2} < f_{m-1, n-1}(1 - p\alpha) \right] = 1 - \alpha \quad (8.4.14)$$

Solving for τ^2/σ^2 in the inequality.

The length of the general confidence interval is considered next.

The (random) length of the general two-sided confidence interval above is

$$L = [f_{m-1, n-1}(1-p\alpha) - f_{m-1, n-1}(\alpha-p\alpha)] \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} \quad (8.4.15)$$

Assuming that $m > 5$ and $n > 1$,

1. L is an decreasing function of α and hence an increasing function of the confidence level.
2. $\mathbb{E}(L) = \frac{\tau^2}{\sigma^2} \frac{m-1}{m-3}$
3. $\text{var}(L) = 2 \frac{\tau^4}{\sigma^4} \left(\frac{m-1}{m-3} \right)^2 \frac{m+n-4}{(n-1)(m-5)}$

Proof

Parts (b) and (c) follow since $\frac{\sigma^2}{\tau^2} \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})}$ as the F distribution with $n-1$ degrees of freedom in the numerator and $m-1$ degrees of freedom in the denominator.

Optimally, we might want to choose p so that $\mathbb{E}(L)$ is minimized. However, this is difficult computationally, and fortunately the equal-tail interval with $p = \frac{1}{2}$ is not too far from optimal when the sample sizes m and n are large.

Estimation in the Bivariate Normal Model

In this subsection, we consider a model that is superficially similar to the two-sample normal model, but is actually much simpler. Suppose that

$$((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n)) \quad (8.4.16)$$

is a random sample of size n from the bivariate normal distribution of a random vector (X, Y) , with $\mathbb{E}(X) = \mu$, $\mathbb{E}(Y) = \nu$, $\text{var}(X) = \sigma^2$, $\text{var}(Y) = \tau^2$, and $\text{cov}(X, Y) = \delta$.

Thus, instead of a *pair of samples*, we have a *sample of pairs*. This type of model frequently arises in *before and after experiments*, in which a measurement of interest is recorded for a sample of n objects from the population, both before and after a treatment. For example, we could record the blood pressure of a sample of n patients, before and after the administration of a certain drug. The critical point is that in this model, X_i and Y_i are measurements made on the same underlying object in the sample. As with the two-sample normal model, the interest is usually in estimating the difference of the means.

We will use our usual notation for the sample means and variances of $\mathbf{X} = (X_1, X_2, \dots, X_n)$ and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$. Recall also that the sample covariance of (\mathbf{X}, \mathbf{Y}) , is

$$S(\mathbf{X}, \mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n [X_i - M(\mathbf{X})][Y_i - M(\mathbf{Y})] \quad (8.4.17)$$

(not to be confused with the pooled estimate of the standard deviation in the two sample model).

The vector of differences $\mathbf{Y} - \mathbf{X} = (Y_1 - X_1, Y_2 - X_2, \dots, Y_n - X_n)$ is a random sample of size n from the distribution of $Y - X$, which is normal with

1. $\mathbb{E}(Y - X) = \nu - \mu$
2. $\text{var}(Y - X) = \sigma^2 + \tau^2 - 2\delta$

The sample mean and variance of the sample of differences are given by

1. $M(\mathbf{Y} - \mathbf{X}) = M(\mathbf{Y}) - M(\mathbf{X})$
2. $S^2(\mathbf{Y} - \mathbf{X}) = S^2(\mathbf{X}) + S^2(\mathbf{Y}) - 2S(\mathbf{X}, \mathbf{Y})$

Thus, the sample of differences $\mathbf{Y} - \mathbf{X}$ fits the normal model for a single variable. The section on Estimation in the Normal Model could be used to obtain confidence sets and intervals for the parameters $(\nu - \mu, \sigma^2 + \tau^2 - 2\delta)$.

In the setting of this subsection, suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ are independent. Mathematically this fits both models—the two-sample normal model and the bivariate normal model. Which procedure would work better for estimating the difference of means $\nu - \mu$?

1. If the standard deviations σ and τ are known.
2. If the standard deviations σ and τ are unknown.

Answer

1. The two methods are equivalent.
2. The bivariate normal model works better.

Although the setting in the last problem fits both models *mathematically*, only one model would make sense in a real problem. Again, the critical point is whether (X_i, Y_i) makes sense as a pair of random variables (measurements) corresponding to a given object in the sample.

Computational Exercises

A new drug is being developed to reduce a certain blood chemical. A sample of 36 patients are given a placebo while a sample of 49 patients are given the drug. Let X denote the measurement for a patient given the placebo and Y the measurement for a patient given the drug (in mg). The statistics are $m(\mathbf{x}) = 87$, $s(\mathbf{x}) = 4$, $m(\mathbf{y}) = 63$, $s(\mathbf{y}) = 6$.

1. Compute the 90% confidence interval for τ/σ .
2. Assuming that $\sigma = \tau$, compute the 90% confidence interval for $\nu - \mu$.
3. Based on (a), is the assumption that $\sigma = \tau$ reasonable?
4. Based on (b), is the drug effective?

Answer

1. (1.149, 1.936)
2. (-24.834, -23.166)
3. Perhaps not.
4. Yes

A company claims that an herbal supplement improves intelligence. A sample of 25 persons are given a standard IQ test before and after taking the supplement. Let X denote the IQ of a subject before taking the supplement and Y the IQ of the subject after the supplement. The before and after statistics are $m(\mathbf{x}) = 105$, $s(\mathbf{x}) = 13$, $m(\mathbf{y}) = 110$, $s(\mathbf{y}) = 17$, $s(\mathbf{x}, \mathbf{y}) = 190$. Do you believe the company's claim?

Answer

A 90% confidence lower bound for the difference in IQ is 2.675. There may be a vary small increase.

In Fisher's iris data, let X denote consider the petal length of a Versicolor iris and Y the petal length of a Virginica iris.

1. Compute the 90% confidence interval for τ/σ .
2. Assuming that $\sigma = \tau$, compute the 90% confidence interval for $\nu - \mu$.
3. Based on (a), is the assumption that $\sigma = \tau$ reasonable?

Answer

1. (0.8, 1.3)
2. (10.5, 14.1)
3. Yes

A plant has two machines that produce a circular rod whose diameter (in cm) is critical. Let X denote the diameter of a rod from the first machine and Y the diameter of a rod from the second machine. A sample of 100 rods from the first machine as mean 10.3 and standard deviation 1.2. A sample of 100 rods from the second machine has mean 9.8 and standard deviation 1.6.

1. Compute the 90% confidence interval for τ/σ .
2. Assuming that $\sigma = \tau$, compute the 90% confidence interval for $\nu - \mu$.
3. Based on (a), is the assumption that $\sigma = \tau$ reasonable?

Answer

1. (1.127, 1.578)
2. (0.832, 0.168)
3. Perhaps not.

This page titled [8.4: Estimation in the Two-Sample Normal Model](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

8.5: Bayesian Set Estimation

Basic Theory

As usual, our starting point is a random experiment with an underlying sample space and a probability measure \mathbb{P} . In the basic statistical model, we have an observable random variable \mathbf{X} taking values in a set S . In general, \mathbf{X} can have quite a complicated structure. For example, if the experiment is to sample n objects from a population and record various measurements of interest, then

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (8.5.1)$$

where X_i is the vector of measurements for the i th object.

Suppose also that the distribution of \mathbf{X} depends on a parameter θ taking values in a parameter space Θ . The parameter may also be vector valued, in which case $\Theta \subseteq \mathbb{R}^k$ for some $k \in \mathbb{N}_+$ and the parameter has the form $\theta = (\theta_1, \theta_2, \dots, \theta_k)$.

The Bayesian Formulation

Recall that in Bayesian analysis, the unknown parameter θ is treated as a random variable. Specifically, suppose that the conditional probability density function of the data vector \mathbf{X} given $\theta \in \Theta$ is denoted $f(\mathbf{x} | \theta)$ for $\mathbf{x} \in S$. Moreover, the parameter θ is given a *prior distribution* with probability density function h on Θ . (The prior distribution is often subjective, and is chosen to reflect our knowledge, if any, of the parameter.) The joint probability density function of the data vector and the parameter is

$$(\mathbf{x}, \theta) \mapsto h(\theta)f(\mathbf{x} | \theta); \quad (\mathbf{x}, \theta) \in S \times \Theta \quad (8.5.2)$$

Next, the (unconditional) probability density function of \mathbf{X} is the function f given by

$$f(\mathbf{x}) = \sum_{\theta \in \Theta} h(\theta)f(\mathbf{x} | \theta), \quad \mathbf{x} \in S \quad (8.5.3)$$

if the parameter has a discrete distribution, or by

$$f(\mathbf{x}) = \int_{\Theta} h(\theta)f(\mathbf{x} | \theta) d\theta, \quad \mathbf{x} \in S \quad (8.5.4)$$

if the parameter has a continuous distribution. Finally, by Bayes' theorem, the *posterior probability density function* of θ given $\mathbf{x} \in S$ is

$$h(\theta | \mathbf{x}) = \frac{h(\theta)f(\mathbf{x} | \theta)}{f(\mathbf{x})}, \quad \theta \in \Theta \quad (8.5.5)$$

In some cases, we can recognize the posterior distribution from the functional form of $\theta \mapsto h(\theta)f(\mathbf{x} | \theta)$ without having to actually compute the normalizing constant $f(\mathbf{x})$, and thus reducing the computational burden significantly. In particular, this is often the case when we have a *conjugate parametric family* of distributions of θ . Recall that this means that when the prior distribution of θ belongs to the family, so does the posterior distribution given $\mathbf{x} \in S$.

Confidence Sets

Now let $C(\mathbf{X})$ be a confidence set (that is, a subset of the parameter space that depends on the data variable \mathbf{X} but no unknown parameters). One possible definition of a $1 - \alpha$ level *Bayesian confidence set* requires that

$$\mathbb{P}[\theta \in C(\mathbf{x}) | \mathbf{X} = \mathbf{x}] = 1 - \alpha \quad (8.5.6)$$

In this definition, only θ is random and thus the probability above is computed using the posterior probability density function $\theta \mapsto h(\theta | \mathbf{x})$. Another possible definition requires that

$$\mathbb{P}[\theta \in C(\mathbf{X})] = 1 - \alpha \quad (8.5.7)$$

In this definition, \mathbf{X} and θ are both random, and so the probability above would be computed using the joint probability density function $(\mathbf{x}, \theta) \mapsto h(\theta)f(\mathbf{x} | \theta)$. Whatever the philosophical arguments may be, the first definition is certainly the easier one from a computational viewpoint, and hence is the one most commonly used.

Let us compare the classical and Bayesian approaches. In the classical approach, the parameter θ is deterministic, but unknown. *Before* the data are collected, the confidence set $C(\mathbf{X})$ (which is random by virtue of \mathbf{X}) will contain the parameter with probability $1 - \alpha$. *After* the data are collected, the computed confidence set $C(\mathbf{x})$ either contains θ or does not, and we will usually never know which. By

contrast in a Bayesian confidence set, the random parameter θ falls in the computed, deterministic confidence set $C(\mathbf{x})$ with probability $1 - \alpha$.

Real Parameters

Suppose that θ is real valued, so that $\Theta \subseteq \mathbb{R}$. For $r \in (0, 1)$, we can compute the $1 - \alpha$ level Bayesian confidence interval as $[U_{(1-r)\alpha}(\mathbf{x}), U_{1-r\alpha}(\mathbf{x})]$ where $U_p(\mathbf{x})$ is the quantile of order p for the posterior distribution of θ given $\mathbf{X} = \mathbf{x}$. As in past sections, r is the fraction of α in the right tail of the posterior distribution and $1 - r$ is the fraction of α in the left tail of the posterior distribution. As usual, $r = \frac{1}{2}$ gives the symmetric, two-sided confidence interval; letting $r \rightarrow 0$ gives the confidence lower bound; and letting $r \rightarrow 1$ gives the confidence upper bound.

Random Samples

In terms of our data vector \mathbf{X} the most important special case arises when we have a basic variable X with values in a set R , and given θ , $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from X . That is, given θ , \mathbf{X} is a sequence of independent, identically distributed variables, each with the same distribution as X given θ . Thus $S = R^n$ and if X has conditional probability density function $g(x | \theta)$, then

$$f(\mathbf{x} | \theta) = g(x_1 | \theta)g(x_2 | \theta) \cdots g(x_n | \theta), \quad \mathbf{x} = (x_1, x_2, \dots, x_n) \in S \quad (8.5.8)$$

Applications

The Bernoulli Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Bernoulli distribution with unknown success parameter $p \in (0, 1)$. In the usual language of reliability, $X_i = 1$ means success on trial i and $X_i = 0$ means failure on trial i . The distribution is named for Jacob Bernoulli. Recall that the Bernoulli distribution has probability density function (given p)

$$g(x | p) = p^x (1 - p)^{1-x}, \quad x \in \{0, 1\} \quad (8.5.9)$$

Note that the number of successes in the n trials is $Y = \sum_{i=1}^n X_i$. Given p , random variable Y has the binomial distribution with parameters n and p .

In our previous discussion of Bayesian estimation, we showed that the beta distribution is conjugate for p . Specifically, if the prior distribution of p is beta with left parameter $a > 0$ and right parameter $b > 0$, then the posterior distribution of p given \mathbf{X} is beta with left parameter $a + Y$ and right parameter $b + (n - Y)$; the left parameter is increased by the number of successes and the right parameter by the number of failure. It follows that a $1 - \alpha$ level Bayesian confidence interval for p is $[U_{\alpha/2}(y), U_{1-\alpha/2}(y)]$ where $U_r(y)$ is the quantile of order r for the posterior beta distribution. In the special case $a = b = 1$ the prior distribution is uniform on $(0, 1)$ and reflects a lack of previous knowledge about p .

Suppose that we have a coin with an unknown probability p of heads, and that we give p the uniform prior, reflecting our lack of knowledge about p . We then toss the coin 50 times, observing 30 heads.

1. Find the posterior distribution of p given the data.
2. Construct the 95% Bayesian confidence interval.
3. Construct the classical Wald confidence interval at the 95% level.

Answer

1. Beta with left parameter 31 and right parameter 21.
2. $[0.461, 0.724]$
3. $[0.464, 0.736]$

The Poisson Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the *Poisson distribution* with parameter $\lambda \in (0, \infty)$. Recall that the Poisson distribution is often used to model the number of “random points” in a region of time or space and is studied in more detail in the chapter on the Poisson Process. The distribution is named for the inimitable Simeon Poisson and given λ , has probability density function

$$g(x | \lambda) = e^{-\lambda} \frac{\lambda^x}{x!}, \quad x \in \mathbb{N} \quad (8.5.10)$$

As usual, we will denote the sum of the sample values by $Y = \sum_{i=1}^n X_i$. Given λ , random variable Y also has a Poisson distribution, but with parameter $n\lambda$.

In our previous discussion of Bayesian estimation, we showed that the gamma distribution is conjugate for λ . Specifically, if the prior distribution of λ is gamma with shape parameter $k > 0$ and rate parameter $r > 0$ (so that the scale parameter is $1/r$), then the posterior distribution of λ given \mathbf{X} is gamma with shape parameter $k + Y$ and rate parameter $r + n$. It follows that a $1 - \alpha$ level Bayesian confidence interval for λ is $[U_{\alpha/2}(y), U_{1-\alpha/2}(y)]$ where $U_p(y)$ is the quantile of order p for the posterior gamma distribution.

Consider the alpha emissions data, which we believe come from a Poisson distribution with unknown parameter λ . Suppose that *a priori*, we believe that λ is about 5, so we give λ a prior gamma distribution with shape parameter 5 and rate parameter 1. (Thus the mean is 5 and the standard deviation $\sqrt{5} = 2.236$.)

1. Find the posterior distribution of λ given the data.
2. Construct the 95% Bayesian confidence interval.
3. Construct the classical t confidence interval at the 95% level.

Answer

1. Gamma with shape parameter 10104 and rate parameter 1208.
2. (8.202, 8.528)
3. (8.324, 8.410)

The Normal Distribution

Suppose that $\mathbf{x} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the normal distribution with unknown mean $\mu \in \mathbb{R}$ and known variance $\sigma^2 \in (0, \infty)$. Of course, the normal distribution plays an especially important role in statistics, in part because of the central limit theorem. The normal distribution is widely used to model physical quantities subject to numerous small, random errors. Recall that the normal probability density function (given the parameters) is

$$g(x | \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left[-\left(\frac{x - \mu}{\sigma} \right)^2 \right], \quad x \in \mathbb{R} \quad (8.5.11)$$

We denote the sum of the sample values by $Y = \sum_{i=1}^n X_i$. Recall that Y also has a normal distribution (given μ and σ), but with mean $n\mu$ and variance $n\sigma^2$.

In our previous discussion of Bayesian estimation, we showed that the normal distribution is conjugate for μ (with σ known). Specifically, if the prior distribution of μ is normal with mean $a \in \mathbb{R}$ and standard deviation $b \in (0, \infty)$, then the posterior distribution of μ given \mathbf{X} is also normal, with

$$\mathbb{E}(\mu | \mathbf{X}) = \frac{Yb^2 + a\sigma^2}{\sigma^2 + nb^2}, \quad \text{var}(\mu | \mathbf{X}) = \frac{\sigma^2 b^2}{\sigma^2 + nb^2} \quad (8.5.12)$$

It follows that a $1 - \alpha$ level Bayesian confidence interval for μ is $[U_{\alpha/2}(y), U_{1-\alpha/2}(y)]$ where $U_p(y)$ is the quantile of order p for the posterior normal distribution. An interesting special case is when $b = \sigma$, so that the standard deviation of the prior distribution of μ is the same as the standard deviation of the sampling distribution. In this case, the posterior mean is $(Y + a)/(n + 1)$ and the posterior variance is $\sigma^2/(n + 1)$.

The length of a certain machined part is supposed to be 10 centimeters but due to imperfections in the manufacturing process, the actual length is a normally distributed with mean μ and variance σ^2 . The variance is due to inherent factors in the process, which remain fairly stable over time. From historical data, it is known that $\sigma = 0.3$. On the other hand, μ may be set by adjusting various parameters in the process and hence may change to an unknown value fairly frequently. Thus, suppose that we give μ with a prior normal distribution with mean 10 and standard deviation 0.03. A sample of 100 parts has mean 10.2.

1. Find the posterior distribution of μ given the data.
2. Construct the 95% Bayesian confidence interval.
3. Construct the classical z confidence interval at the 95% level.

Answer

1. Normal with mean 10.198 and standard deviation 0.0299.
2. (10.14, 10.26)
3. (10.14, 10.26)

The Beta Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the beta distribution with unknown left shape parameter $a \in (0, \infty)$ and right shape parameter $b = 1$. The beta distribution is widely used to model random proportions and probabilities and other variables that take values in bounded intervals. Recall that the probability density function (given a) is

$$g(x | a) = ax^{a-1}, \quad x \in (0, 1) \quad (8.5.13)$$

We denote the product of the sample values by $W = X_1 X_2 \cdots X_n$.

In our previous discussion of Bayesian estimation, we showed that the gamma distribution is conjugate for a . Specifically, if the prior distribution of a is gamma with shape parameter $k > 0$ and rate parameter $r > 0$, then the posterior distribution of a given \mathbf{X} is also gamma, with shape parameter $k + n$ and rate parameter $r + \ln(W)$. It follows that a $1 - \alpha$ level Bayesian confidence interval for a is $[U_{\alpha/2}(w), U_{1-\alpha/2}(w)]$ where $U_p(w)$ is the quantile of order p for the posterior gamma distribution. In the special case that $k = 1$, the prior distribution of a is exponential with rate parameter r .

Suppose that the resistance of an electrical component (in Ohms) has the beta distribution with unknown left parameter a and right parameter $b = 1$. We believe that a may be about 10, so we give a the prior gamma distribution with shape parameter 10 and rate parameter 1. We sample 20 components and observe the data

0.98, 0.93, 0.99, 0.89, 0.79, 0.99, 0.92, 0.97, 0.88, 0.97, 0.86, 0.84, 0.96, 0.97, 0.92, 0.90, 0.98, 0.96, 0.96, 1.00. (8.5.14)

1. Find the posterior distribution of a .
2. Construct the 95% Bayesian confidence interval for a .

Answer

1. Gamma with shape parameter 30 and rate parameter 2.424.
2. (8.349, 17.180)

The Pareto Distribution

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from the Pareto distribution with shape parameter $a \in (0, \infty)$ and scale parameter $b = 1$. The Pareto distribution is used to model certain financial variables and other variables with heavy-tailed distributions, and is named for Vilfredo Pareto. Recall that the probability density function (given a) is

$$g(x | a) = \frac{a}{x^{a+1}}, \quad x \in [1, \infty) \quad (8.5.15)$$

We denote the product of the sample values by $W = X_1 X_2 \cdots X_n$.

In our previous discussion of Bayesian estimation, we showed that the gamma distribution is conjugate for a . Specifically, if the prior distribution of a is gamma with shape parameter $k > 0$ and rate parameter $r > 0$, then the posterior distribution of a given \mathbf{X} is also gamma, with shape parameter $k + n$ and rate parameter $r + \ln(W)$. It follows that a $1 - \alpha$ level Bayesian confidence interval for a is $[U_{\alpha/2}(w), U_{1-\alpha/2}(w)]$ where $U_p(w)$ is the quantile of order p for the posterior gamma distribution. In the special case that $k = 1$, the prior distribution of a is exponential with rate parameter r .

Suppose that a financial variable has the Pareto distribution with unknown shape parameter a and scale parameter $b = 1$. We believe that a may be about 4, so we give a the prior gamma distribution with shape parameter 4 and rate parameter 1. A random sample of size 20 from the variable gives the data

1.09, 1.13, 2.00, 1.43, 1.26, 1.00, 1.36, 1.03, 1.46, 1.18, 2.16, 1.16, 1.22, 1.06, 1.28, 1.23, 1.11, 1.03, 1.04, 1.05. (8.5.16)

1. Find the posterior distribution of a .
2. Construct the 95% Bayesian confidence interval for a .

Answer

1. Gamma with shape parameter 24 and rate parameter 5.223.
2. (2.944, 6.608)

This page titled [8.5: Bayesian Set Estimation](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

9: Hypothesis Testing

Hypothesis testing refers to the process of choosing between competing hypotheses about a probability distribution, based on observed data from the distribution. It is a core topic in mathematical statistics, and indeed is a fundamental part of the language of statistics. In this chapter, we study the basics of hypothesis testing, and explore hypothesis tests in some of the most important parametric models: the normal model and the Bernoulli model.

[9.1: Introduction to Hypothesis Testing](#)

[9.2: Tests in the Normal Model](#)

[9.3: Tests in the Bernoulli Model](#)

[9.4: Tests in the Two-Sample Normal Model](#)

[9.5: Likelihood Ratio Tests](#)

[9.6: Chi-Square Tests](#)

This page titled [9: Hypothesis Testing](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.1: Introduction to Hypothesis Testing

Basic Theory

Preliminaries

As usual, our starting point is a random experiment with an underlying sample space and a probability measure \mathbb{P} . In the basic statistical model, we have an observable random variable \mathbf{X} taking values in a set S . In general, \mathbf{X} can have quite a complicated structure. For example, if the experiment is to sample n objects from a population and record various measurements of interest, then

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (9.1.1)$$

where X_i is the vector of measurements for the i th object. The most important special case occurs when (X_1, X_2, \dots, X_n) are independent and identically distributed. In this case, we have a random sample of size n from the common distribution.

The purpose of this section is to define and discuss the basic concepts of statistical *hypothesis testing*. Collectively, these concepts are sometimes referred to as the *Neyman-Pearson* framework, in honor of Jerzy Neyman and Egon Pearson, who first formalized them.

Hypotheses

A *statistical hypothesis* is a statement about the distribution of \mathbf{X} . Equivalently, a statistical hypothesis specifies a set of possible distributions of \mathbf{X} : the set of distributions for which the statement is true. A hypothesis that specifies a single distribution for \mathbf{X} is called *simple*; a hypothesis that specifies more than one distribution for \mathbf{X} is called *composite*.

In *hypothesis testing*, the goal is to see if there is sufficient statistical evidence to reject a presumed *null hypothesis* in favor of a conjectured *alternative hypothesis*. The null hypothesis is usually denoted H_0 while the alternative hypothesis is usually denoted H_1 .

An hypothesis test is a *statistical decision*; the conclusion will either be to *reject* the null hypothesis in favor of the alternative, or to *fail to reject* the null hypothesis. The decision that we make must, of course, be based on the observed value \mathbf{x} of the data vector \mathbf{X} . Thus, we will find an appropriate subset R of the sample space S and reject H_0 if and only if $\mathbf{x} \in R$. The set R is known as the *rejection region* or the *critical region*. Note the asymmetry between the null and alternative hypotheses. This asymmetry is due to the fact that we *assume* the null hypothesis, in a sense, and then see if there is sufficient evidence in \mathbf{x} to overturn this assumption in favor of the alternative.

An hypothesis test is a statistical analogy to proof by contradiction, in a sense. Suppose for a moment that H_1 is a statement in a mathematical theory and that H_0 is its negation. One way that we can prove H_1 is to assume H_0 and work our way logically to a contradiction. In an hypothesis test, we don't "prove" anything of course, but there are similarities. We assume H_0 and then see if the data \mathbf{x} are sufficiently at odds with that assumption that we feel justified in rejecting H_0 in favor of H_1 .

Often, the critical region is defined in terms of a statistic $w(\mathbf{X})$, known as a *test statistic*, where w is a function from S into another set T . We find an appropriate rejection region $R_T \subseteq T$ and reject H_0 when the observed value $w(\mathbf{x}) \in R_T$. Thus, the rejection region in S is then $R = w^{-1}(R_T) = \{\mathbf{x} \in S : w(\mathbf{x}) \in R_T\}$. As usual, the use of a statistic often allows significant *data reduction* when the dimension of the test statistic is much smaller than the dimension of the data vector.

Errors

The ultimate decision may be correct or may be in error. There are two types of errors, depending on which of the hypotheses is actually true.

Types of errors:

1. A *type 1 error* is rejecting the null hypothesis H_0 when H_0 is true.
2. A *type 2 error* is failing to reject the null hypothesis H_0 when the alternative hypothesis H_1 is true.

Similarly, there are two ways to make a *correct* decision: we could reject H_0 when H_1 is true or we could fail to reject H_0 when H_0 is true. The possibilities are summarized in the following table:

Hypothesis Test

State Decision	Fail to reject H_0	Reject H_0
H_0 True	Correct	Type 1 error
H_1 True	Type 2 error	Correct

Of course, when we observe $\mathbf{X} = \mathbf{x}$ and make our decision, either we will have made the correct decision or we will have committed an error, and usually we will never know which of these events has occurred. *Prior* to gathering the data, however, we can consider the probabilities of the various errors.

If H_0 is true (that is, the distribution of \mathbf{X} is specified by H_0), then $\mathbb{P}(\mathbf{X} \in R)$ is the probability of a type 1 error for this distribution. If H_0 is composite, then H_0 specifies a variety of different distributions for \mathbf{X} and thus there is a set of type 1 error probabilities.

The maximum probability of a type 1 error, over the set of distributions specified by H_0 , is the *significance level* of the test or the *size* of the critical region.

The significance level is often denoted by α . Usually, the rejection region is constructed so that the significance level is a prescribed, small value (typically 0.1, 0.05, 0.01).

If H_1 is true (that is, the distribution of \mathbf{X} is specified by H_1), then $\mathbb{P}(\mathbf{X} \notin R)$ is the probability of a type 2 error for this distribution. Again, if H_1 is composite then H_1 specifies a variety of different distributions for \mathbf{X} , and thus there will be a set of type 2 error probabilities. Generally, there is a tradeoff between the type 1 and type 2 error probabilities. If we reduce the probability of a type 1 error, by making the rejection region R smaller, we necessarily increase the probability of a type 2 error because the complementary region $S \setminus R$ is larger.

The extreme cases can give us some insight. First consider the decision rule in which we *never* reject H_0 , regardless of the evidence \mathbf{x} . This corresponds to the rejection region $R = \emptyset$. A type 1 error is impossible, so the significance level is 0. On the other hand, the probability of a type 2 error is 1 for any distribution defined by H_1 . At the other extreme, consider the decision rule in which we always rejects H_0 regardless of the evidence \mathbf{x} . This corresponds to the rejection region $R = S$. A type 2 error is impossible, but now the probability of a type 1 error is 1 for any distribution defined by H_0 . In between these two worthless tests are meaningful tests that take the evidence \mathbf{x} into account.

Power

If H_1 is true, so that the distribution of \mathbf{X} is specified by H_1 , then $\mathbb{P}(\mathbf{X} \in R)$, the probability of rejecting H_0 is the *power* of the test for that distribution.

Thus the power of the test for a distribution specified by H_1 is the probability of making the correct decision.

Suppose that we have two tests, corresponding to rejection regions R_1 and R_2 , respectively, each having significance level α . The test with region R_1 is *uniformly more powerful* than the test with region R_2 if

$$\mathbb{P}(\mathbf{X} \in R_1) \geq \mathbb{P}(\mathbf{X} \in R_2) \text{ for every distribution of } \mathbf{X} \text{ specified by } H_1 \quad (9.1.2)$$

Naturally, in this case, we would prefer the first test. Often, however, two tests will not be uniformly ordered; one test will be more powerful for some distributions specified by H_1 while the other test will be more powerful for other distributions specified by H_1 .

If a test has significance level α and is uniformly more powerful than any other test with significance level α , then the test is said to be a *uniformly most powerful test* at level α .

Clearly a uniformly most powerful test is the best we can do.

P-value

In most cases, we have a general procedure that allows us to construct a test (that is, a rejection region R_α) for any given significance level $\alpha \in (0, 1)$. Typically, R_α decreases (in the subset sense) as α decreases.

The P -value of the observed value \mathbf{x} of \mathbf{X} , denoted $P(\mathbf{x})$, is defined to be the smallest α for which $\mathbf{x} \in R_\alpha$; that is, the smallest significance level for which H_0 is rejected, given $\mathbf{X} = \mathbf{x}$.

Knowing $P(\mathbf{x})$ allows us to test H_0 at any significance level for the given data \mathbf{x} : If $P(\mathbf{x}) \leq \alpha$ then we would reject H_0 at significance level α ; if $P(\mathbf{x}) > \alpha$ then we fail to reject H_0 at significance level α . Note that $P(\mathbf{X})$ is a *statistic*. Informally, $P(\mathbf{x})$ can often be thought of as the probability of an outcome “as or more extreme” than the observed value \mathbf{x} , where *extreme* is interpreted relative to the null hypothesis H_0 .

Analogy with Justice Systems

There is a helpful analogy between statistical hypothesis testing and the criminal justice system in the US and various other countries. Consider a person charged with a crime. The presumed *null hypothesis* is that the person is innocent of the crime; the conjectured *alternative hypothesis* is that the person is guilty of the crime. The test of the hypotheses is a trial with evidence presented by both sides playing the role of the data. After considering the evidence, the jury delivers the decision as either *not guilty* or *guilty*. Note that *innocent* is not a possible verdict of the jury, because it is not the point of the trial to *prove* the person innocent. Rather, the point of the trial is to see whether there is sufficient evidence to overturn the null hypothesis that the person is innocent in favor of the alternative hypothesis of that the person is guilty. A *type 1 error* is convicting a person who is innocent; a *type 2 error* is acquitting a person who is guilty. Generally, a type 1 error is considered the more serious of the two possible errors, so in an attempt to hold the chance of a type 1 error to a very low level, the standard for conviction in serious criminal cases is *beyond a reasonable doubt*.

Tests of an Unknown Parameter

Hypothesis testing is a very general concept, but an important special class occurs when the distribution of the data variable \mathbf{X} depends on a parameter θ taking values in a parameter space Θ . The parameter may be vector-valued, so that $\theta = (\theta_1, \theta_2, \dots, \theta_n)$ and $\Theta \subseteq \mathbb{R}^k$ for some $k \in \mathbb{N}_+$. The hypotheses generally take the form

$$H_0 : \theta \in \Theta_0 \text{ versus } H_1 : \theta \notin \Theta_0 \quad (9.1.3)$$

where Θ_0 is a prescribed subset of the parameter space Θ . In this setting, the probabilities of making an error or a correct decision depend on the true value of θ . If R is the rejection region, then the *power function* Q is given by

$$Q(\theta) = \mathbb{P}_\theta(\mathbf{X} \in R), \quad \theta \in \Theta \quad (9.1.4)$$

The power function gives a lot of information about the test.

The power function satisfies the following properties:

1. $Q(\theta)$ is the probability of a type 1 error when $\theta \in \Theta_0$.
2. $\max \{Q(\theta) : \theta \in \Theta_0\}$ is the significance level of the test.
3. $1 - Q(\theta)$ is the probability of a type 2 error when $\theta \notin \Theta_0$.
4. $Q(\theta)$ is the power of the test when $\theta \notin \Theta_0$.

If we have two tests, we can compare them by means of their power functions.

Suppose that we have two tests, corresponding to rejection regions R_1 and R_2 , respectively, each having significance level α . The test with rejection region R_1 is uniformly more powerful than the test with rejection region R_2 if $Q_1(\theta) \geq Q_2(\theta)$ for all $\theta \notin \Theta_0$.

Most hypothesis tests of an unknown real parameter θ fall into three special cases:

Suppose that θ is a real parameter and $\theta_0 \in \Theta$ a specified value. The tests below are respectively the *two-sided test*, the *left-tailed test*, and the *right-tailed test*.

1. $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$
2. $H_0 : \theta \geq \theta_0$ versus $H_1 : \theta < \theta_0$
3. $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$

Thus the tests are named after the conjectured alternative. Of course, there may be other unknown parameters besides θ (known as *nuisance parameters*).

Equivalence Between Hypothesis Test and Confidence Sets

There is an equivalence between hypothesis tests and confidence sets for a parameter θ .

Suppose that $C(\mathbf{x})$ is a $1 - \alpha$ level confidence set for θ . The following test has significance level α for the hypothesis $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$: Reject H_0 if and only if $\theta_0 \notin C(\mathbf{x})$

Proof

By definition, $\mathbb{P}[\theta \in C(\mathbf{X})] = 1 - \alpha$. Hence if H_0 is true so that $\theta = \theta_0$, then the probability of a type 1 error is $P[\theta \notin C(\mathbf{X})] = \alpha$.

Equivalently, we *fail* to reject H_0 at significance level α if and only if θ_0 is in the corresponding $1 - \alpha$ level confidence set. In particular, this equivalence applies to interval estimates of a real parameter θ and the common tests for θ given [above](#).

In each case below, the confidence interval has confidence level $1 - \alpha$ and the test has significance level α .

1. Suppose that $[L(\mathbf{X}), U(\mathbf{X})]$ is a two-sided confidence interval for θ . Reject $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$ if and only if $\theta_0 < L(\mathbf{X})$ or $\theta_0 > U(\mathbf{X})$.
2. Suppose that $L(\mathbf{X})$ is a confidence lower bound for θ . Reject $H_0 : \theta \leq \theta_0$ versus $H_1 : \theta > \theta_0$ if and only if $\theta_0 < L(\mathbf{X})$.
3. Suppose that $U(\mathbf{X})$ is a confidence upper bound for θ . Reject $H_0 : \theta \geq \theta_0$ versus $H_1 : \theta < \theta_0$ if and only if $\theta_0 > U(\mathbf{X})$.

Pivot Variables and Test Statistics

Recall that confidence sets of an unknown parameter θ are often constructed through a *pivot variable*, that is, a random variable $W(\mathbf{X}, \theta)$ that depends on the data vector \mathbf{X} and the parameter θ , but whose distribution does not depend on θ and is known. In this case, a natural test statistic for the [basic tests](#) given above is $W(\mathbf{X}, \theta_0)$.

This page titled [9.1: Introduction to Hypothesis Testing](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.2: Tests in the Normal Model

Basic Theory

The Normal Model

The normal distribution is perhaps the most important distribution in the study of mathematical statistics, in part because of the central limit theorem. As a consequence of this theorem, a measured quantity that is subject to numerous small, random errors will have, at least approximately, a normal distribution. Such variables are ubiquitous in statistical experiments, in subjects varying from the physical and biological sciences to the social sciences.

So in this section, we assume that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the normal distribution with mean μ and standard deviation σ . Our goal in this section is to construct hypothesis tests for μ and σ ; these are among of the most important special cases of hypothesis testing. This section parallels the section on Estimation in the Normal Model in the chapter on Set Estimation, and in particular, the duality between interval estimation and hypothesis testing will play an important role. But first we need to review some basic facts that will be critical for our analysis.

Recall that the sample mean M and sample variance S^2 are

$$M = \frac{1}{n} \sum_{i=1}^n X_i, \quad S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - M)^2 \quad (9.2.1)$$

From our study of point estimation, recall that M is an unbiased and consistent estimator of μ while S^2 is an unbiased and consistent estimator of σ^2 . From these basic statistics we can construct the test statistics that will be used to construct our hypothesis tests. The following results were established in the section on Special Properties of the Normal Distribution.

Define

$$Z = \frac{M - \mu}{\sigma/\sqrt{n}}, \quad T = \frac{M - \mu}{S/\sqrt{n}}, \quad V = \frac{n-1}{\sigma^2} S^2 \quad (9.2.2)$$

1. Z has the standard normal distribution.
2. T has the student t distribution with $n-1$ degrees of freedom.
3. V has the chi-square distribution with $n-1$ degrees of freedom.
4. Z and V are independent.

It follows that each of these random variables is a pivot variable for (μ, σ) since the distributions do not depend on the parameters, but the variables themselves functionally depend on one or both parameters. The pivot variables will lead to natural test statistics that can then be used to perform the hypothesis tests of the parameters. To construct our tests, we will need quantiles of these standard distributions. The quantiles can be computed using the special distribution calculator or from most mathematical and statistical software packages. Here is the notation we will use:

Let $p \in (0, 1)$ and $k \in \mathbb{N}_+$.

1. $z(p)$ denotes the quantile of order p for the standard normal distribution.
2. $t_k(p)$ denotes the quantile of order p for the student t distribution with k degrees of freedom.
3. $\chi_k^2(p)$ denotes the quantile of order p for the chi-square distribution with k degrees of freedom

Since the standard normal and student t distributions are symmetric about 0, it follows that $z(1-p) = -z(p)$ and $t_k(1-p) = -t_k(p)$ for $p \in (0, 1)$ and $k \in \mathbb{N}_+$. On the other hand, the chi-square distribution is not symmetric.

Tests for the Mean with Known Standard Deviation

For our first discussion, we assume that the distribution mean μ is unknown but the standard deviation σ is known. This is not always an artificial assumption. There are often situations where σ is stable over time, and hence is at least approximately known, while μ changes because of different “treatments”. Examples are given in the computational exercises below.

For a conjectured $\mu_0 \in \mathbb{R}$, define the test statistic

$$Z = \frac{M - \mu_0}{\sigma/\sqrt{n}} \quad (9.2.3)$$

1. If $\mu = \mu_0$ then Z has the standard normal distribution.
2. If $\mu \neq \mu_0$ then Z has the normal distribution with mean $\frac{\mu - \mu_0}{\sigma/\sqrt{n}}$ and variance 1.

So in case (b), $\frac{\mu - \mu_0}{\sigma/\sqrt{n}}$ can be viewed as a *non-centrality parameter*. The graph of the probability density function of Z is like that of the standard normal probability density function, but shifted to the right or left by the non-centrality parameter, depending on whether $\mu > \mu_0$ or $\mu < \mu_0$.

For $\alpha \in (0, 1)$, each of the following tests has significance level α :

1. Reject $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$ if and only if $Z < -z(1 - \alpha/2)$ or $Z > z(1 - \alpha/2)$ if and only if $M < \mu_0 - z(1 - \alpha/2) \frac{\sigma}{\sqrt{n}}$ or $M > \mu_0 + z(1 - \alpha/2) \frac{\sigma}{\sqrt{n}}$.
2. Reject $H_0 : \mu \leq \mu_0$ versus $H_1 : \mu > \mu_0$ if and only if $Z > z(1 - \alpha)$ if and only if $M > \mu_0 + z(1 - \alpha) \frac{\sigma}{\sqrt{n}}$.
3. Reject $H_0 : \mu \geq \mu_0$ versus $H_1 : \mu < \mu_0$ if and only if $Z < -z(1 - \alpha)$ if and only if $M < \mu_0 - z(1 - \alpha) \frac{\sigma}{\sqrt{n}}$.

Proof

In part (a), H_0 is a simple hypothesis, and under H_0 , Z has the standard normal distribution. So α is probability of falsely rejecting H_0 by definition of the quantiles. In parts (b) and (c), Z has a non-central normal distribution under H_0 as discussed [above](#). So if H_0 is true, the the maximum type 1 error probability α occurs when $\mu = \mu_0$. The decision rules in terms of M are equivalent to the corresponding ones in terms of Z by simple algebra.

Part (a) is the standard two-sided test, while (b) is the right-tailed test and (c) is the left-tailed test. Note that in each case, the hypothesis test is the dual of the corresponding interval estimate constructed in the section on Estimation in the Normal Model.

For each of the tests [above](#), we *fail* to reject H_0 at significance level α if and only if μ_0 is in the corresponding $1 - \alpha$ confidence interval, that is

1. $M - z(1 - \alpha/2) \frac{\sigma}{\sqrt{n}} \leq \mu_0 \leq M + z(1 - \alpha/2) \frac{\sigma}{\sqrt{n}}$
2. $\mu_0 \leq M + z(1 - \alpha) \frac{\sigma}{\sqrt{n}}$
3. $\mu_0 \geq M - z(1 - \alpha) \frac{\sigma}{\sqrt{n}}$

Proof

This follows from the [previous result](#). In each case, we start with the inequality that corresponds to not rejecting H_0 and solve for μ_0 .

The two-sided test in (a) corresponds to $\alpha/2$ in each tail of the distribution of the test statistic Z , under H_0 . This set is said to be *unbiased*. But of course we can construct other biased tests by partitioning the confidence level α between the left and right tails in a non-symmetric way.

For every $\alpha, p \in (0, 1)$, the following test has significance level α : Reject $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$ if and only if $Z < z(\alpha - p\alpha)$ or $Z \geq z(1 - p\alpha)$.

1. $p = \frac{1}{2}$ gives the symmetric, unbiased test.
2. $p \downarrow 0$ gives the left-tailed test.
3. $p \uparrow 1$ gives the right-tailed test.

Proof

As before H_0 is a simple hypothesis, and if H_0 is true, Z has the standard normal distribution. So the probability of falsely rejecting H_0 is α by definition of the quantiles. Parts (a)–(c) follow from properties of the standard normal quantile function.

The P -value of these test can be computed in terms of the standard normal distribution function Φ .

The P -values of the standard tests [above](#) are respectively

1. $2[1 - \Phi(|Z|)]$
2. $1 - \Phi(Z)$
3. $\Phi(Z)$

Recall that the *power function* of a test of a parameter is the probability of rejecting the null hypothesis, as a function of the true value of the parameter. Our next series of results will explore the power functions of the tests [above](#).

The power function of the general two-sided test [above](#) is given by

$$Q(\mu) = \Phi\left(z(\alpha - p\alpha) - \frac{\sqrt{n}}{\sigma}(\mu - \mu_0)\right) + \Phi\left(\frac{\sqrt{n}}{\sigma}(\mu - \mu_0) - z(1 - p\alpha)\right), \quad \mu \in \mathbb{R} \quad (9.2.4)$$

1. Q is decreasing on $(-\infty, m_0)$ and increasing on (m_0, ∞) where $m_0 = \mu_0 + [z(\alpha - p\alpha) + z(1 - p\alpha)] \frac{\sqrt{n}}{2\sigma}$.
2. $Q(\mu_0) = \alpha$.
3. $Q(\mu) \rightarrow 1$ as $\mu \uparrow \infty$ and $Q(\mu) \rightarrow 1$ as $\mu \downarrow -\infty$.
4. If $p = \frac{1}{2}$ then Q is symmetric about μ_0 (and $m_0 = \mu_0$).
5. As p increases, $Q(\mu)$ increases if $\mu > \mu_0$ and decreases if $\mu < \mu_0$.

So by varying p , we can make the test more powerful for some values of μ , but only at the expense of making the test less powerful for other values of μ .

The power function of the left-tailed test [above](#) is given by

$$Q(\mu) = \Phi\left(z(\alpha) + \frac{\sqrt{n}}{\sigma}(\mu - \mu_0)\right), \quad \mu \in \mathbb{R} \quad (9.2.5)$$

1. Q is increasing on \mathbb{R} .
2. $Q(\mu_0) = \alpha$.
3. $Q(\mu) \rightarrow 1$ as $\mu \uparrow \infty$ and $Q(\mu) \rightarrow 0$ as $\mu \downarrow -\infty$.

The power function of the right-tailed test [above](#), is given by

$$Q(\mu) = \Phi\left(z(\alpha) - \frac{\sqrt{n}}{\sigma}(\mu - \mu_0)\right), \quad \mu \in \mathbb{R} \quad (9.2.6)$$

1. Q is decreasing on \mathbb{R} .
2. $Q(\mu_0) = \alpha$.
3. $Q(\mu) \rightarrow 0$ as $\mu \uparrow \infty$ and $Q(\mu) \rightarrow 1$ as $\mu \downarrow -\infty$.

For any of the three tests in [above](#), increasing the sample size n or decreasing the standard deviation σ results in a uniformly more powerful test.

In the mean test experiment, select the normal test statistic and select the normal sampling distribution with standard deviation $\sigma = 2$, significance level $\alpha = 0.1$, sample size $n = 20$, and $\mu_0 = 0$. Run the experiment 1000 times for several values of the true distribution mean μ . For each value of μ , note the relative frequency of the event that the null hypothesis is rejected. Sketch the empirical power function.

In the mean estimate experiment, select the normal pivot variable and select the normal distribution with $\mu = 0$ and standard deviation $\sigma = 2$, confidence level $1 - \alpha = 0.90$, and sample size $n = 10$. For each of the three types of confidence intervals, run the experiment 20 times. State the corresponding hypotheses and significance level, and for each run, give the set of μ_0 for which the null hypothesis would be rejected.

In many cases, the first step is to *design* the experiment so that the significance level is α and so that the test has a given power β for a given alternative μ_1 .

For either of the one-sided tests in [above](#), the sample size n needed for a test with significance level α and power β for the alternative μ_1 is

$$n = \left(\frac{\sigma [z(\beta) - z(\alpha)]}{\mu_1 - \mu_0} \right)^2 \quad (9.2.7)$$

Proof

This follows from setting the power function equal to β and solving for n

For the unbiased, two-sided test, the sample size n needed for a test with significance level α and power β for the alternative μ_1 is approximately

$$n = \left(\frac{\sigma [z(\beta) - z(\alpha/2)]}{\mu_1 - \mu_0} \right)^2 \quad (9.2.8)$$

Proof

In the power function for the two-sided test given [above](#), we can neglect the first term if $\mu_1 < \mu_0$ and neglect the second term if $\mu_1 > \mu_0$.

Tests of the Mean with Unknown Standard Deviation

For our next discussion, we construct tests of μ without requiring the assumption that σ is known. And in applications of course, σ is usually unknown.

For a conjectured $\mu_0 \in \mathbb{R}$, define the test statistic

$$T = \frac{M - \mu_0}{S/\sqrt{n}} \quad (9.2.9)$$

1. If $\mu = \mu_0$, the statistic T has the student t distribution with $n - 1$ degrees of freedom.
2. If $\mu \neq \mu_0$ then T has a non-central t distribution with $n - 1$ degrees of freedom and non-centrality parameter $\frac{\mu - \mu_0}{\sigma/\sqrt{n}}$.

In case (b), the graph of the probability density function of T is much (but not exactly) the same as that of the ordinary t distribution with $n - 1$ degrees of freedom, but shifted to the right or left by the non-centrality parameter, depending on whether $\mu > \mu_0$ or $\mu < \mu_0$.

For $\alpha \in (0, 1)$, each of the following tests has significance level α :

1. Reject $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$ if and only if $T < -t_{n-1}(1 - \alpha/2)$ or $T > t_{n-1}(1 - \alpha/2)$ if and only if $M < \mu_0 - t_{n-1}(1 - \alpha/2)\frac{S}{\sqrt{n}}$ or $T > \mu_0 + t_{n-1}(1 - \alpha/2)\frac{S}{\sqrt{n}}$.
2. Reject $H_0 : \mu \leq \mu_0$ versus $H_1 : \mu > \mu_0$ if and only if $T > t_{n-1}(1 - \alpha)$ if and only if $M > \mu_0 + t_{n-1}(1 - \alpha)\frac{S}{\sqrt{n}}$.
3. Reject $H_0 : \mu \geq \mu_0$ versus $H_1 : \mu < \mu_0$ if and only if $T < -t_{n-1}(1 - \alpha)$ if and only if $M < \mu_0 - t_{n-1}(1 - \alpha)\frac{S}{\sqrt{n}}$.

Proof

In part (a), T has the chi-square distribution with $n - 1$ degrees of freedom under H_0 . So if H_0 is true, the probability of falsely rejecting H_0 is α by definition of the quantiles. In parts (b) and (c), T has a non-central t distribution with $n - 1$ degrees of freedom under H_0 , as discussed [above](#). Hence if H_0 is true, the maximum type 1 error probability α occurs when $\mu = \mu_0$. The decision rules in terms of M are equivalent to the corresponding ones in terms of T by simple algebra.

Part (a) is the standard two-sided test, while (b) is the right-tailed test and (c) is the left-tailed test. Note that in each case, the hypothesis test is the dual of the corresponding interval estimate constructed in the section on Estimation in the Normal Model.

For each of the tests [above](#), we fail to reject H_0 at significance level α if and only if μ_0 is in the corresponding $1 - \alpha$ confidence interval.

1. $M - t_{n-1}(1 - \alpha/2)\frac{S}{\sqrt{n}} \leq \mu_0 \leq M + t_{n-1}(1 - \alpha/2)\frac{S}{\sqrt{n}}$

2. $\mu_0 \leq M + t_{n-1}(1 - \alpha) \frac{S}{\sqrt{n}}$
3. $\mu_0 \geq M - t_{n-1}(1 - \alpha) \frac{S}{\sqrt{n}}$

Proof

This follows from the [previous result](#). In each case, we start with the inequality that corresponds to *not* rejecting H_0 and then solve for μ_0 .

The two-sided test in (a) corresponds to $\alpha/2$ in each tail of the distribution of the test statistic T , under H_0 . This set is said to be *unbiased*. But of course we can construct other biased tests by partitioning the confidence level α between the left and right tails in a non-symmetric way.

For every $\alpha, p \in (0, 1)$, the following test has significance level α : Reject $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$ if and only if $T < t_{n-1}(\alpha - p\alpha)$ or $T \geq t_{n-1}(1 - p\alpha)$ if and only if $M < \mu_0 + t_{n-1}(\alpha - p\alpha) \frac{S}{\sqrt{n}}$ or $M > \mu_0 + t_{n-1}(1 - p\alpha) \frac{S}{\sqrt{n}}$.

1. $p = \frac{1}{2}$ gives the symmetric, unbiased test.
2. $p \downarrow 0$ gives the left-tailed test.
3. $p \uparrow 1$ gives the right-tailed test.

Proof

Once again, H_0 is a simple hypothesis, and under H_0 the test statistic T has the student t distribution with $n - 1$ degrees of freedom. So if H_0 is true, the probability of falsely rejecting H_0 is α by definition of the quantiles. Parts (a)–(c) follow from properties of the quantile function.

The P -value of these test can be computed in terms of the distribution function Φ_{n-1} of the t -distribution with $n - 1$ degrees of freedom.

The P -values of the standard tests [above](#) are respectively

1. $2[1 - \Phi_{n-1}(|T|)]$
2. $1 - \Phi_{n-1}(T)$
3. $\Phi_{n-1}(T)$

In the mean test experiment, select the student test statistic and select the normal sampling distribution with standard deviation $\sigma = 2$, significance level $\alpha = 0.1$, sample size $n = 20$, and $\mu_0 = 1$. Run the experiment 1000 times for several values of the true distribution mean μ . For each value of μ , note the relative frequency of the event that the null hypothesis is rejected. Sketch the empirical power function.

In the mean estimate experiment, select the student pivot variable and select the normal sampling distribution with mean 0 and standard deviation 2. Select confidence level 0.90 and sample size 10. For each of the three types of intervals, run the experiment 20 times. State the corresponding hypotheses and significance level, and for each run, give the set of μ_0 for which the null hypothesis would be rejected.

The power function for the t tests [above](#) can be computed explicitly in terms of the non-central t distribution function. Qualitatively, the graphs of the power functions are similar to the case when σ is known, given above [two-sided](#), [left-tailed](#), and [right-tailed](#) cases.

If an upper bound σ_0 on the standard deviation σ is known, then conservative estimates on the sample size needed for a given confidence level and a given margin of error can be obtained using the methods for the normal pivot variable, in the [two-sided](#) and [one-sided](#) cases.

Tests of the Standard Deviation

For our next discussion, we will construct hypothesis tests for the distribution standard deviation σ . So our assumption is that σ is unknown, and of course almost always, μ would be unknown as well.

For a conjectured value $\sigma_0 \in (0, \infty)$, define the test statistic

$$V = \frac{n-1}{\sigma_0^2} S^2 \quad (9.2.10)$$

1. If $\sigma = \sigma_0$, then V has the chi-square distribution with $n - 1$ degrees of freedom.
2. If $\sigma \neq \sigma_0$ then V has the gamma distribution with shape parameter $(n - 1)/2$ and scale parameter $2\sigma^2/\sigma_0^2$.

Recall that the ordinary chi-square distribution with $n - 1$ degrees of freedom is the gamma distribution with shape parameter $(n - 1)/2$ and scale parameter $\frac{1}{2}$. So in case (b), the ordinary chi-square distribution is scaled by σ^2/σ_0^2 . In particular, the scale factor is greater than 1 if $\sigma > \sigma_0$ and less than 1 if $\sigma < \sigma_0$.

For every $\alpha \in (0, 1)$, the following test has significance level α :

1. Reject $H_0 : \sigma = \sigma_0$ versus $H_1 : \sigma \neq \sigma_0$ if and only if $V < \chi_{n-1}^2(\alpha/2)$ or $V > \chi_{n-1}^2(1 - \alpha/2)$ if and only if $S^2 < \chi_{n-1}^2(\alpha/2) \frac{\sigma_0^2}{n-1}$ or $S^2 > \chi_{n-1}^2(1 - \alpha/2) \frac{\sigma_0^2}{n-1}$
2. Reject $H_0 : \sigma \geq \sigma_0$ versus $H_1 : \sigma < \sigma_0$ if and only if $V < \chi_{n-1}^2(\alpha)$ if and only if $S^2 < \chi_{n-1}^2(\alpha) \frac{\sigma_0^2}{n-1}$
3. Reject $H_0 : \sigma \leq \sigma_0$ versus $H_1 : \sigma > \sigma_0$ if and only if $V > \chi_{n-1}^2(1 - \alpha)$ if and only if $S^2 > \chi_{n-1}^2(1 - \alpha) \frac{\sigma_0^2}{n-1}$

Proof

The logic is largely the same as with our other hypothesis test. In part (a), H_0 is a simple hypothesis, and under H_0 , the test statistic V has the chi-square distribution with $n - 1$ degrees of freedom. So if H_0 is true, the probability of falsely rejecting H_0 is α by definition of the quantiles. In parts (b) and (c), V has the more general gamma distribution under H_0 , as discussed [above](#). If H_0 is true, the maximum type 1 error probability is α and occurs when $\sigma = \sigma_0$.

Part (a) is the unbiased, two-sided test that corresponds to $\alpha/2$ in each tail of the chi-square distribution of the test statistic V , under H_0 . Part (b) is the left-tailed test and part (c) is the right-tailed test. Once again, we have a duality between the hypothesis tests and the interval estimates constructed in the section on Estimation in the Normal Model.

For each of the tests in [above](#), we *fail* to reject H_0 at significance level α if and only if σ_0^2 is in the corresponding $1 - \alpha$ confidence interval. That is

1. $\frac{n-1}{\chi_{n-1}^2(1-\alpha/2)} S^2 \leq \sigma_0^2 \leq \frac{n-1}{\chi_{n-1}^2(\alpha/2)} S^2$
2. $\sigma_0^2 \leq \frac{n-1}{\chi_{n-1}^2(\alpha)} S^2$
3. $\sigma_0^2 \geq \frac{n-1}{\chi_{n-1}^2(1-\alpha)} S^2$

Proof

This follows from the [previous result](#). In each case, we start with the inequality that corresponds to *not* rejecting H_0 and then solve for σ_0^2 .

As before, we can construct more general two-sided tests by partitioning the significance level α between the left and right tails of the chi-square distribution in an arbitrary way.

For every $\alpha, p \in (0, 1)$, the following test has significance level α : Reject $H_0 : \sigma = \sigma_0$ versus $H_1 : \sigma \neq \sigma_0$ if and only if $V \leq \chi_{n-1}^2(\alpha - p\alpha)$ or $V \geq \chi_{n-1}^2(1 - p\alpha)$ if and only if $S^2 < \chi_{n-1}^2(\alpha - p\alpha) \frac{\sigma_0^2}{n-1}$ or $S^2 > \chi_{n-1}^2(1 - p\alpha) \frac{\sigma_0^2}{n-1}$.

1. $p = \frac{1}{2}$ gives the equal-tail test.
2. $p \downarrow 0$ gives the left-tail test.
3. $p \uparrow 1$ gives the right-tail test.

Proof

As before, H_0 is a simple hypothesis, and under H_0 the test statistic V has the chi-square distribution with $n - 1$ degrees of freedom. So if H_0 is true, the probability of falsely rejecting H_0 is α by definition of the quantiles. Parts (a)–(c) follow from properties of the quantile function.

Recall again that the *power function* of a test of a parameter is the probability of rejecting the null hypothesis, as a function of the true value of the parameter. The power functions of the tests for σ can be expressed in terms of the distribution function G_{n-1} of the chi-square distribution with $n - 1$ degrees of freedom.

The power function of the general two-sided test [above](#) is given by the following formula, and satisfies the given properties:

$$Q(\sigma) = 1 - G_{n-1} \left(\frac{\sigma_0^2}{\sigma^2} \chi_{n-1}^2 (1 - p \alpha) \right) + G_{n-1} \left(\frac{\sigma_0^2}{\sigma^2} \chi_{n-1}^2 (\alpha - p \alpha) \right) \quad (9.2.11)$$

1. Q is decreasing on $(-\infty, \sigma_0)$ and increasing on (σ_0, ∞) .
2. $Q(\sigma_0) = \alpha$.
3. $Q(\sigma) \rightarrow 1$ as $\sigma \uparrow \infty$ and $Q(\sigma) \rightarrow 1$ as $\sigma \downarrow 0$.

The power function of the left-tailed test in [above](#) is given by the following formula, and satisfies the given properties:

$$Q(\sigma) = 1 - G_{n-1} \left(\frac{\sigma_0^2}{\sigma^2} \chi_{n-1}^2 (1 - \alpha) \right) \quad (9.2.12)$$

1. Q is increasing on $(0, \infty)$.
2. $Q(\sigma_0) = \alpha$.
3. $Q(\sigma) \rightarrow 1$ as $\sigma \uparrow \infty$ and $Q(\sigma) \rightarrow 0$ as $\sigma \downarrow 0$.

The power function for the right-tailed test [above](#) is given by the following formula, and satisfies the given properties:

$$Q(\sigma) = G_{n-1} \left(\frac{\sigma_0^2}{\sigma^2} \chi_{n-1}^2 (\alpha) \right) \quad (9.2.13)$$

1. Q is decreasing on $(0, \infty)$.
2. $Q(\sigma_0) = \alpha$.
3. $Q(\sigma) \rightarrow 0$ as $\sigma \uparrow \infty$ and $Q(\sigma) \rightarrow 0$ as $\sigma \uparrow \infty$ and as $\sigma \downarrow 0$.

In the variance test experiment, select the normal distribution with mean 0, and select significance level 0.1, sample size 10, and test standard deviation 1.0. For various values of the true standard deviation, run the simulation 1000 times. Record the relative frequency of rejecting the null hypothesis and plot the empirical power curve.

1. Two-sided test
2. Left-tailed test
3. Right-tailed test

In the variance estimate experiment, select the normal distribution with mean 0 and standard deviation 2, and select confidence level 0.90 and sample size 10. Run the experiment 20 times. State the corresponding hypotheses and significance level, and for each run, give the set of test standard deviations for which the null hypothesis would be rejected.

1. Two-sided confidence interval
2. Confidence lower bound
3. Confidence upper bound

Exercises

Robustness

The primary assumption that we made is that the underlying sampling distribution is normal. Of course, in real statistical problems, we are unlikely to know much about the sampling distribution, let alone whether or not it is normal. Suppose in fact that the underlying distribution is not normal. When the sample size n is relatively large, the distribution of the sample mean will still be approximately normal by the central limit theorem, and thus our tests of the mean μ should still be approximately valid. On the

other hand, tests of the variance σ^2 are less robust to deviations from the assumption of normality. The following exercises explore these ideas.

In the mean test experiment, select the gamma distribution with shape parameter 1 and scale parameter 1. For the three different tests and for various significance levels, sample sizes, and values of μ_0 , run the experiment 1000 times. For each configuration, note the relative frequency of rejecting H_0 . When H_0 is true, compare the relative frequency with the significance level.

In the mean test experiment, select the uniform distribution on $[0, 4]$. For the three different tests and for various significance levels, sample sizes, and values of μ_0 , run the experiment 1000 times. For each configuration, note the relative frequency of rejecting H_0 . When H_0 is true, compare the relative frequency with the significance level.

How large n needs to be for the testing procedure to work well depends, of course, on the underlying distribution; the more this distribution deviates from normality, the larger n must be. Fortunately, convergence to normality in the central limit theorem is rapid and hence, as you observed in the exercises, we can get away with relatively small sample sizes (30 or more) in most cases.

In the variance test experiment, select the gamma distribution with shape parameter 1 and scale parameter 1. For the three different tests and for various significance levels, sample sizes, and values of σ_0 , run the experiment 1000 times. For each configuration, note the relative frequency of rejecting H_0 . When H_0 is true, compare the relative frequency with the significance level.

In the variance test experiment, select the uniform distribution on $[0, 4]$. For the three different tests and for various significance levels, sample sizes, and values of μ_0 , run the experiment 1000 times. For each configuration, note the relative frequency of rejecting H_0 . When H_0 is true, compare the relative frequency with the significance level.

Computational Exercises

The length of a certain machined part is supposed to be 10 centimeters. In fact, due to imperfections in the manufacturing process, the actual length is a random variable. The standard deviation is due to inherent factors in the process, which remain fairly stable over time. From historical data, the standard deviation is known with a high degree of accuracy to be 0.3. The mean, on the other hand, may be set by adjusting various parameters in the process and hence may change to an unknown value fairly frequently. We are interested in testing $H_0 : \mu = 10$ versus $H_1 : \mu \neq 10$.

1. Suppose that a sample of 100 parts has mean 10.1. Perform the test at the 0.1 level of significance.
2. Compute the P -value for the data in (a).
3. Compute the power of the test in (a) at $\mu = 10.05$.
4. Compute the approximate sample size needed for significance level 0.1 and power 0.8 when $\mu = 10.05$.

Answer

1. Test statistic 3.33, critical values ± 1.645 Reject H_0 .
2. $P = 0.0010$
3. The power of the test at 10.05 is approximately 0.0509.
4. Sample size 223

A bag of potato chips of a certain brand has an advertised weight of 250 grams. Actually, the weight (in grams) is a random variable. Suppose that a sample of 75 bags has mean 248 and standard deviation 5. At the 0.05 significance level, perform the following tests:

1. $H_0 : \mu \geq 250$ versus $H_1 : \mu < 250$
2. $H_0 : \sigma \geq 7$ versus $H_1 : \sigma < 7$

Answer

1. Test statistic -3.464 critical value -1.665 Reject H_0 .
2. $P < 0.0001$ so reject H_0 .

At a telemarketing firm, the length of a telephone solicitation (in seconds) is a random variable. A sample of 50 calls has mean 310 and standard deviation 25. At the 0.1 level of significance, can we conclude that

1. $\mu > 300$?
2. $\sigma > 20$?

Answer

1. Test statistic 2.828, critical value 1.2988. Reject H_0 .
2. $P = 0.0071$ so reject H_0 .

At a certain farm the weight of a peach (in ounces) at harvest time is a random variable. A sample of 100 peaches has mean 8.2 and standard deviation 1.0. At the 0.01 level of significance, can we conclude that

1. $\mu > 8$?
2. $\sigma < 1.5$?

Answer

1. Test statistic 2.0, critical value 2.363. Fail to reject H_0 .
2. $P < 0.0001$ so reject H_0 .

The hourly wage for a certain type of construction work is a random variable with standard deviation 1.25. For sample of 25 workers, the mean wage was \$6.75. At the 0.01 level of significance, can we conclude that $\mu < 7.00$?

Answer

Test statistic -1 , critical value -2.328 Fail to reject H_0 .

Data Analysis Exercises

Using Michelson's data, test to see if the velocity of light is greater than 730 (+299000) km/sec, at the 0.005 significance level.

Answer

Test statistic 15.49, critical value 2.6270. Reject H_0 .

Using Cavendish's data, test to see if the density of the earth is less than 5.5 times the density of water, at the 0.05 significance level.

Answer

Test statistic -1.269 , critical value -1.7017 Fail to reject H_0 .

Using Short's data, test to see if the parallax of the sun differs from 9 seconds of a degree, at the 0.1 significance level.

Answer

Test statistic -3.730 , critical value ± 1.6749 Reject H_0 .

Using Fisher's iris data, perform the following tests, at the 0.1 level:

1. The mean petal length of Setosa irises differs from 15 mm.
2. The mean petal length of Verginica irises is greater than 52 mm.
3. The mean petal length of Versicolor irises is less than 44 mm.

Answer

1. Test statistic -1.563 , critical values ± 1.672 Fail to reject H_0 .
2. Test statistic 4.556, critical value 1.2988. Reject H_0 .
3. Test statistic -1.028 , critical value -1.2988 Fail to Reject H_0 .

9.3: Tests in the Bernoulli Model

Basic Tests

Preliminaries

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Bernoulli distribution with unknown parameter $p \in (0, 1)$. Thus, these are independent random variables taking the values 1 and 0 with probabilities p and $1 - p$ respectively. In the usual language of reliability, 1 denotes *success* and 0 denotes *failure*, but of course these are generic terms. Often this model arises in one of the following contexts:

1. There is an *event* of interest in a basic experiment, with unknown probability p . We replicate the experiment n times and define $X_i = 1$ if and only if the event occurred on run i .
2. We have a population of objects of several different types; p is the unknown proportion of objects of a particular type of interest. We select n objects at random from the population and let $X_i = 1$ if and only if object i is of the type of interest. When the sampling is *with* replacement, these variables really do form a random sample from the Bernoulli distribution. When the sampling is *without* replacement, the variables are dependent, but the Bernoulli model may still be approximately valid if the population size is very large compared to the sample size n . For more on these points, see the discussion of sampling with and without replacement in the chapter on Finite Sampling Models.

In this section, we will construct hypothesis tests for the parameter p . The parameter space for p is the interval $(0, 1)$, and all hypotheses define subsets of this space. This section parallels the section on Estimation in the Bernoulli Model in the Chapter on Interval Estimation.

The Binomial Test

Recall that the number of successes $Y = \sum_{i=1}^n X_i$ has the binomial distribution with parameters n and p , and has probability density function given by

$$\mathbb{P}(Y = y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (9.3.1)$$

Recall also that the mean is $\mathbb{E}(Y) = np$ and variance is $\text{var}(Y) = np(1-p)$. Moreover Y is sufficient for p and hence is a natural candidate to be a test statistic for hypothesis tests about p . For $\alpha \in (0, 1)$, let $b_{n,p}(\alpha)$ denote the quantile of order α for the binomial distribution with parameters n and p . Since the binomial distribution is discrete, only certain (exact) quantiles are possible. For the remainder of this discussion, $p_0 \in (0, 1)$ is a conjectured value of p .

For every $\alpha \in (0, 1)$, the following tests have approximate significance level α :

1. Reject $H_0 : p = p_0$ versus $H_1 : p \neq p_0$ if and only if $Y \leq b_{n,p_0}(\alpha/2)$ or $Y \geq b_{n,p_0}(1 - \alpha/2)$.
2. Reject $H_0 : p \geq p_0$ versus $H_1 : p < p_0$ if and only if $Y \leq b_{n,p_0}(\alpha)$.
3. Reject $H_0 : p \leq p_0$ versus $H_1 : p > p_0$ if and only if $Y \geq b_{n,p_0}(1 - \alpha)$.

Proof

In part (a), H_0 is a simple hypothesis, and under H_0 the test statistic Y has the binomial distribution with parameter n and p_0 . Thus, if H_0 is true, then α is (approximately) the probability of falsely rejecting H_0 by definition of the quantiles. In parts (b) and (c), H_0 specifies a range of values of p . But if H_0 is true, the maximum type 1 probability is (approximately) α and occurs when $p = p_0$.

The test in (a) is the standard, symmetric, two-sided test, corresponding to probability $\alpha/2$ (approximately) in both tails of the binomial distribution under H_0 . The test in (b) is the left-tailed test and the test in (c) is the right-tailed test. As usual, we can generalize the two-sided test by partitioning α between the left and right tails of the binomial distribution in an arbitrary manner.

For any $\alpha, r \in (0, 1)$, the following test has (approximate) significance level α : Reject $H_0 : p = p_0$ versus $H_1 : p \neq p_0$ if and only if $Y \leq b_{n,p_0}(\alpha - r\alpha)$ or $Y \geq b_{n,p_0}(1 - r\alpha)$.

1. $r = \frac{1}{2}$ gives the standard symmetric two-sided test.
2. $r \downarrow 0$ gives the left-tailed test.

3. $r \uparrow 1$ gives the right-tailed test.

Proof

Once again, H_0 is a simple hypothesis and under H_0 , the test statistic Y has the binomial distribution with parameters n and p_0 . Thus if H_0 is true then the probability of falsely rejecting H_0 is α by definition of the quantiles. Parts (a)–(c) follow from properties of the quantile function.

An Approximate Normal Test

When n is large, the distribution of Y is approximately normal, by the central limit theorem, so we can construct an approximate normal test.

Suppose that the sample size n is large. For a conjectured $p_0 \in (0, 1)$, define the test statistic

$$Z = \frac{Y - np_0}{\sqrt{np_0(1-p_0)}} \quad (9.3.2)$$

1. If $p = p_0$, then Z has approximately a standard normal distribution.
2. If $p \neq p_0$, then Z has approximately a normal distribution with mean $\sqrt{n} \frac{p-p_0}{\sqrt{p_0(1-p_0)}}$ and variance $\frac{p(1-p)}{p_0(1-p_0)}$

Proof

1. This follows from the DeMoivre-Laplace theorem, the special case of the central limit theorem applied to the binomial distribution. Note that Z is simply the standard score associated with Y .
2. With some fairly simple algebra, we can write

$$Z = \sqrt{n} \frac{p - p_0}{\sqrt{p_0(1-p_0)}} + \sqrt{\frac{p(1-p)}{p_0(1-p_0)}} \frac{Y - np}{\sqrt{np(1-p)}} \quad (9.3.3)$$

The second factor in the second term is again simply the standard score associated with Y and hence this factor has approximately a standard normal distribution. So the result follows from the basic linearity property of the normal distribution.

As usual, for $\alpha \in (0, 1)$, let $z(\alpha)$ denote the quantile of order α for the standard normal distribution. For selected values of α , $z(\alpha)$ can be obtained from the special distribution calculator, or from most statistical software packages. Recall also by symmetry that $z(1-\alpha) = -z(\alpha)$.

For every $\alpha \in (0, 1)$, the following tests have approximate significance level α :

1. Reject $H_0 : p = p_0$ versus $H_1 : p \neq p_0$ if and only if $Z < -z(1-\alpha/2)$ or $Z > z(1-\alpha/2)$.
2. Reject $H_0 : p \geq p_0$ versus $H_1 : p < p_0$ if and only if $Z < -z(1-\alpha)$.
3. Reject $H_0 : p \leq p_0$ versus $H_1 : p \geq p_0$ if and only if $Z > z(1-\alpha)$.

Proof

In part (a), H_0 is a simple hypothesis and under H_0 the test statistic Z has approximately a standard normal distribution. Hence if H_0 is true then the probability of falsely rejecting H_0 is approximately α by definition of the quantiles. In parts (b) and (c), H_0 specifies a range of values of p , and under H_0 the test statistic Z has a nonstandard normal distribution, as described [above](#). The maximum type one error probability is α and occurs when $p = p_0$.

The test in (a) is the symmetric, two-sided test that corresponds to $\alpha/2$ in both tails of the distribution of Z , under H_0 . The test in (b) is the left-tailed test and the test in (c) is the right-tailed test. As usual, we can construct a more general two-sided test by partitioning α between the left and right tails of the standard normal distribution in an arbitrary manner.

For every $\alpha, r \in (0, 1)$, the following test has approximate significance level α : Reject $H_0 : p = p_0$ versus $H_1 : p \neq p_0$ if and only if $Z < z(\alpha - r\alpha)$ or $Z > z(1 - r\alpha)$.

1. $r = \frac{1}{2}$ gives the standard, symmetric two-sided test.
2. $r \downarrow 0$ gives the left-tailed test.

3. $r \uparrow 1$ gives the right-tailed test.

Proof

In part (a), H_0 is again a simple hypothesis, and under H_0 the test statistic Z has approximately a standard normal distribution. So if H_0 is true, the probability of falsely rejecting H_0 is α by definition of the quantiles.

Simulation Exercises

In the proportion test experiment, set $H_0 : p = p_0$, and select sample size 10, significance level 0.1, and $p_0 = 0.5$. For each $p \in \{0.1, 0.2, \dots, 0.9\}$ run the experiment 1000 times and then note the relative frequency of rejecting the null hypothesis. Graph the empirical power function.

In the proportion test experiment, repeat the previous exercise with sample size 20.

In the proportion test experiment, set $H_0 : p \leq p_0$, and select sample size 15, significance level 0.05, and $p_0 = 0.3$. For each $p \in \{0.1, 0.2, \dots, 0.9\}$ run the experiment 1000 times and note the relative frequency of rejecting the null hypothesis. Graph the empirical power function.

In the proportion test experiment, repeat the previous exercise with sample size 30.

In the proportion test experiment, set $H_0 : p \geq p_0$, and select sample size 20, significance level 0.01, and $p_0 = 0.6$. For each $p \in \{0.1, 0.2, \dots, 0.9\}$ run the experiment 1000 times and then note the relative frequency of rejecting the null hypothesis. Graph the empirical power function.

In the proportion test experiment, repeat the previous exercise with sample size 50.

Computational Exercises

In a pole of 1000 registered voters in a certain district, 427 prefer candidate X. At the 0.1 level, is the evidence sufficient to conclude that more than 40% of the registered voters prefer X?

Answer

Test statistic 1.743, critical value 1.282. Reject H_0 .

A coin is tossed 500 times and results in 302 heads. At the 0.05 level, test to see if the coin is unfair.

Answer

Test statistic 4.651, critical values ± 1.961 . Reject H_0 ; the coin is almost certainly unfair.

A sample of 400 memory chips from a production line are tested, and 32 are defective. At the 0.05 level, test to see if the proportion of defective chips is less than 0.1.

Answer

Test statistic -1.333 , critical value -1.645 . Fail to reject H_0 .

A new drug is administered to 50 patients and the drug is effective in 42 cases. At the 0.1 level, test to see if the success rate for the new drug is greater than 0.8.

Answer

Test statistic 0.707, critical value 1.282. Fail to reject H_0 .

Using the M&M data, test the following alternative hypotheses at the 0.1 significance level:

1. The proportion of red M&Ms differs from $\frac{1}{6}$.
2. The proportion of green M&Ms is less than $\frac{1}{6}$.

3. The proportion of yellow M&M is greater than $\frac{1}{6}$.

Answer

1. Test statistic 0.162, critical values ± 1.645 Fail to reject H_0 .
2. Test statistic -4.117 , critical value -1.282 Reject H_0 .
3. Test statistic 8.266, critical value 1.282. Reject H_0 .

The Sign Test

Derivation

Suppose now that we have a basic random experiment with a real-valued random variable U of interest. We assume that U has a continuous distribution with support on an interval of $S \subseteq \mathbb{R}$. Let m denote the quantile of a specified order $p_0 \in (0, 1)$ for the distribution of U . Thus, by definition,

$$p_0 = \mathbb{P}(U \leq m) \quad (9.3.4)$$

In general of course, m is unknown, even though p_0 is specified, because we don't know the distribution of U . Suppose that we want to construct hypothesis tests for m . For a given test value m_0 , let

$$p = \mathbb{P}(U \leq m_0) \quad (9.3.5)$$

Note that p is unknown even though m_0 is specified, because again, we don't know the distribution of U .

Relations

1. $m = m_0$ if and only if $p = p_0$.
2. $m < m_0$ if and only if $p > p_0$.
3. $m > m_0$ if and only if $p < p_0$.

Proof

These results follow since we are assuming that the distribution of U is continuous and is supported on the interval S .

As usual, we repeat the basic experiment n times to generate a random sample $\mathbf{U} = (U_1, U_2, \dots, U_n)$ of size n from the distribution of U . Let $X_i = \mathbf{1}(U_i \leq m_0)$ be the indicator variable of the event $\{U_i \leq m_0\}$ for $i \in \{1, 2, \dots, n\}$.

Note that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a statistic (an observable function of the data vector \mathbf{U}) and is a random sample of size n from the Bernoulli distribution with parameter p .

From the last two results it follows that tests of the unknown quantile m can be converted to tests of the Bernoulli parameter p , and thus the tests developed above apply. This procedure is known as the *sign test*, because essentially, only the sign of $U_i - m_0$ is recorded for each i . This procedure is also an example of a *nonparametric test*, because no assumptions about the distribution of U are made (except for continuity). In particular, we do not need to assume that the distribution of U belongs to a particular parametric family.

The most important special case of the sign test is the case where $p_0 = \frac{1}{2}$; this is the sign test of the median. If the distribution of U is known to be symmetric, the median and the mean agree. In this case, sign tests of the median are also tests of the mean.

Simulation Exercises

In the sign test experiment, set the sampling distribution to normal with mean 0 and standard deviation 2. Set the sample size to 10 and the significance level to 0.1. For each of the 9 values of m_0 , run the simulation 1000 times.

1. When $m = m_0$, give the empirical estimate of the significance level of the test and compare with 0.1.
2. In the other cases, give the empirical estimate of the power of the test.

In the sign test experiment, set the sampling distribution to uniform on the interval $[0, 5]$. Set the sample size to 20 and the significance level to 0.05. For each of the 9 values of m_0 , run the simulation 1000 times.

1. When $m = m_0$, give the empirical estimate of the significance level of the test and compare with 0.05.

2. In the other cases, give the empirical estimate of the power of the test.

In the sign test experiment, set the sampling distribution to gamma with shape parameter 2 and scale parameter 1. Set the sample size to 30 and the significance level to 0.025. For each of the 9 values of m_0 , run the simulation 1000 times.

1. When $m = m_0$, give the empirical estimate of the significance level of the test and compare with 0.025.
2. In the other cases, give the empirical estimate of the power of the test.

Computational Exercises

Using the M&M data, test to see if the median weight exceeds 47.9 grams, at the 0.1 level.

Answer

Test statistic 3.286, critical value 1.282. Reject H_0 .

Using Fisher's iris data, perform the following tests, at the 0.1 level:

1. The median petal length of Setosa irises differs from 15 mm.
2. The median petal length of Verginica irises is less than 52 mm.
3. The median petal length of Versicolor irises is less than 42 mm.

Answer

1. Test statistic 3.394, critical values ± 1.645 Reject H_0 .
2. Test statistic -1.980 , critical value -1.282 Reject H_0 .
3. Test statistic -0.566 , critical value -1.282 Fail to reject H_0 .

This page titled [9.3: Tests in the Bernoulli Model](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.4: Tests in the Two-Sample Normal Model

In this section, we will study hypothesis tests in the two-sample normal model and in the bivariate normal model. This section parallels the section on Estimation in the Two Sample Normal Model in the chapter on Interval Estimation.

The Two-Sample Normal Model

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size m from the normal distribution with mean μ and standard deviation σ , and that $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ is a random sample of size n from the normal distribution with mean ν and standard deviation τ . Moreover, suppose that the samples \mathbf{X} and \mathbf{Y} are independent.

This type of situation arises frequently when the random variables represent a measurement of interest for the objects of the population, and the samples correspond to two different treatments. For example, we might be interested in the blood pressure of a certain population of patients. The \mathbf{X} vector records the blood pressures of a control sample, while the \mathbf{Y} vector records the blood pressures of the sample receiving a new drug. Similarly, we might be interested in the yield of an acre of corn. The \mathbf{X} vector records the yields of a sample receiving one type of fertilizer, while the \mathbf{Y} vector records the yields of a sample receiving a different type of fertilizer.

Usually our interest is in a comparison of the parameters (either the mean or variance) for the two sampling distributions. In this section we will construct tests for the difference of the means and the ratio of the variances. As with previous estimation problems we have studied, the procedures vary depending on what parameters are known or unknown. Also as before, key elements in the construction of the tests are the sample means and sample variances and the special properties of these statistics when the sampling distribution is normal.

We will use the following notation for the sample mean and sample variance of a generic sample $\mathbf{U} = (U_1, U_2, \dots, U_k)$:

$$M(\mathbf{U}) = \frac{1}{k} \sum_{i=1}^k U_i, \quad S^2(\mathbf{U}) = \frac{1}{k-1} \sum_{i=1}^k [U_i - M(\mathbf{U})]^2 \quad (9.4.1)$$

Tests of the Difference in the Means with Known Standard Deviations

Our first discussion concerns tests for the difference in the means $\nu - \mu$ under the assumption that the standard deviations σ and τ are known. This is often, but not always, an unrealistic assumption. In some statistical problems, the variances are stable, and are at least approximately known, while the means may be different because of different treatments. Also this is a good place to start because the analysis is fairly easy.

For a conjectured difference of the means $\delta \in \mathbb{R}$, define the test statistic

$$Z = \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - \delta}{\sqrt{\sigma^2/m + \tau^2/n}} \quad (9.4.2)$$

1. If $\nu - \mu = \delta$ then Z has the standard normal distribution.
2. If $\nu - \mu \neq \delta$ then Z has the normal distribution with mean $[(\nu - \mu) - \delta] / \sqrt{\sigma^2/m + \tau^2/n}$ and variance 1.

Proof

From properties of normal samples, $M(\mathbf{X})$ has a normal distribution with mean μ and variance σ^2/m and similarly $M(\mathbf{Y})$ has a normal distribution with mean ν and variance τ^2/n . Since the samples are independent, $M(\mathbf{X})$ and $M(\mathbf{Y})$ are independent, so $M(\mathbf{Y}) - M(\mathbf{X})$ has a normal distribution with mean $\nu - \mu$ and variance $\sigma^2/m + \tau^2/n$. The final result then follows since Z is a linear function of $M(\mathbf{Y}) - M(\mathbf{X})$.

Of course (b) actually subsumes (a), but we separate them because the two cases play an important role in the hypothesis tests. In part (b), the non-zero mean can be viewed as a *non-centrality parameter*.

As usual, for $p \in (0, 1)$, let $z(p)$ denote the quantile of order p for the standard normal distribution. For selected values of p , $z(p)$ can be obtained from the special distribution calculator or from most statistical software packages. Recall also by symmetry that $z(1-p) = -z(p)$.

For every $\alpha \in (0, 1)$, the following tests have significance level α :

1. Reject $H_0 : \nu - \mu = \delta$ versus $H_1 : \nu - \mu \neq \delta$ if and only if $Z < -z(1 - \alpha/2)$ or $Z > z(1 - \alpha/2)$ if and only if $M(\mathbf{Y}) - M(\mathbf{X}) > \delta + z(1 - \alpha/2)\sqrt{\sigma^2/m + \tau^2/n}$ or $M(\mathbf{Y}) - M(\mathbf{X}) < \delta - z(1 - \alpha/2)\sqrt{\sigma^2/m + \tau^2/n}$.
2. Reject $H_0 : \nu - \mu \geq \delta$ versus $H_1 : \nu - \mu < \delta$ if and only if $Z < -z(1 - \alpha)$ if and only if $M(\mathbf{Y}) - M(\mathbf{X}) < \delta - z(1 - \alpha)\sqrt{\sigma^2/m + \tau^2/n}$.
3. Reject $H_0 : \nu - \mu \leq \delta$ versus $H_1 : \nu - \mu > \delta$ if and only if $Z > z(1 - \alpha)$ if and only if $M(\mathbf{Y}) - M(\mathbf{X}) > \delta + z(1 - \alpha)\sqrt{\sigma^2/m + \tau^2/n}$.

Proof

This follows the same logic that we have seen before. In part (a), H_0 is a simple hypothesis, and under this hypothesis Z has the standard normal distribution. Thus, if H_0 is true then the probability of falsely rejecting H_0 is α by definition of the quantiles. In parts (b) and (c), H_0 specifies a range of values of $\nu - \mu$, and under H_0 , Z has a nonstandard normal distribution, as described [above](#). But the largest type 1 error probability is α and occurs when $\nu - \mu = \delta$. The decision rules in terms of $M(\mathbf{Y}) - M(\mathbf{X})$ are equivalent to those in terms of Z by simple algebra.

For each of the tests above, we *fail* to reject H_0 at significance level α if and only if δ is in the corresponding $1 - \alpha$ level confidence interval.

1. $[M(\mathbf{Y}) - M(\mathbf{X})] - z(1 - \alpha/2)\sqrt{\sigma^2/m + \tau^2/n} \leq \delta \leq [M(\mathbf{Y}) - M(\mathbf{X})] + z(1 - \alpha/2)\sqrt{\sigma^2/m + \tau^2/n}$
2. $\delta \leq [M(\mathbf{Y}) - M(\mathbf{X})] + z(1 - \alpha)\sqrt{\sigma^2/m + \tau^2/n}$
3. $\delta \geq [M(\mathbf{Y}) - M(\mathbf{X})] - z(1 - \alpha)\sqrt{\sigma^2/m + \tau^2/n}$

Proof

These results follow from the previous results [above](#). In each case, we start with the inequality that corresponds to *not* rejecting the null hypothesis and solve for δ .

Tests of the Difference of the Means with Unknown Standard Deviations

Next we will construct tests for the difference in the means $\nu - \mu$ under the more realistic assumption that the standard deviations σ and τ are unknown. In this case, it is more difficult to find a suitable test statistic, but we can do the analysis in the special case that the standard deviations are the same. Thus, we will assume that $\sigma = \tau$, and the common value σ is unknown. This assumption is reasonable if there is an inherent variability in the measurement variables that does not change even when different treatments are applied to the objects in the population. Recall that the *pooled estimate* of the common variance σ^2 is the weighted average of the sample variances, with the degrees of freedom as the weight factors:

$$S^2(\mathbf{X}, \mathbf{Y}) = \frac{(m-1)S^2(\mathbf{X}) + (n-1)S^2(\mathbf{Y})}{m+n-2} \quad (9.4.3)$$

The statistic $S^2(\mathbf{X}, \mathbf{Y})$ is an unbiased and consistent estimator of the common variance σ^2 .

For a conjectured $\delta \in \mathbb{R}$ define the test statistic

$$T = \frac{[M(\mathbf{Y}) - M(\mathbf{X})] - \delta}{S(\mathbf{X}, \mathbf{Y})\sqrt{1/m + 1/n}} \quad (9.4.4)$$

1. If $\nu - \mu = \delta$ then T has the t distribution with $m + n - 2$ degrees of freedom,
2. If $\nu - \mu \neq \delta$ then T has a non-central t distribution with $m + n - 2$ degrees of freedom and non-centrality parameter

$$\frac{(\nu - \mu) - \delta}{\sigma\sqrt{1/m + 1/n}} \quad (9.4.5)$$

Proof

Part (b) actually subsumes part (a), since the ordinary t distribution is a special case of the non-central t distribution, with non-centrality parameter 0. With some basic algebra, we can write T in the form

$$T = \frac{Z + a}{\sqrt{V/(m+n-2)}} \quad (9.4.6)$$

where Z is the standard score of $M(\mathbf{Y}) - M(\mathbf{X})$, a is the non-centrality parameter given in the theorem, and $V = \frac{m+n-2}{\sigma^2} S^2(\mathbf{X}, \mathbf{Y})$. So Z has the standard normal distribution, V has the chi-square distribution with $m+n-2$ degrees of freedom, and Z and V are independent. Thus by definition, T has the non-central t distribution with $m+n-2$ degrees of freedom and non-centrality parameter a .

As usual, for $k > 0$ and $p \in (0, 1)$, let $t_k(p)$ denote the quantile of order p for the t distribution with k degrees of freedom. For selected values of k and p , values of $t_k(p)$ can be computed from the special distribution calculator, or from most statistical software packages. Recall also that, by symmetry, $t_k(1-p) = -t_k(p)$.

The following tests have significance level α :

1. Reject $H_0 : \nu - \mu = \delta$ versus $H_1 : \nu - \mu \neq \delta$ if and only if $T < -t_{m+n-2}(1-\alpha/2)$ or $T > t_{m+n-2}(1-\alpha/2)$ if and only if $M(\mathbf{Y}) - M(\mathbf{X}) > \delta + t_{m+n-2}(1-\alpha/2)\sqrt{\sigma^2/m + \tau^2/n}$ or $M(\mathbf{Y}) - M(\mathbf{X}) < \delta - t_{m+n-2}(1-\alpha/2)\sqrt{\sigma^2/m + \tau^2/n}$
2. Reject $H_0 : \nu - \mu \geq \delta$ versus $H_1 : \nu - \mu < \delta$ if and only if $T \leq -t_{m+n-2}(1-\alpha)$ if and only if $M(\mathbf{Y}) - M(\mathbf{X}) < \delta - t_{m+n-2}(1-\alpha)\sqrt{\sigma^2/m + \tau^2/n}$
3. Reject $H_0 : \nu - \mu \leq \delta$ versus $H_1 : \nu - \mu > \delta$ if and only if $T \geq t_{m+n-2}(1-\alpha)$ if and only if $M(\mathbf{Y}) - M(\mathbf{X}) > \delta + t_{m+n-2}(1-\alpha)\sqrt{\sigma^2/m + \tau^2/n}$

Proof

This follows the same logic that we have seen before. In part (a), H_0 is a simple hypothesis, and under this hypothesis T has the t distribution with $m+n-2$ degrees of freedom. Thus, if H_0 is true then the probability of falsely rejecting H_0 is α by definition of the quantiles. In parts (b) and (c), H_0 specifies a range of values of $\nu - \mu$, and under H_0 , T has a non-central t distribution, as described [above](#). But the largest type 1 error probability is α and occurs when $\nu - \mu = \delta$. The decision rules in terms of $M(\mathbf{Y}) - M(\mathbf{X})$ are equivalent to those in terms of T by simple algebra.

For each of the tests above, we fail to reject H_0 at significance level α if and only if δ is in the corresponding $1-\alpha$ level confidence interval.

1. $[M(\mathbf{Y}) - M(\mathbf{X})] - t_{m+n-2}(1-\alpha/2)\sqrt{\sigma^2/m + \tau^2/n} \leq \delta \leq [M(\mathbf{Y}) - M(\mathbf{X})] + t_{m+n-2}(1-\alpha/2)\sqrt{\sigma^2/m + \tau^2/n}$
2. $\delta \leq [M(\mathbf{Y}) - M(\mathbf{X})] + t_{m+n-2}(1-\alpha)\sqrt{\sigma^2/m + \tau^2/n}$
3. $\delta \geq [M(\mathbf{Y}) - M(\mathbf{X})] - t_{m+n-2}(1-\alpha)\sqrt{\sigma^2/m + \tau^2/n}$

Proof

These results follow from the previous results [above](#). In each case, we start with the inequality that corresponds to *not* rejecting the null hypothesis and solve for δ .

Tests of the Ratio of the Variances

Next we will construct tests for the ratio of the distribution variances τ^2/σ^2 . So the basic assumption is that the variances, and of course the means μ and ν are unknown.

For a conjectured $\rho \in (0, \infty)$, define the test statistics

$$F = \frac{S^2(\mathbf{X})}{S^2(\mathbf{Y})} \rho \quad (9.4.7)$$

1. If $\tau^2/\sigma^2 = \rho$ then F has the F distribution with $m-1$ degrees of freedom in the numerator and $n-1$ degrees of freedom in the denominator.
2. If $\tau^2/\sigma^2 \neq \rho$ then F has a scaled F distribution with $m-1$ degrees of freedom in the numerator, $n-1$ degrees of freedom in the denominator, and scale factor $\rho \frac{\sigma^2}{\tau^2}$.

Proof

Part (b) actually subsumes part (a) when $\rho = \tau^2/\sigma^2$, so we will just prove (b). Note that

$$F = \left(\frac{S^2(\mathbf{X})/\sigma^2}{S^2(\mathbf{Y})/\tau^2} \right) \rho \frac{\sigma^2}{\tau^2} \quad (9.4.8)$$

But $S^2(\mathbf{X})/\sigma^2$ has the chi-square distribution with $m - 1$ degrees of freedom, $S^2(\mathbf{Y})/\tau^2$ has the chi-square distribution with $n - 1$ degrees of freedom, and the variables are independent. Hence the ratio has the F distribution with $m - 1$ degrees of freedom in the numerator and $n - 1$ degrees of freedom in the denominator

The following tests have significance level α :

1. Reject $H_0 : \tau^2/\sigma^2 = \rho$ versus $H_1 : \tau^2/\sigma^2 \neq \rho$ if and only if $F > f_{m-1, n-1}(1 - \alpha/2)$ or $F < f_{m-1, n-1}(\alpha/2)$.
2. Reject $H_0 : \tau^2/\sigma^2 \leq \rho$ versus $H_1 : \tau^2/\sigma^2 > \rho$ if and only if $F < f_{m-1, n-1}(\alpha)$.
3. Reject $H_0 : \tau^2/\sigma^2 \geq \rho$ versus $H_1 : \tau^2/\sigma^2 < \rho$ if and only if $F > f_{m-1, n-1}(1 - \alpha)$.

Proof

The proof is the usual argument. In part (a), H_0 is a simple hypothesis, and under this hypothesis F has the f distribution with $m - 1$ degrees of freedom in the numerator $n - 1$ degrees of freedom in the denominator. Thus, if H_0 is true then the probability of falsely rejecting H_0 is α by definition of the quantiles. In parts (b) and (c), H_0 specifies a range of values of τ^2/σ^2 , and under H_0 , F has a scaled F distribution, as described [above](#). But the largest type 1 error probability is α and occurs when $\tau^2/\sigma^2 = \rho$.

For each of the tests above, we *fail* to reject H_0 at significance level α if and only if ρ_0 is in the corresponding $1 - \alpha$ level confidence interval.

1. $\frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} F_{m-1, n-1}(\alpha/2) \leq \rho \leq \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} F_{m-1, n-1}(1 - \alpha/2)$
2. $\rho \leq \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} F_{m-1, n-1}(\alpha)$
3. $\rho \geq \frac{S^2(\mathbf{Y})}{S^2(\mathbf{X})} F_{m-1, n-1}(1 - \alpha)$

Proof

These results follow from the previous results [above](#). In each case, we start with the inequality that corresponds to *not* rejecting the null hypothesis and solve for ρ .

Tests in the Bivariate Normal Model

In this subsection, we consider a model that is superficially similar to the two-sample normal model, but is actually much simpler. Suppose that

$$((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n)) \quad (9.4.9)$$

is a random sample of size n from the bivariate normal distribution of (X, Y) with $\mathbb{E}(X) = \mu$, $\mathbb{E}(Y) = \nu$, $\text{var}(X) = \sigma^2$, $\text{var}(Y) = \tau^2$, and $\text{cov}(X, Y) = \delta$.

Thus, instead of a *pair of samples*, we have a *sample of pairs*. The fundamental difference is that in this model, variables X and Y are measured on the *same* objects in a sample drawn from the population, while in the previous model, variables X and Y are measured on two distinct samples drawn from the population. The bivariate model arises, for example, in *before and after experiments*, in which a measurement of interest is recorded for a sample of n objects from the population, both before and after a treatment. For example, we could record the blood pressure of a sample of n patients, before and after the administration of a certain drug.

We will use our usual notation for the sample means and variances of $\mathbf{X} = (X_1, X_2, \dots, X_n)$ and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$. Recall also that the sample covariance of (\mathbf{X}, \mathbf{Y}) is

$$S(\mathbf{X}, \mathbf{Y}) = \frac{1}{n-1} \sum_{i=1}^n [X_i - M(\mathbf{X})][Y_i - M(\mathbf{Y})] \quad (9.4.10)$$

(not to be confused with the pooled estimate of the standard deviation in the two-sample model above).

The sequence of differences $\mathbf{Y} - \mathbf{X} = (Y_1 - X_1, Y_2 - X_2, \dots, Y_n - X_n)$ is a random sample of size n from the distribution of $Y - X$. The sampling distribution is normal with

1. $\mathbb{E}(Y - X) = \nu - \mu$
2. $\text{var}(Y - X) = \sigma^2 + \tau^2 - 2\delta$

The sample mean and variance of the sample of differences are

1. $M(\mathbf{Y} - \mathbf{X}) = M(\mathbf{Y}) - M(\mathbf{X})$
2. $S^2(\mathbf{Y} - \mathbf{X}) = S^2(\mathbf{X}) + S^2(\mathbf{Y}) - 2S(\mathbf{X}, \mathbf{Y})$

The sample of differences $\mathbf{Y} - \mathbf{X}$ fits the normal model for a single variable. The section on Tests in the Normal Model could be used to perform tests for the distribution mean $\nu - \mu$ and the distribution variance $\sigma^2 + \tau^2 - 2\delta$.

Computational Exercises

A new drug is being developed to reduce a certain blood chemical. A sample of 36 patients are given a placebo while a sample of 49 patients are given the drug. The statistics (in mg) are $m_1 = 87$, $s_1 = 4$, $m_2 = 63$, $s_2 = 6$. Test the following at the 10% significance level:

1. $H_0 : \sigma_1 = \sigma_2$ versus $H_1 : \sigma_1 \neq \sigma_2$.
2. $H_0 : \mu_1 \leq \mu_2$ versus $H_1 : \mu_1 > \mu_2$ (assuming that $\sigma_1 = \sigma_2$).
3. Based on (b), is the drug effective?

Answer

1. Test statistic 0.4, critical values 0.585, 1.667. Reject H_0 .
2. Test statistic 1.0, critical values ± 1.6625 Fail to reject H_0 .
3. Probably not

A company claims that an herbal supplement improves intelligence. A sample of 25 persons are given a standard IQ test before and after taking the supplement. The before and after statistics are $m_1 = 105$, $s_1 = 13$, $m_2 = 110$, $s_2 = 17$, $s_{1,2} = 190$. At the 10% significance level, do you believe the company's claim?

Answer

Test statistic 2.8, critical value 1.3184. Reject H_0 .

In Fisher's iris data, consider the petal length variable for the samples of Versicolor and Virginica irises. Test the following at the 10% significance level:

1. $H_0 : \sigma_1 = \sigma_2$ versus $H_1 : \sigma_1 \neq \sigma_2$.
2. $H_0 : \mu_1 \leq \mu_2$ versus $H_1 : \mu_1 > \mu_2$ (assuming that $\sigma_1 = \sigma_2$).

Answer

1. Test statistic 1.1, critical values 0.6227, 1.6072. Fail to reject H_0 .
2. Test statistic -11.4 , critical value -1.6602 Reject H_0 .

A plant has two machines that produce a circular rod whose diameter (in cm) is critical. A sample of 100 rods from the first machine has mean 10.3 and standard deviation 1.2. A sample of 100 rods from the second machine has mean 9.8 and standard deviation 1.6. Test the following hypotheses at the 10% level.

1. $H_0 : \sigma_1 = \sigma_2$ versus $H_1 : \sigma_1 \neq \sigma_2$.
2. $H_0 : \mu_1 = \mu_2$ versus $H_1 : \mu_1 \neq \mu_2$ (assuming that $\sigma_1 = \sigma_2$).

Answer

1. Test statistic 0.56, critical values 0.7175, 1.3942. Reject H_0 .
2. Test statistic -4.97 , critical values ± 1.645 Reject H_0 .

This page titled [9.4: Tests in the Two-Sample Normal Model](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.5: Likelihood Ratio Tests

Basic Theory

As usual, our starting point is a random experiment with an underlying sample space, and a probability measure \mathbb{P} . In the basic statistical model, we have an observable random variable \mathbf{X} taking values in a set S . In general, \mathbf{X} can have quite a complicated structure. For example, if the experiment is to sample n objects from a population and record various measurements of interest, then

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (9.5.1)$$

where X_i is the vector of measurements for the i th object. The most important special case occurs when (X_1, X_2, \dots, X_n) are independent and identically distributed. In this case, we have a random sample of size n from the common distribution.

In the previous sections, we developed tests for parameters based on natural test statistics. However, in other cases, the tests may not be parametric, or there may not be an obvious statistic to start with. Thus, we need a more general method for constructing test statistics. Moreover, we do not yet know if the tests constructed so far are the best, in the sense of maximizing the power for the set of alternatives. In this and the next section, we investigate both of these ideas. Likelihood functions, similar to those used in maximum likelihood estimation, will play a key role.

Tests of Simple Hypotheses

Suppose that \mathbf{X} has one of two possible distributions. Our simple hypotheses are

- $H_0 : \mathbf{X}$ has probability density function f_0 .
- $H_1 : \mathbf{X}$ has probability density function f_1 .

We will use subscripts on the probability measure \mathbb{P} to indicate the two hypotheses, and we assume that f_0 and f_1 are positive on S . The test that we will construct is based on the following simple idea: if we observe $\mathbf{X} = \mathbf{x}$, then the condition $f_1(\mathbf{x}) > f_0(\mathbf{x})$ is evidence in favor of the alternative; the opposite inequality is evidence against the alternative.

The *likelihood ratio function* $L : S \rightarrow (0, \infty)$ is defined by

$$L(\mathbf{x}) = \frac{f_0(\mathbf{x})}{f_1(\mathbf{x})}, \quad \mathbf{x} \in S \quad (9.5.2)$$

The statistic $L(\mathbf{X})$ is the *likelihood ratio statistic*.

Restating our earlier observation, note that small values of L are evidence in favor of H_1 . Thus it seems reasonable that the likelihood ratio statistic may be a good test statistic, and that we should consider tests in which we reject H_0 if and only if $L \leq l$, where l is a constant to be determined:

The significance level of the test is $\alpha = \mathbb{P}_0(L \leq l)$.

As usual, we can try to construct a test by choosing l so that α is a prescribed value. If \mathbf{X} has a discrete distribution, this will only be possible when α is a value of the distribution function of $L(\mathbf{X})$.

An important special case of this model occurs when the distribution of \mathbf{X} depends on a parameter θ that has two possible values. Thus, the parameter space is $\{\theta_0, \theta_1\}$, and f_0 denotes the probability density function of \mathbf{X} when $\theta = \theta_0$ and f_1 denotes the probability density function of \mathbf{X} when $\theta = \theta_1$. In this case, the hypotheses are equivalent to $H_0 : \theta = \theta_0$ versus $H_1 : \theta = \theta_1$.

As noted earlier, another important special case is when $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size n from a distribution an underlying random variable X taking values in a set R . In this case, $S = R^n$ and the probability density function f of \mathbf{X} has the form

$$f(x_1, x_2, \dots, x_n) = g(x_1)g(x_2) \cdots g(x_n), \quad (x_1, x_2, \dots, x_n) \in S \quad (9.5.3)$$

where g is the probability density function of X . So the hypotheses simplify to

- $H_0 : X$ has probability density function g_0 .

- $H_1 : X$ has probability density function g_1 .

and the likelihood ratio statistic is

$$L(X_1, X_2, \dots, X_n) = \prod_{i=1}^n \frac{g_0(X_i)}{g_1(X_i)} \quad (9.5.4)$$

In this special case, it turns out that under H_1 , the likelihood ratio statistic, as a function of the sample size n , is a martingale.

The Neyman-Pearson Lemma

The following theorem is the *Neyman-Pearson Lemma*, named for Jerzy Neyman and Egon Pearson. It shows that the test given above is most powerful. Let

$$R = \{\mathbf{x} \in S : L(\mathbf{x}) \leq l\} \quad (9.5.5)$$

and recall that the *size* of a rejection region is the significance of the test with that rejection region.

Consider the tests with rejection regions R given above and arbitrary $A \subseteq S$. If the size of R is at least as large as the size of A then the test with rejection region R is more powerful than the test with rejection region A . That is, if $\mathbb{P}_0(\mathbf{X} \in R) \geq \mathbb{P}_0(\mathbf{X} \in A)$ then $\mathbb{P}_1(\mathbf{X} \in R) \geq \mathbb{P}_1(\mathbf{X} \in A)$.

Proof

First note that from the definitions of L and R that the following inequalities hold:

$$\mathbb{P}_0(\mathbf{X} \in A) \leq l \mathbb{P}_1(\mathbf{X} \in A) \text{ for } A \subseteq R \quad (9.5.6)$$

$$\mathbb{P}_0(\mathbf{X} \in A) \geq l \mathbb{P}_1(\mathbf{X} \in A) \text{ for } A \subseteq R^c \quad (9.5.7)$$

Now for arbitrary $A \subseteq S$, write $R = (R \cap A) \cup (R \setminus A)$ and $A = (A \cap R) \cup (A \setminus R)$. From the additivity of probability and the inequalities above, it follows that

$$\mathbb{P}_1(\mathbf{X} \in R) - \mathbb{P}_1(\mathbf{X} \in A) \geq \frac{1}{l} [\mathbb{P}_0(\mathbf{X} \in R) - \mathbb{P}_0(\mathbf{X} \in A)] \quad (9.5.8)$$

Hence if $\mathbb{P}_0(\mathbf{X} \in R) \geq \mathbb{P}_0(\mathbf{X} \in A)$ then $\mathbb{P}_1(\mathbf{X} \in R) \geq \mathbb{P}_1(\mathbf{X} \in A)$.

The Neyman-Pearson lemma is more useful than might be first apparent. In many important cases, the *same* most powerful test works for a range of alternatives, and thus is a *uniformly* most powerful test for this range. Several special cases are discussed below.

Generalized Likelihood Ratio

The likelihood ratio statistic can be generalized to composite hypotheses. Suppose again that the probability density function f_θ of the data variable \mathbf{X} depends on a parameter θ , taking values in a parameter space Θ . Consider the hypotheses $\theta \in \Theta_0$ versus $\theta \notin \Theta_0$, where $\Theta_0 \subseteq \Theta$.

Define

$$L(\mathbf{x}) = \frac{\sup \{f_\theta(\mathbf{x}) : \theta \in \Theta_0\}}{\sup \{f_\theta(\mathbf{x}) : \theta \in \Theta\}} \quad (9.5.9)$$

The function L is the *likelihood ratio function* and $L(\mathbf{X})$ is the *likelihood ratio statistic*.

By the same reasoning as before, small values of $L(\mathbf{x})$ are evidence in favor of the alternative hypothesis.

Examples and Special Cases

Tests for the Exponential Model

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size $n \in \mathbb{N}_+$ from the exponential distribution with scale parameter $b \in (0, \infty)$. The sample variables might represent the lifetimes from a sample of devices of a certain type. We are interested in testing the simple hypotheses $H_0 : b = b_0$ versus $H_1 : b = b_1$, where $b_0, b_1 \in (0, \infty)$ are distinct specified values.

Recall that the sum of the variables is a sufficient statistic for b :

$$Y = \sum_{i=1}^n X_i \quad (9.5.10)$$

Recall also that Y has the gamma distribution with shape parameter n and scale parameter b . For $\alpha > 0$, we will denote the quantile of order α for this distribution by $\gamma_{n,b}(\alpha)$.

The likelihood ratio statistic is

$$L = \left(\frac{b_1}{b_0}\right)^n \exp\left[\left(\frac{1}{b_1} - \frac{1}{b_0}\right) Y\right] \quad (9.5.11)$$

Proof

Recall that the PDF g of the exponential distribution with scale parameter $b \in (0, \infty)$ is given by $g(x) = (1/b)e^{-x/b}$ for $x \in (0, \infty)$. If g_j denotes the PDF when $b = b_j$ for $j \in \{0, 1\}$ then

$$\frac{g_0(x)}{g_1(x)} = \frac{(1/b_0)e^{-x/b_0}}{(1/b_1)e^{-x/b_1}} = \frac{b_1}{b_0} e^{(1/b_1 - 1/b_0)x}, \quad x \in (0, \infty) \quad (9.5.12)$$

Hence the likelihood ratio function is

$$L(x_1, x_2, \dots, x_n) = \prod_{i=1}^n \frac{g_0(x_i)}{g_1(x_i)} = \left(\frac{b_1}{b_0}\right)^n e^{(1/b_1 - 1/b_0)y}, \quad (x_1, x_2, \dots, x_n) \in (0, \infty)^n \quad (9.5.13)$$

where $y = \sum_{i=1}^n x_i$.

The following tests are most powerful test at the α level

1. Suppose that $b_1 > b_0$. Reject $H_0 : b = b_0$ versus $H_1 : b = b_1$ if and only if $Y \geq \gamma_{n,b_0}(1 - \alpha)$.
2. Suppose that $b_1 < b_0$. Reject $H_0 : b = b_0$ versus $H_1 : b = b_1$ if and only if $Y \leq \gamma_{n,b_0}(\alpha)$.

Proof

Under H_0 , Y has the gamma distribution with parameters n and b_0 .

1. If $b_1 > b_0$ then $1/b_1 < 1/b_0$. From simple algebra, a rejection region of the form $L(\mathbf{X}) \leq l$ becomes a rejection region of the form $Y \geq y$. The precise value of y in terms of l is not important. For the test to have significance level α we must choose $y = \gamma_{n,b_0}(1 - \alpha)$.
2. If $b_1 < b_0$ then $1/b_1 > 1/b_0$. From simple algebra, a rejection region of the form $L(\mathbf{X}) \leq l$ becomes a rejection region of the form $Y \leq y$. Again, the precise value of y in terms of l is not important. For the test to have significance level α we must choose $y = \gamma_{n,b_0}(\alpha)$.

Note that these tests do not depend on the value of b_1 . This fact, together with the monotonicity of the power function can be used to show that the tests are uniformly most powerful for the usual one-sided tests.

Suppose that $b_0 \in (0, \infty)$.

1. The decision rule in part (a) above is uniformly most powerful for the test $H_0 : b \leq b_0$ versus $H_1 : b > b_0$.
2. The decision rule in part (b) above is uniformly most powerful for the test $H_0 : b \geq b_0$ versus $H_1 : b < b_0$.

Tests for the Bernoulli Model

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size $n \in \mathbb{N}_+$ from the Bernoulli distribution with success parameter p . The sample could represent the results of tossing a coin n times, where p is the probability of heads. We wish to test the simple hypotheses $H_0 : p = p_0$ versus $H_1 : p = p_1$, where $p_0, p_1 \in (0, 1)$ are distinct specified values. In the coin tossing model, we know that the probability of heads is either p_0 or p_1 , but we don't know which.

Recall that the number of successes is a sufficient statistic for p :

$$Y = \sum_{i=1}^n X_i \quad (9.5.14)$$

Recall also that Y has the binomial distribution with parameters n and p . For $\alpha \in (0, 1)$, we will denote the quantile of order α for this distribution by $b_{n,p}(\alpha)$; although since the distribution is discrete, only certain values of α are possible.

The likelihood ratio statistic is

$$L = \left(\frac{1-p_0}{1-p_1} \right)^n \left[\frac{p_0(1-p_1)}{p_1(1-p_0)} \right]^Y \quad (9.5.15)$$

Proof

Recall that the PDF g of the Bernoulli distribution with parameter $p \in (0, 1)$ is given by $g(x) = p^x(1-p)^{1-x}$ for $x \in \{0, 1\}$. If g_j denotes the PDF when $p = p_j$ for $j \in \{0, 1\}$ then

$$\frac{g_0(x)}{g_1(x)} = \frac{p_0^x(1-p_0)^{1-x}}{p_1^x(1-p_1)^{1-x}} = \left(\frac{p_0}{p_1} \right)^x \left(\frac{1-p_0}{1-p_1} \right)^{1-x} = \left(\frac{1-p_0}{1-p_1} \right) \left[\frac{p_0(1-p_1)}{p_1(1-p_0)} \right]^x, \quad x \in \{0, 1\} \quad (9.5.16)$$

Hence the likelihood ratio function is

$$L(x_1, x_2, \dots, x_n) = \prod_{i=1}^n \frac{g_0(x_i)}{g_1(x_i)} = \left(\frac{1-p_0}{1-p_1} \right)^n \left[\frac{p_0(1-p_1)}{p_1(1-p_0)} \right]^y, \quad (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \quad (9.5.17)$$

where $y = \sum_{i=1}^n x_i$.

The following tests are most powerful test at the α level

1. Suppose that $p_1 > p_0$. Reject $H_0 : p = p_0$ versus $H_1 : p = p_1$ if and only if $Y \geq b_{n,p_0}(1-\alpha)$.
2. Suppose that $p_1 < p_0$. Reject $p = p_0$ versus $p = p_1$ if and only if $Y \leq b_{n,p_0}(\alpha)$.

Proof

Under H_0 , Y has the binomial distribution with parameters n and p_0 .

1. If $p_1 > p_0$ then $p_0(1-p_1)/p_1(1-p_0) < 1$. From simple algebra, a rejection region of the form $L(\mathbf{X}) \leq l$ becomes a rejection region of the form $Y \geq y$. The precise value of y in terms of l is not important. For the test to have significance level α we must choose $y = b_{n,p_0}(1-\alpha)$.
2. If $p_1 < p_0$ then $p_0(1-p_1)/p_1(1-p_0) > 1$. From simple algebra, a rejection region of the form $L(\mathbf{X}) \leq l$ becomes a rejection region of the form $Y \leq y$. Again, the precise value of y in terms of l is not important. For the test to have significance level α we must choose $y = b_{n,p_0}(\alpha)$.

Note that these tests do not depend on the value of p_1 . This fact, together with the monotonicity of the power function can be used to show that the tests are uniformly most powerful for the usual one-sided tests.

Suppose that $p_0 \in (0, 1)$.

1. The decision rule in part (a) above is uniformly most powerful for the test $H_0 : p \leq p_0$ versus $H_1 : p > p_0$.
2. The decision rule in part (b) above is uniformly most powerful for the test $H_0 : p \geq p_0$ versus $H_1 : p < p_0$.

Tests in the Normal Model

The one-sided tests that we derived in the normal model, for μ with σ known, for μ with σ unknown, and for σ with μ unknown are all uniformly most powerful. On the other hand, none of the two-sided tests are uniformly most powerful.

A Nonparametric Example

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample of size $n \in \mathbb{N}_+$, either from the Poisson distribution with parameter 1 or from the geometric distribution on \mathbb{N} with parameter $p = \frac{1}{2}$. Note that both distributions have mean 1 (although the Poisson distribution has variance 1 while the geometric distribution has variance 2). So, we wish to test the hypotheses

- $H_0 : X$ has probability density function $g_0(x) = e^{-1} \frac{1}{x!}$ for $x \in \mathbb{N}$.

- $H_1 : X$ has probability density function $g_1(x) = \left(\frac{1}{2}\right)^{x+1}$ for $x \in \mathbb{N}$.

The likelihood ratio statistic is

$$L = 2^n e^{-n} \frac{2^Y}{U} \text{ where } Y = \sum_{i=1}^n X_i \text{ and } U = \prod_{i=1}^n X_i! \quad (9.5.18)$$

Proof

Note that

$$\frac{g_0(x)}{g_1(x)} = \frac{e^{-1}/x!}{(1/2)^{x+1}} = 2e^{-1} \frac{2^x}{x!}, \quad x \in \mathbb{N} \quad (9.5.19)$$

Hence the likelihood ratio function is

$$L(x_1, x_2, \dots, x_n) = \prod_{i=1}^n \frac{g_0(x_i)}{g_1(x_i)} = 2^n e^{-n} \frac{2^y}{u}, \quad (x_1, x_2, \dots, x_n) \in \mathbb{N}^n \quad (9.5.20)$$

where $y = \sum_{i=1}^n x_i$ and $u = \prod_{i=1}^n x_i!$.

The most powerful tests have the following form, where d is a constant: reject H_0 if and only if $\ln(2)Y - \ln(U) \leq d$.

Proof

A rejection region of the form $L(\mathbf{X}) \leq l$ is equivalent to

$$\frac{2^Y}{U} \leq \frac{le^n}{2^n} \quad (9.5.21)$$

Taking the natural logarithm, this is equivalent to $\ln(2)Y - \ln(U) \leq d$ where $d = n + \ln(l) - n \ln(2)$

This page titled [9.5: Likelihood Ratio Tests](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

9.6: Chi-Square Tests

In this section, we will study a number of important hypothesis tests that fall under the general term *chi-square tests*. These are named, as you might guess, because in each case the test statistics has (in the limit) a chi-square distribution. Although there are several different tests in this general category, they all share some common themes:

- In each test, there are one or more underlying multinomial samples. Of course, the multinomial model includes the Bernoulli model as a special case.
- Each test works by comparing the *observed frequencies* of the various outcomes with *expected frequencies* under the null hypothesis.
- If the model is *incompletely specified*, some of the expected frequencies must be estimated; this reduces the degrees of freedom in the limiting chi-square distribution.

We will start with the simplest case, where the derivation is the most straightforward; in fact this test is equivalent to a test we have already studied. We then move to successively more complicated models.

The One-Sample Bernoulli Model

Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a random sample from the Bernoulli distribution with unknown success parameter $p \in (0, 1)$. Thus, these are independent random variables taking the values 1 and 0 with probabilities p and $1 - p$ respectively. We want to test $H_0 : p = p_0$ versus $H_1 : p \neq p_0$, where $p_0 \in (0, 1)$ is specified. Of course, we have already studied such tests in the Bernoulli model. But keep in mind that our methods in this section will generalize to a variety of new models that we have not yet studied.

Let $O_1 = \sum_{j=1}^n X_j$ and $O_0 = n - O_1 = \sum_{j=1}^n (1 - X_j)$. These statistics give the number of times (frequency) that outcomes 1 and 0 occur, respectively. Moreover, we know that each has a binomial distribution; O_1 has parameters n and p , while O_0 has parameters n and $1 - p$. In particular, $\mathbb{E}(O_1) = np$, $\mathbb{E}(O_0) = n(1 - p)$, and $\text{var}(O_1) = \text{var}(O_0) = np(1 - p)$. Moreover, recall that O_1 is sufficient for p . Thus, any good test statistic should be a function of O_1 . Next, recall that when n is large, the distribution of O_1 is approximately normal, by the central limit theorem. Let

$$Z = \frac{O_1 - np_0}{\sqrt{np_0(1 - p_0)}} \quad (9.6.1)$$

Note that Z is the standard score of O_1 under H_0 . Hence if n is large, Z has approximately the standard normal distribution under H_0 , and therefore $V = Z^2$ has approximately the chi-square distribution with 1 degree of freedom under H_0 . As usual, let χ_k^2 denote the quantile function of the chi-square distribution with k degrees of freedom.

An approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi_1^2(1 - \alpha)$.

The test above is equivalent to the unbiased test with test statistic Z (the approximate normal test) derived in the section on Tests in the Bernoulli model.

For purposes of generalization, the critical result in the next exercise is a special representation of V . Let $e_0 = n(1 - p_0)$ and $e_1 = np_0$. Note that these are the expected frequencies of the outcomes 0 and 1, respectively, under H_0 .

V can be written in terms of the observed and expected frequencies as follows:

$$V = \frac{(O_0 - e_0)^2}{e_0} + \frac{(O_1 - e_1)^2}{e_1} \quad (9.6.2)$$

This representation shows that our test statistic V measures the discrepancy between the expected frequencies, under H_0 , and the observed frequencies. Of course, large values of V are evidence in favor of H_1 . Finally, note that although there are two terms in the expansion of V in Exercise 3, there is only one degree of freedom since $O_0 + O_1 = n$. The observed and expected frequencies could be stored in a 1×2 table.

The Multi-Sample Bernoulli Model

Suppose now that we have samples from several (possibly) different, independent Bernoulli trials processes. Specifically, suppose that $\mathbf{X}_i = (X_{i,1}, X_{i,2}, \dots, X_{i,n_i})$ is a random sample of size n_i from the Bernoulli distribution with unknown success parameter $p_i \in (0, 1)$ for each $i \in \{1, 2, \dots, m\}$. Moreover, the samples $(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_m)$ are independent. We want to test hypotheses about the unknown parameter vector $\mathbf{p} = (p_1, p_2, \dots, p_m)$. There are two common cases that we consider below, but first let's set up the essential notation

that we will need for both cases. For $i \in \{1, 2, \dots, m\}$ and $j \in \{0, 1\}$, let $O_{i,j}$ denote the number of times that outcome j occurs in sample \mathbf{X}_i . The observed frequency $O_{i,j}$ has a binomial distribution; $O_{i,1}$ has parameters n_i and p_i while $O_{i,0}$ has parameters n_i and $1 - p_i$.

The Completely Specified Case

Consider a specified parameter vector $\mathbf{p}_0 = (p_{0,1}, p_{0,2}, \dots, p_{0,m}) \in (0, 1)^m$. We want to test the null hypothesis $H_0 : \mathbf{p} = \mathbf{p}_0$, versus $H_1 : \mathbf{p} \neq \mathbf{p}_0$. Since the null hypothesis specifies the value of p_i for each i , this is called the *completely specified case*. Now let $e_{i,0} = n_i(1 - p_{i,0})$ and let $e_{i,1} = n_i p_{i,0}$. These are the expected frequencies of the outcomes 0 and 1, respectively, from sample \mathbf{X}_i under H_0 .

If n_i is large for each i , then under H_0 the following test statistic has approximately the chi-square distribution with m degrees of freedom:

$$V = \sum_{i=1}^m \sum_{j=0}^1 \frac{(O_{i,j} - e_{i,j})^2}{e_{i,j}} \quad (9.6.3)$$

Proof

This follows from the result [above](#) and independence.

As a rule of thumb, “large” means that we need $e_{i,j} \geq 5$ for each $i \in \{1, 2, \dots, m\}$ and $j \in \{0, 1\}$. But of course, the larger these expected frequencies the better.

Under the large sample assumption, an approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi_m^2(1 - \alpha)$.

Once again, note that the test statistic V measures the discrepancy between the expected and observed frequencies, over all outcomes and all samples. There are $2m$ terms in the expansion of V in Exercise 4, but only m degrees of freedom, since $O_{i,0} + O_{i,1} = n_i$ for each $i \in \{1, 2, \dots, m\}$. The observed and expected frequencies could be stored in an $m \times 2$ table.

The Equal Probability Case

Suppose now that we want to test the null hypothesis $H_0 : p_1 = p_2 = \dots = p_m$ that all of the success probabilities are the same, versus the complementary alternative hypothesis H_1 that the probabilities are not all the same. Note, in contrast to the previous model, that the null hypothesis does not specify the value of the common success probability p . But note also that under the null hypothesis, the m samples can be combined to form one large sample of Bernoulli trials with success probability p . Thus, a natural approach is to estimate p and then define the test statistic that measures the discrepancy between the expected and observed frequencies, just as before. The challenge will be to find the distribution of the test statistic.

Let $n = \sum_{i=1}^m n_i$ denote the total sample size when the samples are combined. Then the overall sample mean, which in this context is the overall sample proportion of successes, is

$$P = \frac{1}{n} \sum_{i=1}^m \sum_{j=1}^{n_i} X_{i,j} = \frac{1}{n} \sum_{i=1}^m O_{i,1} \quad (9.6.4)$$

The sample proportion P is the best estimate of p , in just about any sense of the word. Next, let $E_{i,0} = n_i(1 - P)$ and $E_{i,1} = n_i P$. These are the *estimated* expected frequencies of 0 and 1, respectively, from sample \mathbf{X}_i under H_0 . Of course these estimated frequencies are now *statistics* (and hence random) rather than parameters. Just as before, we define our test statistic

$$V = \sum_{i=1}^m \sum_{j=0}^1 \frac{(O_{i,j} - E_{i,j})^2}{E_{i,j}} \quad (9.6.5)$$

It turns out that under H_0 , the distribution of V converges to the chi-square distribution with $m - 1$ degrees of freedom as $n \rightarrow \infty$.

An approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi_{m-1}^2(1 - \alpha)$.

Intuitively, we lost a degree of freedom over the [completely specified case](#) because we had to estimate the unknown common success probability p . Again, the observed and expected frequencies could be stored in an $m \times 2$ table.

The One-Sample Multinomial Model

Our next model generalizes the [one-sample Bernoulli model](#) in a different direction. Suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of multinomial trials. Thus, these are independent, identically distributed random variables, each taking values in a set S with k elements. If we want, we can assume that $S = \{0, 1, \dots, k-1\}$; the one-sample Bernoulli model then corresponds to $k = 2$. Let f denote the common probability density function of the sample variables on S , so that $f(j) = \mathbb{P}(X_i = j)$ for $i \in \{1, 2, \dots, n\}$ and $j \in S$. The values of f are assumed unknown, but of course we must have $\sum_{j \in S} f(j) = 1$, so there are really only $k-1$ unknown parameters. For a given probability density function f_0 on S we want to test $H_0 : f = f_0$ versus $H_1 : f \neq f_0$.

By this time, our general approach should be clear. We let O_j denote the number of times that outcome $j \in S$ occurs in sample \mathbf{X} :

$$O_j = \sum_{i=1}^n \mathbf{1}(X_i = j) \quad (9.6.6)$$

Note that O_j has the binomial distribution with parameters n and $f(j)$. Thus, $e_j = n f_0(j)$ is the expected number of times that outcome j occurs, under H_0 . Our test statistic, of course, is

$$V = \sum_{j \in S} \frac{(O_j - e_j)^2}{e_j} \quad (9.6.7)$$

It turns out that under H_0 , the distribution of V converges to the chi-square distribution with $k-1$ degrees of freedom as $n \rightarrow \infty$. Note that there are k terms in the expansion of V , but only $k-1$ degrees of freedom since $\sum_{j \in S} O_j = n$.

An approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi_{k-1}^2(1-\alpha)$.

Again, as a rule of thumb, we need $e_j \geq 5$ for each $j \in S$, but the larger the expected frequencies the better.

The Multi-Sample Multinomial Model

As you might guess, our final generalization is to the multi-sample multinomial model. Specifically, suppose that $\mathbf{X}_i = (X_{i,1}, X_{i,2}, \dots, X_{i,n_i})$ is a random sample of size n_i from a distribution on a set S with k elements, for each $i \in \{1, 2, \dots, m\}$. Moreover, we assume that the samples $(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_m)$ are independent. Again there is no loss in generality if we take $S = \{0, 1, \dots, k-1\}$. Then $k=2$ reduces to the [multi-sample Bernoulli model](#), and $m=1$ corresponds to the [one-sample multinomial model](#).

Let f_i denote the common probability density function of the variables in sample \mathbf{X}_i , so that $f_i(j) = \mathbb{P}(X_{i,l} = j)$ for $i \in \{1, 2, \dots, m\}$, $l \in \{1, 2, \dots, n_i\}$, and $j \in S$. These are generally unknown, so that our vector of parameters is the vector of probability density functions: $\mathbf{f} = (f_1, f_2, \dots, f_m)$. Of course, $\sum_{j \in S} f_i(j) = 1$ for $i \in \{1, 2, \dots, m\}$ so there are actually $m(k-1)$ unknown parameters. We are interested in testing hypotheses about \mathbf{f} . As in the multi-sample Bernoulli model, there are two common cases that we consider below, but first let's set up the essential notation that we will need for both cases. For $i \in \{1, 2, \dots, m\}$ and $j \in S$, let $O_{i,j}$ denote the number of times that outcome j occurs in sample \mathbf{X}_i . The observed frequency $O_{i,j}$ has a binomial distribution with parameters n_i and $f_i(j)$.

The Completely Specified Case

Consider a given vector of probability density functions on S , denoted $\mathbf{f}_0 = (f_{0,1}, f_{0,2}, \dots, f_{0,m})$. We want to test the null hypothesis $H_0 : \mathbf{f} = \mathbf{f}_0$, versus $H_1 : \mathbf{f} \neq \mathbf{f}_0$. Since the null hypothesis specifies the value of $f_i(j)$ for each i and j , this is called the *completely specified case*. Let $e_{i,j} = n_i f_{0,i}(j)$. This is the expected frequency of outcome j in sample \mathbf{X}_i under H_0 .

If n_i is large for each i , then under H_0 , the test statistic V below has approximately the chi-square distribution with $m(k-1)$ degrees of freedom:

$$V = \sum_{i=1}^m \sum_{j \in S} \frac{(O_{i,j} - e_{i,j})^2}{e_{i,j}} \quad (9.6.8)$$

Proof

This follows from the [one-sample multinomial case](#) and independence.

As usual, our rule of thumb is that we need $e_{i,j} \geq 5$ for each $i \in \{1, 2, \dots, m\}$ and $j \in S$. But of course, the larger these expected frequencies the better.

Under the large sample assumption, an approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi_{m(k-1)}^2(1-\alpha)$.

As always, the test statistic V measures the discrepancy between the expected and observed frequencies, over all outcomes and all samples. There are mk terms in the expansion of V in Exercise 8, but we lose m degrees of freedom, since $\sum_{j \in S} O_{i,j} = n_i$ for each $i \in \{1, 2, \dots, m\}$.

The Equal PDF Case

Suppose now that we want to test the null hypothesis $H_0 : f_1 = f_2 = \dots = f_m$ that all of the probability density functions are the same, versus the complementary alternative hypothesis H_1 that the probability density functions are not all the same. Note, in contrast to the previous model, that the null hypothesis does not specify the value of the common success probability density function f . But note also that under the null hypothesis, the m samples can be combined to form one large sample of multinomial trials with probability density function f . Thus, a natural approach is to estimate the values of f and then define the test statistic that measures the discrepancy between the expected and observed frequencies, just as before.

Let $n = \sum_{i=1}^m n_i$ denote the total sample size when the samples are combined. Under H_0 , our best estimate of $f(j)$ is

$$P_j = \frac{1}{n} \sum_{i=1}^m O_{i,j} \quad (9.6.9)$$

Hence our estimate of the expected frequency of outcome j in sample \mathbf{X}_i under H_0 is $E_{i,j} = n_i P_j$. Again, this estimated frequency is now a *statistic* (and hence random) rather than a parameter. Just as before, we define our test statistic

$$V = \sum_{i=1}^m \sum_{j \in S} \frac{(O_{i,j} - E_{i,j})^2}{E_{i,j}} \quad (9.6.10)$$

As you no doubt expect by now, it turns out that under H_0 , the distribution of V converges to a chi-square distribution as $n \rightarrow \infty$. But let's see if we can determine the degrees of freedom heuristically.

The limiting distribution of V has $(k-1)(m-1)$ degrees of freedom.

Proof

There are km terms in the expansion of V . We lose m degrees of freedom since $\sum_{j \in S} O_{i,j} = n_i$ for each $i \in \{1, 2, \dots, m\}$. We must estimate all but one of the probabilities $f(j)$ for $j \in S$, thus losing $k-1$ degrees of freedom.

An approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi_{(k-1)(m-1)}^2(1-\alpha)$.

A Goodness of Fit Test

A *goodness of fit test* is an hypothesis test that an unknown sampling distribution is a particular, specified distribution or belongs to a parametric family of distributions. Such tests are clearly fundamental and important. The [one-sample multinomial model](#) leads to a quite general goodness of fit test.

To set the stage, suppose that we have an observable random variable X for an experiment, taking values in a general set S . Random variable X might have a continuous or discrete distribution, and might be single-variable or multi-variable. We want to test the null hypothesis that X has a given, completely specified distribution, or that the distribution of X belongs to a particular parametric family.

Our first step, in either case, is to sample from the distribution of X to obtain a sequence of independent, identically distributed variables $\mathbf{X} = (X_1, X_2, \dots, X_n)$. Next, we select $k \in \mathbb{N}_+$ and partition S into k (disjoint) subsets. We will denote the partition by $\{A_j : j \in J\}$ where $\#(J) = k$. Next, we define the sequence of random variables $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ by $Y_i = j$ if and only if $X_i \in A_j$ for $i \in \{1, 2, \dots, n\}$ and $j \in J$.

\mathbf{Y} is a multinomial trials sequence with parameters n and f , where $f(j) = \mathbb{P}(X \in A_j)$ for $j \in J$.

The Completely Specified Case

Let H denote the statement that X has a given, completely specified distribution. Let f_0 denote the probability density function on J defined by $f_0(j) = \mathbb{P}(X \in A_j | H)$ for $j \in J$. To test hypothesis H , we can formally test $H_0 : f = f_0$ versus $H_1 : f \neq f_0$, which of course, is precisely the problem we solved in the [one-sample multinomial model](#).

Generally, we would partition the space S into as many subsets as possible, subject to the restriction that the expected frequencies all be at least 5.

The Partially Specified Case

Often we don't really want to test whether X has a completely specified distribution (such as the normal distribution with mean 5 and variance 9), but rather whether the distribution of X belongs to a specified parametric family (such as the normal). A natural course of action in this case would be to estimate the unknown parameters and then proceed just as above. As we have seen before, the expected frequencies would be statistics E_j because they would be based on the estimated parameters. As a rule of thumb, we lose a degree of freedom in the chi-square statistic V for each parameter that we estimate, although the precise mathematics can be complicated.

A Test of Independence

Suppose that we have observable random variables X and Y for an experiment, where X takes values in a set S with k elements, and Y takes values in a set T with m elements. Let f denote the joint probability density function of (X, Y) , so that $f(i, j) = \mathbb{P}(X = i, Y = j)$ for $i \in S$ and $j \in T$. Recall that the marginal probability density functions of X and Y are the functions g and h respectively, where

$$g(i) = \sum_{j \in T} f(i, j), \quad i \in S \quad (9.6.11)$$

$$h(j) = \sum_{i \in S} f(i, j), \quad j \in T \quad (9.6.12)$$

Usually, of course, f , g , and h are unknown. In this section, we are interested in testing whether X and Y are independent, a basic and important test. Formally then we want to test the null hypothesis

$$H_0 : f(i, j) = g(i) h(j), \quad (i, j) \in S \times T \quad (9.6.13)$$

versus the complementary alternative H_1 .

Our first step, of course, is to draw a random sample $(\mathbf{X}, \mathbf{Y}) = ((X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n))$ from the distribution of (X, Y) . Since the state spaces are finite, this sample forms a sequence of multinomial trials. Thus, with our usual notation, let $O_{i,j}$ denote the number of times that (i, j) occurs in the sample, for each $(i, j) \in S \times T$. This statistic has the binomial distribution with trial parameter n and success parameter $f(i, j)$. Under H_0 , the success parameter is $g(i) h(j)$. However, since we don't know the success parameters, we must estimate them in order to compute the expected frequencies. Our best estimate of $f(i, j)$ is the sample proportion $\frac{1}{n} O_{i,j}$. Thus, our best estimates of $g(i)$ and $h(j)$ are $\frac{1}{n} N_i$ and $\frac{1}{n} M_j$, respectively, where N_i is the number of times that i occurs in sample \mathbf{X} and M_j is the number of times that j occurs in sample \mathbf{Y} :

$$N_i = \sum_{j \in T} O_{i,j} \quad (9.6.14)$$

$$M_j = \sum_{i \in S} O_{i,j} \quad (9.6.15)$$

Thus, our estimate of the expected frequency of (i, j) under H_0 is

$$E_{i,j} = n \frac{1}{n} N_i \frac{1}{n} M_j = \frac{1}{n} N_i M_j \quad (9.6.16)$$

Of course, we define our test statistic by

$$V = \sum_{i \in S} \sum_{j \in T} \frac{(O_{i,j} - E_{i,j})^2}{E_{i,j}} \quad (9.6.17)$$

As you now expect, the distribution of V converges to a chi-square distribution as $n \rightarrow \infty$. But let's see if we can determine the appropriate degrees of freedom on heuristic grounds.

The limiting distribution of V has $(k-1)(m-1)$ degrees of freedom.

Proof

There are km terms in the expansion of V . We lose one degree of freedom since $\sum_{i \in S} \sum_{j \in T} O_{i,j} = n$. We must estimate all but one of the probabilities $g(i)$ for $i \in S$, thus losing $k-1$ degrees of freedom. We must estimate all but one of the probabilities $h(j)$ for $j \in T$, thus losing $m-1$ degrees of freedom.

An approximate test of H_0 versus H_1 at the α level of significance is to reject H_0 if and only if $V > \chi^2_{(k-1)(m-1)}(1-\alpha)$.

The observed frequencies are often recorded in a $k \times m$ table, known as a *contingency table*, so that $O_{i,j}$ is the number in row i and column j . In this setting, note that N_i is the sum of the frequencies in the i th row and M_j is the sum of the frequencies in the j th column. Also, for historical reasons, the random variables X and Y are sometimes called *factors* and the possible values of the variables *categories*.

Computational and Simulation Exercises

Computational Exercises

In each of the following exercises, specify the number of degrees of freedom of the chi-square statistic, give the value of the statistic and compute the P -value of the test.

A coin is tossed 100 times, resulting in 55 heads. Test the null hypothesis that the coin is fair.

Answer

1 degree of freedom, $V = 1$, $P = 0.3173$

Suppose that we have 3 coins. The coins are tossed, yielding the data in the following table:

	Heads	Tails
Coin 1	29	21
Coin 2	23	17
Coin 3	42	18

1. Test the null hypothesis that all 3 coin are fair.
2. Test the null hypothesis that coin 1 has probability of heads $\frac{3}{5}$; coin 2 is fair; and coin 3 has probability of heads $\frac{2}{3}$.
3. Test the null hypothesis that the 3 coins have the same probability of heads.

Answer

1. 3 degree of freedom, $V = 11.78$, $P = 0.008$.
2. 3 degree of freedom, $V = 1.283$, $P = 0.733$.
3. 2 degree of freedom, $V = 2.301$, $P = 0.316$.

A die is thrown 240 times, yielding the data in the following table:

Score	1	2	3	4	5	6
Frequency	57	39	28	28	36	52

1. Test the null hypothesis that the die is fair.
2. Test the null hypothesis that the die is an ace-six flat die (faces 1 and 6 have probability $\frac{1}{4}$ each while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each).

Answer

1. 5 degree of freedom, $V = 18.45$, $P = 0.0024$
2. 5 degree of freedom, $V = 5.383$, $P = 0.3709$

Two dice are thrown, yielding the data in the following table:

Score	1	2	3	4	5	6
Die 1	22	17	22	13	22	24
Die 2	44	24	19	19	18	36

1. Test the null hypothesis that die 1 is fair and die 2 is an ace-six flat.
2. Test the null hypothesis that all the dice have have the same probability distribuiton.

Answer

1. 10 degree of freedom, $V = 6.2$, $P = 0.798$.
2. 5 degree of freedom, $V = 7.103$, $P = 0.213$.

A university classifies faculty by rank as *instructors*, *assistant professors*, *associate professors*, and *full professors*. The data, by faculty rank and gender, are given in the following contingency table. Test to see if faculty rank and gender are independent.

Faculty	Instructor	Assistant Professor	Associate Professor	Full Professor
Male	62	238	185	115
Female	118	122	123	37

Answer

3 degrees of freedom, $V = 70.111$, $P \approx 0$.

Data Analysis Exercises

The Buffon trial data set gives the results of 104 repetitions of Buffon's needle experiment. The number of crack crossings is 56. In theory, this data set should correspond to 104 Bernoulli trials with success probability $p = \frac{2}{\pi}$. Test to see if this is reasonable.

Answer

1 degree of freedom, $V = 4.332$, $P = 0.037$.

Test to see if the alpha emissions data come from a Poisson distribution.

Answer

We partition of \mathbb{N} into 17 subsets: $\{0, 1\}$, $\{x\}$ for $x \in \{2, 3, \dots, 16\}$ and $\{17, 18, \dots\}$ There are 15 degrees of freedom. The estimated Poisson parameter is 8.367, $V = 9.644$, $P = 0.842$.

Test to see if Michelson's velocity of light data come from a normal distribution.

Answer

Using the following partition of \mathbb{R} : $\{(-\infty, 750), [750, 775), [775, 800), [800, 825), [825, 850), [850, 875), [875, 900), [900, 925), [925, 950), [950, 975), [975, 1000), [1000, \infty)\}$ have 8 degrees of freedom, $V = 11.443$, $P = 0.178$.

Simulation Exercises

In the simulation exercises below, you will be able to explore the goodness of fit test empirically.

In the dice goodness of fit experiment, set the sampling distribution to fair, the sample size to 50, and the significance level to 0.1. Set the test distribution as indicated below and in each case, run the simulation 1000 times. In case (a), give the empirical estimate of the significance level of the test and compare with 0.1. In the other cases, give the empirical estimate of the power of the test. Rank the distributions in (b)-(d) in increasing order of apparent power. Do your results seem reasonable?

1. fair
2. ace-six flats
3. the symmetric, unimodal distribution
4. the distribution skewed right

In the dice goodness of fit experiment, set the sampling distribution to ace-six flats, the sample size to 50, and the significance level to 0.1. Set the test distribution as indicated below and in each case, run the simulation 1000 times. In case (a), give the empirical estimate of the significance level of the test and compare with 0.1. In the other cases, give the empirical estimate of the power of the test. Rank the distributions in (b)-(d) in increasing order of apparent power. Do your results seem reasonable?

1. fair

2. ace-six flats
3. the symmetric, unimodal distribution
4. the distribution skewed right

In the dice goodness of fit experiment, set the sampling distribution to the symmetric, unimodal distribution, the sample size to 50, and the significance level to 0.1. Set the test distribution as indicated below and in each case, run the simulation 1000 times. In case (a), give the empirical estimate of the significance level of the test and compare with 0.1. In the other cases, give the empirical estimate of the power of the test. Rank the distributions in (b)-(d) in increasing order of apparent power. Do your results seem reasonable?

1. the symmetric, unimodal distribution
2. fair
3. ace-six flats
4. the distribution skewed right

In the dice goodness of fit experiment, set the sampling distribution to the distribution skewed right, the sample size to 50, and the significance level to 0.1. Set the test distribution as indicated below and in each case, run the simulation 1000 times. In case (a), give the empirical estimate of the significance level of the test and compare with 0.1. In the other cases, give the empirical estimate of the power of the test. Rank the distributions in (b)-(d) in increasing order of apparent power. Do your results seem reasonable?

1. the distribution skewed right
2. fair
3. ace-six flats
4. the symmetric, unimodal distribution

Suppose that D_1 and D_2 are different distributions. Is the power of the test with sampling distribution D_1 and test distribution D_2 the same as the power of the test with sampling distribution D_2 and test distribution D_1 ? Make a conjecture based on your results in the previous three exercises.

In the dice goodness of fit experiment, set the sampling and test distributions to fair and the significance level to 0.05. Run the experiment 1000 times for each of the following sample sizes. In each case, give the empirical estimate of the significance level and compare with 0.05.

1. $n = 10$
2. $n = 20$
3. $n = 40$
4. $n = 100$

In the dice goodness of fit experiment, set the sampling distribution to fair, the test distributions to ace-six flats, and the significance level to 0.05. Run the experiment 1000 times for each of the following sample sizes. In each case, give the empirical estimate of the power of the test. Do the powers seem to be converging?

1. $n = 10$
2. $n = 20$
3. $n = 40$
4. $n = 100$

CHAPTER OVERVIEW

10: Geometric Models

In this chapter, we explore several problems in geometric probability. These problems are interesting, conceptually clear, and the analysis is relatively simple. Thus, they are good problems for the student of probability. In addition, Buffon's problems and Bertrand's problem are historically famous, and contributed significantly to the early development of probability theory.

[10.1: Buffon's Problems](#)

[10.2: Bertrand's Paradox](#)

[10.3: Random Triangles](#)

This page titled [10: Geometric Models](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

10.1: Buffon's Problems

Buffon's experiments are very old and famous random experiments, named after comte de Buffon. These experiments are considered to be among the first problems in geometric probability.

Buffon's Coin Experiment

Buffon's coin experiment consists of dropping a coin randomly on a floor covered with identically shaped tiles. The event of interest is that the coin crosses a crack between tiles. We will model Buffon's coin problem with square tiles of side length 1—assuming the side length is 1 is equivalent to taking the side length as the unit of measurement.

Assumptions

First, let us define the experiment mathematically. As usual, we will idealize the physical objects by assuming that the coin is a perfect circle with radius r and that the cracks between tiles are line segments. A natural way to describe the outcome of the experiment is to record the center of the coin relative to the center of the tile where the coin happens to fall. More precisely, we will construct coordinate axes so that the tile where the coin falls occupies the square $S = [-\frac{1}{2}, \frac{1}{2}]^2$.

Now when the coin is tossed, we will denote the center of the coin by $(X, Y) \in S$ so that S is our sample space and X and Y are our basic random variables. Finally, we will assume that $r < \frac{1}{2}$ so that it is at least possible for the coin to fall inside the square without touching a crack.




Figure 10.1.1: Buffon's floor

Next, we need to define an appropriate probability measure that describes our basic random vector (X, Y) . If the coin falls “randomly” on the floor, then it is natural to assume that (X, Y) is uniformly distributed on S . By definition, this means that

$$\mathbb{P}[(X, Y) \in A] = \frac{\text{area}(A)}{\text{area}(S)}, \quad A \subseteq S \quad (10.1.1)$$

Run Buffon's coin experiment with the default settings. Watch how the points seem to fill the sample space S in a uniform manner.

The Probability of a Crack Crossing

Our interest is in the probability of the event C that the coin crosses a crack.

The probability of a crack crossing is $\mathbb{P}(C) = 1 - (1 - 2r)^2$.

Proof




Figure 10.1.2: $\mathbb{P}(C)$ as a function of r

In Buffon's coin experiment, change the radius with the scroll bar and watch how the events C and C^c change. Run the experiment with various values of r and compare the physical experiment with the points in the scatterplot. Compare the relative frequency of C to the probability of C .

The convergence of the relative frequency of an event (as the experiment is repeated) to the probability of the event is a special case of the law of large numbers.

Solve Buffon's coin problem with rectangular tiles that have height h and width w .

Answer

$$1 - \frac{(h - 2r)(w - 2r)}{hw}, \quad r < \min \left\{ \frac{h}{2}, \frac{w}{2} \right\} \quad (10.1.2)$$

Solve Buffon's coin problem with equilateral triangular tiles that have side length 1.

Recall that random numbers are simulation of independent random variables, each with the standard uniform distribution, that is, the continuous uniform distribution on the interval $(0, 1)$.

Show how to simulate the center of the coin (X, Y) in Buffon's coin experiment using random numbers.

Answer

$X = U - \frac{1}{2}$, $Y = V - \frac{1}{2}$, where U and V are random numbers.

Buffon's Needle Problem

Buffon's needle experiment consists of dropping a needle on a hardwood floor. The main event of interest is that the needle crosses a crack between floorboards. Strangely enough, the probability of this event leads to a statistical estimate of the number π !

Assumptions

Our first step is to define the experiment mathematically. Again we idealize the physical objects by assuming that the floorboards are uniform and that each has width 1. We will also assume that the needle has length $L < 1$ so that the needle cannot cross more than one crack. Finally, we assume that the cracks between the floorboards and the needle are line segments.

When the needle is dropped, we want to record its orientation relative to the floorboard cracks. One way to do this is to record the angle X that the top half of the needle makes with the line through the center of the needle, parallel to the floorboards, and the distance Y from the center of the needle to the bottom crack. These will be the basic random variables of our experiment, and thus the sample space of the experiment is

$$S = [0, \pi) \times [0, 1) = \{(x, y) : 0 \leq x < \pi, 0 \leq y < 1\} \quad (10.1.3)$$




Figure 10.1.3: Buffon's needle problem

Again, our main modeling assumption is that the needle is tossed “randomly” on the floor. Thus, a reasonable mathematical assumption might be that the basic random vector (X, Y) is uniformly distributed over the sample space. By definition, this means that

$$\mathbb{P}[(X, Y) \in A] = \frac{\text{area}(A)}{\text{area}(S)}, \quad A \subseteq S \quad (10.1.4)$$

Run Buffon's needle experiment with the default settings and watch the outcomes being plotted in the sample space. Note how the points in the scatterplot seem to fill the sample space S in a uniform way.

The Probability of a Crack Crossing

Our main interest is in the event C that the needle crosses a crack between the floorboards.

The event C can be written in terms of the basic angle and distance variables as follows:

$$C = \left\{ Y < \frac{L}{2} \sin(X) \right\} \cup \left\{ Y > 1 - \frac{L}{2} \sin(X) \right\} \quad (10.1.5)$$

The curves $y = \frac{L}{2} \sin(x)$ and $y = 1 - \frac{L}{2} \sin(x)$ on the interval $0 \leq x < \pi$ are shown in blue in the scatterplot of Buffon's needle experiment, and hence event C is the union of the regions below the lower curve and above the upper curve. Thus, the needle crosses a crack precisely when a point falls in this region.

The probability of a crack crossing is $\mathbb{P}(C) = 2L/\pi$.

Proof




Figure 10.1.4: $\mathbb{P}(C)$ as a function of L

In the Buffon's needle experiment, vary the needle length L with the scroll bar and watch how the event C changes. Run the experiment with various values of L and compare the physical experiment with the points in the scatterplot. Compare the relative frequency of C to the probability of C .

The convergence of the relative frequency of an event (as the experiment is repeated) to the probability of the event is a special case of the law of large numbers.

Find the probabilities of the following events in Buffon's needle experiment. In each case, sketch the event as a subset of the sample space.

1. $\{0 < X < \pi/2, 0 < Y < 1/3\}$
2. $\{1/4 < Y < 2/3\}$
3. $\{X < Y\}$
4. $\{X + Y < 2\}$

Answer

1. $\frac{1}{6}$
2. $\frac{5}{12}$
3. $\frac{1}{2\pi}$
4. $\frac{3}{2\pi}$

The Estimate of π

Suppose that we run Buffon's needle experiment a large number of times. By the law of large numbers, the *proportion* of crack crossings should be about the same as the *probability* of a crack crossing. More precisely, we will denote the number of crack crossings in the first n runs by N_n . Note that N_n is a random variable for the compound experiment that consists of n replications of the basic needle experiment. Thus, if n is large, we should have $\frac{N_n}{n} \approx \frac{2L}{\pi}$ and hence

$$\pi \approx \frac{2Ln}{N_n} \quad (10.1.6)$$

This is Buffon's famous estimate of π . In the simulation of Buffon's needle experiment, this estimate is computed on each run and shown numerically in the second table and visually in a graph.

Run the Buffon's needle experiment with needle lengths $L \in \{0.3, 0.5, 0.7, 1\}$ In each case, watch the estimate of π as the simulation runs.

Let us analyze the estimation problem more carefully. On each run j we have an indicator variable I_j , where $I_j = 1$ if the needle crosses a crack on run j and $I_j = 0$ if the needle does not cross a crack on run j . These indicator variables are independent, and identically distributed, since we are assuming independent replications of the experiment. Thus, the sequence forms a Bernoulli trials process.

The number of crack crossings in the first n runs of the experiment is

$$N_n = \sum_{j=1}^n I_j \quad (10.1.7)$$

which has the binomial distribution with parameters n and $2L/\pi$.

The mean and variance of N_n are

1. $\mathbb{E}(N_n) = n \frac{2L}{\pi}$
2. $\text{var}(N_n) = n \frac{2L}{\pi} \left(1 - \frac{2L}{\pi}\right)$

With probability 1, $\frac{N_n}{2Ln} \rightarrow \frac{1}{\pi}$ as $n \rightarrow \infty$ and $\frac{2Ln}{N_n} \rightarrow \pi$ as $n \rightarrow \infty$.

Proof

a

These results follow from the strong law of large numbers.

Thus, we have two basic estimators: $\frac{N_n}{2Ln}$ as an estimator of $\frac{1}{\pi}$ and $\frac{2Ln}{N_n}$ as an estimator of π . The estimator of $\frac{1}{\pi}$ has several important statistical properties. First, it is *unbiased* since the expected value of the estimator is the parameter being estimated:

The estimator of $\frac{1}{\pi}$ is *unbiased*:

$$\mathbb{E}\left(\frac{N_n}{2Ln}\right) = \frac{1}{\pi} \quad (10.1.8)$$

Proof

This follows from the results above for the [binomial distribution](#) and properties of expected value.

Since this estimator is unbiased, the variance gives the mean square error:

$$\text{var}\left(\frac{N_n}{2Ln}\right) = \mathbb{E}\left[\left(\frac{N_n}{2Ln} - \frac{1}{\pi}\right)^2\right] \quad (10.1.9)$$

The mean square error of the estimator of $\frac{1}{\pi}$ is

$$\text{var}\left(\frac{N_n}{2Ln}\right) = \frac{\pi - 2L}{2Ln\pi^2} \quad (10.1.10)$$

The variance is a decreasing function of the needle length L .

Thus, the estimator of $\frac{1}{\pi}$ improves as the needle length increases. On the other hand, the estimator of π is biased; it tends to overestimate π :

The estimator of π is positively biased:

$$\mathbb{E}\left(\frac{2Ln}{N_n}\right) \geq \pi \quad (10.1.11)$$

Proof

Use Jensen's inequality.

The estimator of π also tends to improve as the needle length increases. This is not easy to see mathematically. However, you can see it empirically.

In the Buffon's needle experiment, run the simulation 5000 times each with $L = 0.3$, $L = 0.5$, $L = 0.7$, and $L = 0.9$. Note how well the estimator seems to work in each case.

Finally, we should note that as a practical matter, Buffon's needle experiment is not a very efficient method of approximating π . According to Richard Durrett, to estimate π to four decimal places with $L = \frac{1}{2}$ would require about 100 million tosses!

Run the Buffon's needle experiment until the estimates of π seem to be consistently correct to two decimal places. Note the number of runs required. Try this for needle lengths $L = 0.3$, $L = 0.5$, $L = 0.7$, and $L = 0.9$ and compare the results.

Show how to simulate the angle X and distance Y in Buffon's needle experiment using random numbers.

Answer

$X = \pi U$, $Y = V$, where U and V are random numbers.

Notes

Buffon's needle problem is essentially solved by *Monte-Carlo integration*. In general, *Monte-Carlo methods* use statistical sampling to approximate the solutions of problems that are difficult to solve analytically. The modern theory of Monte-Carlo methods began with Stanislaw Ulam, who used the methods on problems associated with the development of the hydrogen bomb.

The original needle problem has been extended in many ways, starting with Simon Laplace who considered a floor with rectangular tiles. Indeed, variations on the problem are active research problems even today.

Neil Weiss has pointed out that our computer simulation of Buffon's needle experiment is *circular*, in the sense the program assumes knowledge of π (you can see this from the [simulation result](#) above).

Try to write a computer algorithm for Buffon's needle problem, without assuming the value of π or any other transcendental numbers.

This page titled [10.1: Buffon's Problems](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

10.2: Bertrand's Paradox

Preliminaries

Statement of the Problem

Bertrand's problem is to find the probability that a “random chord” on a circle will be longer than the length of a side of the inscribed equilateral triangle. The problem is named after the French mathematician Joseph Louis Bertrand, who studied the problem in 1889.

It turns out, as we will see, that there are (at least) three answers to Bertrand's problem, depending on how one interprets the phrase “random chord”. The lack of a unique answer was considered a paradox at the time, because it was assumed (naively, in hindsight) that there should be a single *natural* answer.

Run Bertrand's experiment 100 times for each of the following models. Do not be concerned with the exact meaning of the models, but see if you can detect a difference in the behavior of the outcomes

1. Uniform distance
2. Uniform angle
3. Uniform endpoint

Mathematical Formulation

To formulate the problem mathematically, let us take $(0, 0)$ as the center of the circle and take the radius of the circle to be 1. These assumptions entail no loss of generality because they amount to measuring distances relative to the center of the circle, and taking the radius of the circle as the unit of length. Now consider a chord on the circle. By rotating the circle, we can assume that one point of the chord is $(1, 0)$ and the other point is (X, Y) where $Y > 0$ and $X^2 + Y^2 = 1$.

With these assumptions, the chord is completely specified by giving any one of the following variables

1. The (perpendicular) distance D from the center of the circle to the midpoint of the chord. Note that $0 \leq D \leq 1$.
2. The angle A between the x -axis and the line from the center of the circle to the midpoint of the chord. Note that $0 \leq A \leq \pi/2$.
3. The horizontal coordinate X . Note that $-1 \leq X \leq 1$.




Figure 10.2.1: A chord in the circle

The variables are related as follows:

1. $D = \cos(A)$
2. $X = 2D^2 - 1$
3. $Y = 2D\sqrt{1 - D^2}$

The inverse relations are given below. Note again that there are one-to-one correspondences between X , A , and D .

1. $A = \arccos(D)$
2. $D = \sqrt{\frac{1}{2}(x + 1)}$
3. $D = \sqrt{\frac{1}{2} \pm \frac{1}{2}\sqrt{1 - y^2}}$

If the chord is generated in a probabilistic way, D , A , X , and Y become random variables. In light of the previous results, specifying the distribution of any of the variables D , A , or X completely determines the distribution of all four variables.

The angle A is also the angle between the chord and the tangent line to the circle at $(1, 0)$.

Now consider the equilateral triangle inscribed in the circle so that one of the vertices is $(1, 0)$. Consider the chord defined by the upper side of the triangle.

For this chord, the angle, distance, and coordinate variables are given as follows:

1. $a = \pi/3$
2. $d = 1/2$
3. $x = -1/2$
4. $y = \sqrt{3}/2$




Figure 10.2.2: The inscribed equilateral triangle

Now suppose that a chord is chosen in probabilistic way.

The length of the chord is greater than the length of a side of the inscribed equilateral triangle if and only if the following equivalent conditions occur:

1. $0 < D < 1/2$
2. $\pi/3 < A < \pi/2$
3. $-1 < X < -1/2$

Models

When an object is generated “at random”, a sequence of “natural” variables that determines the object should be given an appropriate uniform distribution. The coordinates of the coin center are such a sequence in Buffon's coin experiment; the angle and distance variables are such a sequence in Buffon's needle experiment. The crux of Bertrand's paradox is the fact that the distance D , the angle A , and the coordinate X each seems to be a natural variable that determine the chord, but different models are obtained, depending on which is given the uniform distribution.

The Model with Uniform Distance

Suppose that D is uniformly distributed on the interval $[0, 1]$.

The solution of Bertrand's problem is

$$\mathbb{P}\left(D < \frac{1}{2}\right) = \frac{1}{2} \quad (10.2.1)$$

In Bertrand's experiment, select the uniform distance model. Run the experiment 1000 times and compare the relative frequency function of the chord event to the true probability.

The angle A has probability density function

$$g(a) = \sin(a), \quad 0 < a < \frac{\pi}{2} \quad (10.2.2)$$

Proof

This follows from the standard the change of variables formula.

The coordinate X has probability density function

$$h(x) = \frac{1}{\sqrt{8(x+1)}}, \quad -1 < x < 1 \quad (10.2.3)$$

Proof

This follows from the standard the change of variables formula.

Note that A and X do not have uniform distributions. Recall that a *random number* is a simulation of a variable with the standard uniform distribution, that is the continuous uniform distribution on the interval $[0, 1]$.

Show how to simulate D , A , X , and Y using a random number.

Answer

$A = \arccos(D)$, $X = 2D^2 - 1$, $Y = 2D\sqrt{1 - D^2}$, where D is a random number

The Model with Uniform Angle

Suppose that A is uniformly distributed on the interval $(0, \pi/2)$.

The solution of Bertrand's problem is

$$\mathbb{P}\left(A > \frac{\pi}{3}\right) = \frac{1}{3} \quad (10.2.4)$$

In Bertrand's experiment, select the uniform angle model. Run the experiment 1000 times and compare the relative frequency function of the chord event to the true probability.

The distance D has probability density function

$$f(d) = \frac{2}{\pi\sqrt{1-d^2}}, \quad 0 < d < 1 \quad (10.2.5)$$

Proof

This follows from the standard change of variables formula.

The coordinate X has probability density function

$$h(x) = \frac{1}{\pi\sqrt{1-x^2}}, \quad -1 < x < 1 \quad (10.2.6)$$

Proof

This follows from the change of variables formula.

Note that D and X do not have uniform distributions.

Show how to simulate D , A , X , and Y using a random number.

Answer

$A = \frac{\pi}{2}U$, $D = \cos(A)$, $X = 2D^2 - 1$, $Y = 2D\sqrt{1 - D^2}$, where U is a random number.

The Model with Uniform Endpoint

Suppose that X is uniformly distributed on the interval $(-1, 1)$.

The solution of Bertrand's problem is

$$\mathbb{P}\left(-1 < X < -\frac{1}{2}\right) = \frac{1}{4} \quad (10.2.7)$$

In Bertrand's experiment, select the uniform endpoint model. Run the experiment 1000 times and compare the relative frequency function of the chord event to the true probability.

The distance D has probability density function

$$f(d) = 2d, \quad 0 < d < 1 \quad (10.2.8)$$

Proof

This follows from the change of variables formula.

The angle A has probability density function

$$g(a) = 2 \sin(a) \cos(a), \quad 0 < a < \frac{\pi}{2} \quad (10.2.9)$$

Proof

This follows from the change of variables formula.

Note that D and A do not have uniform distributions; in fact, D has a beta distribution with left parameter 2 and right parameter 1.

Physical Experiments

Suppose that a random chord is generated by tossing a coin of radius 1 on a table ruled with parallel lines that are distance 2 apart. Which of the models (if any) would apply to this physical experiment?

Answer

Uniform distance

Suppose that a needle is attached to the edge of disk of radius 1. A random chord is generated by spinning the needle. Which of the models (if any) would apply to this physical experiment?

Answer

Uniform angle

Suppose that a thin trough is constructed on the edge of a disk of radius 1. Rolling a ball in the trough generates a random point on the circle, so a random chord is generated by rolling the ball twice. Which of the models (if any) would apply to this physical experiment?

Answer

Uniform angle

This page titled [10.2: Bertrand's Paradox](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

10.3: Random Triangles

Preliminaries

Statement of the Problem

Suppose that a stick is randomly broken in two places. What is the probability that the three pieces form a triangle?

Without looking below, make a guess.

Run the triangle experiment 50 times. Do not be concerned with all of the information displayed in the app, but just note whether the pieces form a triangle. Would you like to revise your guess?

Mathematical Formulation

As usual, the first step is to model the random experiment mathematically. We will take the length of the stick as our unit of length, so that we can identify the stick with the interval $[0, 1]$. To break the stick into three pieces, we just need to select two points in the interval. Thus, let X denote the first point chosen and Y the second point chosen. Note that X and Y are random variables and hence the sample space of our experiment is $S = [0, 1]^2$. Now, to model the statement that the points are chosen *at random*, let us assume, as in the previous sections, that X and Y are independent and each is uniformly distributed on $[0, 1]$.

The random point (X, Y) is uniformly distributed on $S = [0, 1]^2$.

Hence

$$\mathbb{P}[(X, Y) \in A] = \frac{\text{area}(A)}{\text{area}(S)} \quad (10.3.1)$$

Triangles

The Probability of a Triangle

The three pieces form a triangle if and only if the *triangle inequalities* hold: the sum of the lengths of any two pieces must be greater than the length of the third piece.

The event that the pieces form a triangle is $T_1 \cup T_2$ where

1. $T_1 = \{(x, y) \in S : y > \frac{1}{2}, x < \frac{1}{2}, y - x < \frac{1}{2}\}$
2. $T_2 = \{(x, y) \in S : x > \frac{1}{2}, y < \frac{1}{2}, x - y < \frac{1}{2}\}$

A sketch of the event T is given below. Curiously, T is composed of triangles!




Figure 10.3.1: The event $T = T_1 \cup T_2$ that the pieces form a triangle

The probability that the pieces form a triangle is $\mathbb{P}(T) = \frac{1}{4}$.

How close did you come with your initial guess? The relative low value of $\mathbb{P}(T)$ is a bit surprising.

Run the triangle experiment 1000 times and compare the empirical probability of T^c to the true probability.

Triangles of Different Types

Now let us compute the probability that the pieces form a triangle of a given type. Recall that in an *acute* triangle all three angles are less than 90° , while an *obtuse* triangle has one angle (and only one) that is greater than 90° . A *right* triangle, of course, has one 90° angle.

Suppose that a triangle has side lengths a , b , and c , where c is the largest of these. The triangle is

1. acute if and only if $c^2 < a^2 + b^2$.
2. obtuse if and only if $c^2 > a^2 + b^2$.
3. right if and only if $c^2 = a^2 + b^2$.

Part (c), of course, is the famous Pythagorean theorem, named for the ancient Greek mathematician Pythagoras.

The right triangle equations for the stick pieces are

1. $(y - x)^2 = x^2 + (1 - y)^2$ in T_1
2. $(1 - x)^2 = x^2 + (y - x)^2$ in T_1
3. $x^2 = (y - x)^2 + (1 - y)^2$ in T_1
4. $(x - y)^2 = y^2 + (1 - x)^2$ in T_2
5. $(1 - x)^2 = y^2 + (x - y)^2$ in T_2
6. $y^2 = (x - y)^2 + (1 - x)^2$ in T_2




Figure 10.3.2: The events that the pieces form acute and obtuse triangles

Let R denote the event that the pieces form a right triangle. Then $\mathbb{P}(R) = 0$.

The event that the pieces form an acute triangle is $A = A_1 \cup A_2$ where

1. A_1 is the region inside curves (a), (b), and (c) of the [right triangle equations](#).
2. A_2 is the region inside curves (d), (e), and (f) of the [right triangle equations](#).

The event that the pieces form an obtuse triangle is $B = B_1 \cup B_2 \cup B_3 \cup B_4 \cup B_5 \cup B_6$ where

1. B_1 , B_2 , and B_3 are the regions inside T_1 and outside of curves (a), (b), and (c) of the [right triangle equations](#), respectively.
2. B_4 , B_5 , and B_6 are the regions inside T_2 and outside of curves (d), (e), and (f) of the [right triangle equations](#), respectively.

The probability that the pieces form an obtuse triangle is

$$\mathbb{P}(B) = \frac{9}{4} - 3 \ln(2) \approx 0.1706 \quad (10.3.2)$$

Proof

Simple calculus shows that $\mathbb{P}(B_i) = 3/8 - \ln(2)/2$ for each $i \in \{1, 2, 3, 4, 5, 6\}$. For example

$$\mathbb{P}(B_1) = \int_0^{1/2} \frac{x(1-2x)}{2-2x} dx \quad (10.3.3)$$

$$\mathbb{P}(B_3) = \int_0^{1/2} \left(y + \frac{1}{2y} - \frac{3}{2} \right) dy \quad (10.3.4)$$

From symmetry it also follows that $\mathbb{P}(B_i)$ is the same for each i .

The probability that the pieces form an acute triangle is

$$\mathbb{P}(A) = 3 \ln(2) - 2 \approx 0.07944 \quad (10.3.5)$$

Proof

Note that $A \cup B \cup R = T$, and A , B , and R are disjoint.

Run the triangle experiment 1000 times and compare the empirical probabilities to the true probabilities.

This page titled [10.3: Random Triangles](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

11: Bernoulli Trials

The Bernoulli trials process is one of the simplest, yet most important, of all random processes. It is an essential topic in any course in probability or mathematical statistics. The process consists of independent trials with two outcomes and with constant probabilities from trial to trial. Thus it is the mathematical abstraction of coin tossing. The process leads to several important probability distributions: the binomial, geometric, and negative binomial.

[11.1: Introduction to Bernoulli Trials](#)

[11.2: The Binomial Distribution](#)

[11.3: The Geometric Distribution](#)

[11.4: The Negative Binomial Distribution](#)

[11.5: The Multinomial Distribution](#)

[11.6: The Simple Random Walk](#)

[11.7: The Beta-Bernoulli Process](#)

This page titled [11: Bernoulli Trials](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.1: Introduction to Bernoulli Trials

Basic Theory

Definition

The *Bernoulli trials process*, named after Jacob Bernoulli, is one of the simplest yet most important random processes in probability. Essentially, the process is the mathematical abstraction of coin tossing, but because of its wide applicability, it is usually stated in terms of a sequence of generic “trials”.

A sequence of *Bernoulli trials* satisfies the following assumptions:

1. Each trial has two possible outcomes, in the language of reliability called *success* and *failure*.
2. The trials are independent. Intuitively, the outcome of one trial has no influence over the outcome of another trial.
3. On each trial, the probability of success is p and the probability of failure is $1 - p$ where $p \in [0, 1]$ is the *success parameter* of the process.

Random Variables

Mathematically, we can describe the Bernoulli trials process with a sequence of indicator random variables:

$$\mathbf{X} = (X_1, X_2, \dots) \quad (11.1.1)$$

An indicator variable is a random variable that takes only the values 1 and 0, which in this setting denote success and failure, respectively. Indicator variable X_i simply records the outcome of trial i . Thus, the indicator variables are independent and have the same probability density function:

$$\mathbb{P}(X_i = 1) = p, \quad \mathbb{P}(X_i = 0) = 1 - p \quad (11.1.2)$$

The distribution defined by this probability density function is known as the *Bernoulli distribution*. In statistical terms, the Bernoulli trials process corresponds to sampling from the Bernoulli distribution. In particular, the first n trials (X_1, X_2, \dots, X_n) form a random sample of size n from the Bernoulli distribution. Note again that the Bernoulli trials process is characterized by a single parameter p .

The joint probability density function of (X_1, X_2, \dots, X_n) trials is given by

$$f_n(x_1, x_2, \dots, x_n) = p^{x_1 + x_2 + \dots + x_n} (1 - p)^{n - (x_1 + x_2 + \dots + x_n)}, \quad (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \quad (11.1.3)$$

Proof

This follows from the basic assumptions of independence and the constant probabilities of 1 and 0.

Note that the exponent of p in the probability density function is the number of successes in the n trials, while the exponent of $1 - p$ is the number of failures.

If $\mathbf{X} = (X_1, X_2, \dots)$ is a Bernoulli trials process with parameter p then $\mathbf{1} - \mathbf{X} = (1 - X_1, 1 - X_2, \dots)$ is a Bernoulli trials sequence with parameter $1 - p$.

Suppose that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent random variables, each with the uniform distribution on the interval $[0, 1]$. For $p \in [0, 1]$ and $i \in \mathbb{N}_+$, let $X_i(p) = \mathbf{1}(U_i \leq p)$. Then $\mathbf{X}(p) = (X_1(p), X_2(p), \dots)$ is a Bernoulli trials process with probability p .

Note that in the previous result, the Bernoulli trials processes for all possible values of the parameter p are defined on a common probability space. This type of construction is sometimes referred to as *coupling*. This result also shows how to simulate a Bernoulli trials process with random numbers. All of the other random process studied in this chapter are functions of the Bernoulli trials sequence, and hence can be simulated as well.

Moments

Let X be an indicator variable with $\mathbb{P}(X = 1) = p$, where $p \in [0, 1]$. Thus, X is the result of a generic Bernoulli trial and has the Bernoulli distribution with parameter p . The following results give the mean, variance and some of the higher moments. A helpful fact is that if we take a positive power of an indicator variable, nothing happens; that is, $X^n = X$ for $n > 0$

The mean and variance of X are

1. $\mathbb{E}(X) = p$
2. $\text{var}(X) = p(1 - p)$

Proof

1. $\mathbb{E}(X) = 1 \cdot p(0 \cdot (1 - p)) = p$
2. $\text{var}(X) = \mathbb{E}(X^2) - [\mathbb{E}(X)]^2 = \mathbb{E}(X) - [\mathbb{E}(X)]^2 = p - p^2$

Note that the graph of $\text{var}(X)$, as a function of $p \in [0, 1]$ is a parabola opening downward. In particular the largest value is $\frac{1}{4}$ when $p = \frac{1}{2}$, and the smallest value is 0 when $p = 0$ or $p = 1$. Of course, in the last two cases, X is deterministic, taking the single value 0 when $p = 0$ and the single value 1 when $p = 1$




Figure 11.1.1: $\text{var}(X)$ as a function of p

Suppose that $p \in (0, 1)$. The skewness and kurtosis of X are

1. $\text{skew}(X) = \frac{1 - 2p}{\sqrt{p(1 - p)}}$
2. $\text{kurt}(X) = -3 + \frac{1}{p(1 - p)}$

The probability generating function of X is $P(t) = \mathbb{E}(t^X) = (1 - p) + pt$ for $t \in \mathbb{R}$.

Examples and Applications

Coins

As we noted earlier, the most obvious example of Bernoulli trials is coin tossing, where success means heads and failure means tails. The parameter p is the probability of heads (so in general, the coin is biased).

In the basic coin experiment, set $n = 100$ and For each $p \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ run the experiment and observe the outcomes.

Generic Examples

In a sense, the most general example of Bernoulli trials occurs when an experiment is replicated. Specifically, suppose that we have a basic random experiment and an event of interest A . Suppose now that we create a compound experiment that consists of independent replications of the basic experiment. Define success on trial i to mean that event A occurred on the i th run, and define failure on trial i to mean that event A did not occur on the i th run. This clearly defines a Bernoulli trials process with parameter $p = \mathbb{P}(A)$.

Bernoulli trials are also formed when we sample from a dichotomous population. Specifically, suppose that we have a population of two types of objects, which we will refer to as type 0 and type 1. For example, the objects could be persons, classified as *male* or *female*, or the objects could be components, classified as *good* or *defective*. We select n objects *at random* from the population; by definition, this means that each object in the population at the time of the draw is equally likely to be chosen. If the sampling is *with replacement*, then each object drawn is replaced before the next draw. In this case, successive draws are independent, so the

types of the objects in the sample form a sequence of Bernoulli trials, in which the parameter p is the proportion of type 1 objects in the population. If the sampling is *without replacement*, then the successive draws are dependent, so the types of the objects in the sample do not form a sequence of Bernoulli trials. However, if the population size is large compared to the sample size, the dependence caused by not replacing the objects may be negligible, so that for all practical purposes, the types of the objects in the sample can be treated as a sequence of Bernoulli trials. Additional discussion of sampling from a dichotomous population is in the chapter Finite Sampling Models.

Suppose that a student takes a multiple choice test. The test has 10 questions, each of which has 4 possible answers (only one correct). If the student blindly guesses the answer to each question, do the questions form a sequence of Bernoulli trials? If so, identify the trial outcomes and the parameter p .

Answer

Yes, probably so. The outcomes are *correct* and *incorrect* and $p = \frac{1}{4}$.

Candidate A is running for office in a certain district. Twenty persons are selected at random from the population of registered voters and asked if they prefer candidate A . Do the responses form a sequence of Bernoulli trials? If so identify the trial outcomes and the meaning of the parameter p .

Answer

Yes, approximately, assuming that the number of registered voters is large, compared to the sample size of 20. The outcomes are *prefer A* and *do not prefer A*; p is the proportion of voters in the entire district who prefer A .

An American roulette wheel has 38 slots; 18 are red, 18 are black, and 2 are green. A gambler plays roulette 15 times, betting on red each time. Do the outcomes form a sequence of Bernoulli trials? If so, identify the trial outcomes and the parameter p .

Answer

Yes, the outcomes are *red* and *black*, and $p = \frac{18}{38}$.

Roulette is discussed in more detail in the chapter on Games of Chance.

Two tennis players play a set of 6 games. Do the games form a sequence of Bernoulli trials? If so, identify the trial outcomes and the meaning of the parameter p .

Answer

No, probably not. The games are almost certainly dependent, and the win probably depends on who is serving and thus is not constant from game to game.

Reliability

Recall that in the standard model of structural reliability, a system is composed of n components that operate independently of each other. Let X_i denote the state of component i , where 1 means working and 0 means failure. If the components are all of the same type, then our basic assumption is that the *state vector*

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (11.1.4)$$

is a sequence of Bernoulli trials. The state of the system, (again where 1 means working and 0 means failed) depends only on the states of the components, and thus is a random variable

$$Y = s(X_1, X_2, \dots, X_n) \quad (11.1.5)$$

where $s: \{0, 1\}^n \rightarrow \{0, 1\}$ is the *structure function*. Generally, the probability that a device is working is the *reliability* of the device, so the parameter p of the Bernoulli trials sequence is the common reliability of the components. By independence, the *system reliability* r is a function of the component reliability:

$$r(p) = \mathbb{P}_p(Y = 1), \quad p \in [0, 1] \quad (11.1.6)$$

where we are emphasizing the dependence of the probability measure \mathbb{P} on the parameter p . Appropriately enough, this function is known as the *reliability function*. Our challenge is usually to find the reliability function, given the structure function.

A *series system* is working if and only if each component is working.

1. The state of the system is $Y = X_1 X_2 \cdots X_n = \min\{X_1, X_2, \dots, X_n\}$.
2. The reliability function is $r(p) = p^n$ for $p \in [0, 1]$.

A *parallel system* is working if and only if at least one component is working.

1. The state of the system is $Y = 1 - (1 - X_1)(1 - X_2) \cdots (1 - X_n) = \max\{X_1, X_2, \dots, X_n\}$.
2. The reliability function is $r(p) = 1 - (1 - p)^n$ for $p \in [0, 1]$.

Recall that in some cases, the system can be represented as a *graph* or *network*. The edges represent the components and the vertices the connections between the components. The system functions if and only if there is a working path between two designated vertices, which we will denote by a and b .

Find the reliability of the *Wheatstone bridge network* shown below (named for Charles Wheatstone).




Figure 11.1.2: The Wheatstone bridge network

Answer

$$r(p) = p(2p - p^2)^2 + (1 - p)(2p^2 - p^4)$$

The Pooled Blood Test

Suppose that each person in a population, independently of all others, has a certain disease with probability $p \in (0, 1)$. Thus, with respect to the disease, the persons in the population form a sequence of Bernoulli trials. The disease can be identified by a blood test, but of course the test has a cost.

For a group of k persons, we will compare two strategies. The first is to test the k persons individually, so that of course, k tests are required. The second strategy is to *pool* the blood samples of the k persons and test the pooled sample first. We assume that the test is negative if and only if all k persons are free of the disease; in this case just one test is required. On the other hand, the test is positive if and only if at least one person has the disease, in which case we then have to test the persons individually; in this case $k + 1$ tests are required. Thus, let Y denote the number of tests required for the pooled strategy.

The number of tests Y has the following properties:

1. $\mathbb{P}(Y = 1) = (1 - p)^k$, $\mathbb{P}(Y = k + 1) = 1 - (1 - p)^k$
2. $\mathbb{E}(Y) = 1 + k[1 - (1 - p)^k]$
3. $\text{var}(Y) = k^2(1 - p)^k[1 - (1 - p)^k]$

In terms of expected value, the pooled strategy is better than the basic strategy if and only if

$$p < 1 - \left(\frac{1}{k}\right)^{1/k} \quad (11.1.7)$$

The graph of the critical value $p_k = 1 - (1/k)^{1/k}$ as a function of $k \in [2, 20]$ is shown in the graph below:




Figure 11.1.3: The graph of p_k as a function of k

The critical value p_k satisfies the following properties:

1. The maximum value of p_k occurs at $k = 3$ and $p_3 \approx 0.307$.
2. $p_k \rightarrow 0$ as $k \rightarrow \infty$.

It follows that if $p \geq 0.307$, pooling never makes sense, regardless of the size of the group k . At the other extreme, if p is very small, so that the disease is quite rare, pooling is better unless the group size k is very large.

Now suppose that we have n persons. If $k \mid n$ then we can partition the population into n/k groups of k each, and apply the pooled strategy to each group. Note that $k = 1$ corresponds to individual testing, and $k = n$ corresponds to the pooled strategy on the entire population. Let Y_i denote the number of tests required for group i .

The random variables $(Y_1, Y_2, \dots, Y_{n/k})$ are independent and each has the distribution given [above](#).

The total number of tests required for this partitioning scheme is $Z_{n,k} = Y_1 + Y_2 + \dots + Y_{n/k}$.

The expected total number of tests is

$$\mathbb{E}(Z_{n,k}) = \begin{cases} n, & k = 1 \\ n \left[\left(1 + \frac{1}{k}\right) - (1-p)^k \right], & k > 1 \end{cases} \quad (11.1.8)$$

The variance of the total number of tests is

$$\text{var}(Z_{n,k}) = \begin{cases} 0, & k = 1 \\ n k (1-p)^k [1 - (1-p)^k], & k > 1 \end{cases} \quad (11.1.9)$$

Thus, in terms of expected value, the optimal strategy is to group the population into n/k groups of size k , where k minimizes the [expected value function](#) above. It is difficult to get a closed-form expression for the optimal value of k , but this value can be determined numerically for specific n and p .

For the following values of n and p , find the optimal pooling size k and the expected number of tests. (Restrict your attention to values of k that divide n .)

1. $n = 100, p = 0.01$
2. $n = 1000, p = 0.05$
3. $n = 1000, p = 0.001$

Answer

1. $k = 10, \mathbb{E}(Y_k) = 19.56$
2. $k = 5, \mathbb{E}(Y_k) = 426.22$
3. $k = 40, \mathbb{E}(Y_k) = 64.23$

If k does not divide n , then we could divide the population of n persons into $\lfloor n/k \rfloor$ groups of k each and one “remainder” group with $n \bmod k$ members. This clearly complicates the analysis, but does not introduce any new ideas, so we will leave this extension to the interested reader.

This page titled [11.1: Introduction to Bernoulli Trials](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.2: The Binomial Distribution

Basic Theory

Definitions

Our random experiment is to perform a sequence of Bernoulli trials $\mathbf{X} = (X_1, X_2, \dots)$. Recall that \mathbf{X} is a sequence of independent, identically distributed indicator random variables, and in the usual language of reliability, 1 denotes success and 0 denotes failure. The common probability of success $p = \mathbb{P}(X_i = 1)$, is the basic parameter of the process. In statistical terms, the first n trials (X_1, X_2, \dots, X_n) form a random sample of size n from the Bernoulli distribution.

In this section we will study the random variable that gives the number of successes in the first n trials and the random variable that gives the proportion of successes in the first n trials. The underlying distribution, the binomial distribution, is one of the most important in probability theory, and so deserves to be studied in considerable detail. As you will see, some of the results in this section have two or more proofs. In almost all cases, note that the proof from Bernoulli trials is the simplest and most elegant.

For $n \in \mathbb{N}$, the number of successes in the first n trials is the random variable

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (11.2.1)$$

The distribution of Y_n is the *binomial distribution* with trial parameter n and success parameter p .

Note that $\mathbf{Y} = (Y_0, Y_1, \dots)$ is the partial sum process associated with the Bernoulli trials sequence \mathbf{X} . In particular, $Y_0 = 0$, so point mass at 0 is considered a degenerate form of the binomial distribution.

Distribution Functions

The probability density function f_n of Y_n is given by

$$f_n(y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (11.2.2)$$

Proof

Recall that if $(x_1, x_2, \dots) \in \{0, 1\}^n$ with $\sum_{i=1}^n x_i = y$ (that is, a bit string of length n with 1 occurring exactly y times), then by independence,

$$\mathbb{P}[(X_1, X_2, \dots, X_n) = (x_1, x_2, \dots, x_n)] = p^y (1-p)^{n-y} \quad (11.2.3)$$

Moreover, the number of bit strings of length n with 1 occurring exactly y times is the binomial coefficient $\binom{n}{y}$. By the additive property of probability

$$\mathbb{P}(Y_n = y) = \binom{n}{y} p^y (1-p)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (11.2.4)$$

Check that f_n is a valid PDF

Clearly $f_n(y) \geq 0$ for $y \in \{0, 1, \dots, n\}$. From the binomial theorem

$$\sum_{y=0}^n f_n(y) = \sum_{y=0}^n \binom{n}{y} p^y (1-p)^{n-y} = [p + (1-p)]^n = 1 \quad (11.2.5)$$

In the binomial coin experiment, vary n and p with the scrollbars, and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the relative frequency function to the probability density function.

The binomial distribution is unimodal: For $k \in \{1, 2, \dots, n\}$,

1. $f_n(k) > f_n(k-1)$ if and only if $k < (n+1)p$.
2. $f_n(k) = f_n(k-1)$ if and only if $k = (n+1)p$ is an integer between 1 and n .

Proof

1. $f_n(k) > f_n(k-1)$ if and only if $\binom{n}{k} p^k (1-p)^{n-k} > \binom{n}{k-1} p^{k-1} (1-p)^{n-k+1}$ if and only if $\frac{p}{k} > \frac{1-p}{n-k+1}$ if and only if $k < (n+1)p$
2. As in (a), $f_n(k) = f_n(k-1)$ if and only if $k = (n+1)p$, which must be an integer.

Thus, the density function at first increases and then decreases, reaching its maximum value at $\lfloor (n+1)p \rfloor$. This integer is a *mode* of the distribution. In the case that $m = (n+1)p$ is an integer between 1 and n , there are two consecutive modes, at $m-1$ and m .

Now let F_n denote the distribution function of Y_n , so that

$$F_n(y) = \mathbb{P}(Y_n \leq y) = \sum_{k=0}^y f_n(k) = \sum_{k=0}^y \binom{n}{k} p^k (1-p)^{n-k}, \quad y \in \{0, 1, \dots, n\} \quad (11.2.6)$$

The distribution function F_n and the quantile function F_n^{-1} do not have simple, closed forms, but values of these functions can be computed from mathematical and statistical software.

Open the special distribution calculator and select the binomial distribution and set the view to CDF. Vary n and p and note the shape and location of the distribution/quantile function. For various values of the parameters, compute the median and the first and third quartiles.

The binomial distribution function also has a nice relationship to the beta distribution function.

The distribution function F_n can be written in the form

$$F_n(k) = \frac{n!}{(n-k-1)!k!} \int_0^{1-p} x^{n-k-1} (1-x)^k dx, \quad k \in \{0, 1, \dots, n\} \quad (11.2.7)$$

Proof

Let $G_n(k)$ denote the expression on the right. Substitution and simple integration shows that $G_n(0) = (1-p)^n = f_n(0) = F_n(0)$. For $k \in \{1, 2, \dots, n\}$, integrating by parts with $u = (1-x)^k$ and $dv = x^{n-k-1} dx$ gives

$$G_n(k) = \binom{n}{k} p^k (1-p)^{n-k} + \frac{n!}{(n-k)!(k-1)!} \int_0^{1-p} x^{n-k} (1-x)^k dx = f_n(k) + G_n(k-1) \quad (11.2.8)$$

It follows that $G_n(k) = \sum_{j=0}^k f_n(j) = F_n(k)$ for $k \in \{0, 1, \dots, n\}$.

The expression on the right in the previous theorem is the beta distribution function, with left parameter $n-k$ and right parameter $k+1$, evaluated at $1-p$.

Moments

The mean, variance and other moments of the binomial distribution can be computed in several different ways. Again let $Y_n = \sum_{i=1}^n X_i$ where $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter p .

The mean and variance of Y_n are

1. $E(Y_n) = np$
2. $\text{var}(Y_n) = np(1-p)$

Proof from Bernoulli trials

1. Recall that $\mathbb{E}(X_i) = p$ for each i . Hence from the additive property of expected value,

$$\mathbb{E}(Y_n) = \sum_{i=1}^n \mathbb{E}(X_i) = \sum_{i=1}^n p = np \quad (11.2.9)$$

2. Recall that $\text{var}(X_i) = p(1-p)$ for each i . Hence from the additive property of variance for *independent* variables,

$$\text{var}(Y_n) = \sum_{i=1}^n \text{var}(X_i) = \sum_{i=1}^n p(1-p) = np(1-p) \quad (11.2.10)$$

Direct proof of (a)

Recall the identity $y \binom{n}{y} = n \binom{n-1}{y-1}$ for $n, y \in \mathbb{N}_+$. Using the binomial theorem,

$$\mathbb{E}(Y_n) = \sum_{y=0}^n y \binom{n}{y} p^y (1-p)^{n-y} = \sum_{y=1}^n y \binom{n}{y} p^y (1-p)^{n-y} \quad (11.2.11)$$

$$= \sum_{y=1}^n n \binom{n-1}{y-1} p^y (1-p)^{n-y} = np \sum_{y=1}^n \binom{n-1}{y-1} p^{y-1} (1-p)^{(n-1)-(y-1)} \quad (11.2.12)$$

$$= np \sum_{k=0}^{n-1} \binom{n-1}{k} p^k (1-p)^{n-1-k} = np [p + (1-p)]^{n-1} = np \quad (11.2.13)$$

A similar, but more complicated proof can be used for part (b).

The expected value of Y_n also makes intuitive sense, since p should be approximately the proportion of successes in a large number of trials. We will discuss the point further in the subsection below on the [proportion of successes](#). Note that the graph of $\text{var}(Y_n)$ as a function of $p \in [0, 1]$ is parabola opening downward. In particular the maximum value of the variance is $n/4$ when $p = 1/2$, and the minimum value is 0 when $p = 0$ or $p = 1$. Of course, in the last two cases, Y_n is deterministic, taking just the value 0 if $p = 0$ and just the value n when $p = 1$.

In the binomial coin experiment, vary n and p with the scrollbars and note the location and size of the mean \pm standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

The probability generating function of Y_n is $P_n(t) = \mathbb{E}(t^{Y_n}) = [(1-p) + pt]^n$ for $t \in \mathbb{R}$.

Proof from Bernoulli trials

Recall that the probability generating function of a sum of independent variables is the product of the probability generating functions of the terms. Recall also that the PGF of X_i is $P(t) = \mathbb{E}(t^{X_i}) = (1-p) + pt$ for each i . Hence $P_n(t) = P^n(t)$.

Direct Proof

Using the binomial theorem yet again,

$$\mathbb{E}(t^{Y_n}) = \sum_{y=0}^n t^y \binom{n}{y} p^y (1-p)^{n-y} = \sum_{y=0}^n \binom{n}{y} (pt)^y (1-p)^{n-y} = [pt + (1-p)]^n \quad (11.2.14)$$

Recall that for $x \in \mathbb{R}$ and $k \in \mathbb{N}$, the *falling power* of x of order k is $x^{(k)} = x(x-1) \cdots (x-k+1)$. If X is a random variable and $k \in \mathbb{N}$, then $\mathbb{E}[X^{(k)}]$ is the *factorial moment* of X of order k . The probability generating function provides an easy way to compute the factorial moments of the binomial distribution.

$$\mathbb{E}[Y_n^{(k)}] = n^{(k)} p^k \text{ for } k \in \mathbb{N}.$$

Proof

Recall that $P_n^{(k)}(1) = \mathbb{E}[Y_n^{(k)}]$ where $P_n^{(k)}$ denotes the k th derivative of P_n . By simple calculus, $P_n^{(k)}(t) = n^{(k)} [(1-p) + pt]^{n-k} p^k$, so $P_n^{(k)} = n^{(k)} p^k$.

Our next result gives a recursion equation and initial conditions for the moments of the binomial distribution.

Recursion relation and initial conditions

1. $\mathbb{E}(Y_n^k) = np \mathbb{E}[(Y_{n-1} + 1)^{k-1}]$ for $n, k \in \mathbb{N}_+$
2. $\mathbb{E}(Y_n^0) = 1$ for $n \in \mathbb{N}$
3. $\mathbb{E}(Y_0^k) = 0$ for $k \in \mathbb{N}_+$

Proof

Recall again the identity $y \binom{n}{y} = n \binom{n-1}{y-1}$.

$$\mathbb{E}(Y_n^k) = \sum_{y=0}^n y^k \binom{n}{y} p^y (1-p)^{n-y} = \sum_{y=1}^n y^{k-1} y \binom{n}{y} p^y (1-p)^{n-y} \quad (11.2.15)$$

$$= \sum_{y=1}^n y^{k-1} n \binom{n-1}{y-1} p^y (1-p)^{n-y} = np \sum_{y=1}^n y^{k-1} \binom{n-1}{y-1} p^{y-1} (1-p)^{(n-1)-(y-1)} \quad (11.2.16)$$

$$= np \sum_{j=0}^{n-1} (j+1)^{k-1} \binom{n-1}{j} p^j (1-p)^{n-1-j} = np \mathbb{E}[(Y_{n-1} + 1)^{k-1}] \quad (11.2.17)$$

The ordinary (raw) moments of Y_n can be computed from the factorial moments and from the recursion relation. Here are the first four, which will be needed below.

The first four moments of Y_n are

1. $\mathbb{E}(Y_n) = np$
2. $\mathbb{E}(Y_n^2) = n^{(2)} p^2 + np$
3. $\mathbb{E}(Y_n^3) = n^{(3)} p^3 + 3n^{(2)} p^2 + np$
4. $\mathbb{E}(Y_n^4) = n^{(4)} p^4 + 6n^{(3)} p^3 + 7n^{(2)} p^2 + np$

Our final moment results gives the skewness and kurtosis of the binomial distribution.

For $p \in (0, 1)$, the skewness of Y_n is

$$\text{skew}(Y_n) = \frac{1-2p}{\sqrt{np(1-p)}} \quad (11.2.18)$$

1. $\text{skew}(Y_n) > 0$ if $p < \frac{1}{2}$, $\text{skew}(Y_n) < 0$ if $p > \frac{1}{2}$, and $\text{skew}(Y_n) = 0$ if $p = \frac{1}{2}$

2. For fixed n , $\text{skew}(Y_n) \rightarrow \infty$ as $p \downarrow 0$ and as $p \uparrow 1$
3. For fixed p , $\text{skew}(Y_n) \rightarrow 0$ as $n \rightarrow \infty$

Proof

These result follow from the standard computational formulas for skewness and kurtosis and the first three moments of the binomial distribution.

Open the binomial timeline experiment. For each of the following values of n , vary p from 0 to 1 and note the shape of the probability density function in light of the previous results on skewness.

1. $n = 10$
2. $n = 20$
3. $n = 100$

For $p \in (0, 1)$, the kurtosis of Y_n is

$$\text{kurt}(Y_n) = 3 - \frac{6}{n} + \frac{1}{np(1-p)} \quad (11.2.19)$$

1. For fixed n , $\text{kurt}(Y_n)$ decreases and then increases as a function of p , with minimum value $3 - \frac{2}{n}$ at the point of symmetry $p = \frac{1}{2}$
2. For fixed n , $\text{kurt}(Y_n) \rightarrow \infty$ as $p \downarrow 0$ and as $p \uparrow 1$
3. For fixed p , $\text{kurt}(Y_n) \rightarrow 3$ as $n \rightarrow \infty$

Proof

These result follow from the standard computational formulas for skewness and kurtosis and the first four moments of the binomial distribution.

Note that the *excess kurtosis* is $\text{kurt}(Y_n) - 3 = \frac{1}{np(1-p)} - \frac{6}{n} \rightarrow 0$ as $n \rightarrow \infty$. This is related to the [convergence of the binomial distribution to the normal](#), which we will discuss below.

The Partial Sum Process

Several important properties of the random process $\mathbf{Y} = (Y_0, Y_1, Y_2, \dots)$ stem from the fact that it is a partial sum process corresponding to the sequence $\mathbf{X} = (X_1, X_2, \dots)$ of independent, identically distributed indicator variables.

\mathbf{Y} has *stationary, independent increments*:

1. If m and n are positive integers with $m \leq n$ then $Y_n - Y_m$ has the same distribution as Y_{n-m} , namely binomial with parameters $n - m$ and p .
2. If $n_1 \leq n_2 \leq n_3 \leq \dots$ then $(Y_{n_1}, Y_{n_2} - Y_{n_1}, Y_{n_3} - Y_{n_2}, \dots)$ is a sequence of independent variables.

Proof

Every partial sum process corresponding to a sequence of independent, identically distributed variables has stationary, independent increments.

The following result gives the finite dimensional distributions of \mathbf{Y} .

The joint probability density functions of the sequence \mathbf{Y} are given as follows:

$$\mathbb{P}(Y_{n_1} = y_1, Y_{n_2} = y_2, \dots, Y_{n_k} = y_k) = \binom{n_1}{y_1} \binom{n_2 - n_1}{y_2 - y_1} \dots \binom{n_k - n_{k-1}}{y_k - y_{k-1}} p^{y_k} (1-p)^{n_k - y_k} \quad (11.2.20)$$

where $n_1, n_2, \dots, n_k \in \mathbb{N}_+$ with $n_1 < n_2 < \dots < n_k$ and where $y_1, y_2, \dots, y_k \in \mathbb{N}$ with $0 \leq y_j - y_{j-1} \leq n_j - n_{j-1}$ for each $j \in \{1, 2, \dots, k\}$.

Proof

From the stationary and independent increments properties,

$$\mathbb{P}(Y_{n_1} = y_1, Y_{n_2} = y_2, \dots, Y_{n_k} = y_k) = f_{n_1}(y_1) f_{n_2 - n_1}(y_2 - y_1) \dots f_{n_k - n_{k-1}}(y_k - y_{k-1}) \quad (11.2.21)$$

The result then follows from substitution and simplification.

Transformations that Preserve the Binomial Distribution

There are two simple but important transformations that preserve the binomial distribution.

If U is a random variable having the binomial distribution with parameters n and p , then $n - U$ has the binomial distribution with parameters n and $1 - p$.

Proof from Bernoulli trials

Recall that if (X_1, X_2, \dots) is a Bernoulli trials sequence with parameter p , then $(1 - X_1, 1 - X_2, \dots)$ is a Bernoulli trials sequence with parameter $1 - p$. Also U has the same distribution as $\sum_{i=1}^n X_i$ (binomial with parameters n and p) so $n - U$ has the same distribution as $\sum_{i=1}^n (1 - X_i)$ (binomial with parameters n and $1 - p$).

Proof from density functions

Note that $\mathbb{P}(n - U = k) = \mathbb{P}(U = n - k) = \binom{n}{k} p^{n-k} (1-p)^k$ for $k \in \{0, 1, \dots, n\}$

The sum of two independent binomial variables *with the same success parameter* also has a binomial distribution.

Suppose that U and V are independent random variables, and that U has the binomial distribution with parameters m and p , and V has the binomial distribution with parameters n and p . Then $U + V$ has the binomial distribution with parameters $m + n$ and p .

Proof from Bernoulli trials

Let (X_1, X_2, \dots) be a Bernoulli trials sequence with parameter p , and let $Y_k = \sum_{i=1}^k X_i$ for $k \in \mathbb{N}$. Then U has the same distribution as Y_m and V has the same distribution as $Y_{m+n} - Y_m$. Since Y_m and $Y_{m+n} - Y_m$ are independent, $U + V$ has the same distribution as $Y_m + (Y_{m+n} - Y_m) = Y_{m+n}$.

Proof from convolution powers

Let f denote the PDF of an indicator variable X with parameter p , so that $f(x) = p^x (1-p)^{1-x}$ for $x \in \{0, 1\}$. The binomial distribution with parameters $k \in \mathbb{N}_+$ and p has PDF $f_k = f^{*k}$, the k -fold convolution power of f . In particular, U has PDF f^{*m} , V has PDF f^{*n} and hence $U + V$ has PDF $f^{*m} * f^{*n} = f^{*(m+n)}$.

Proof from generating functions

U and V have PGFs $P_m(t) = (1 - p + pt)^m$ and $P_n(t) = (1 - p + pt)^n$ for $t \in \mathbb{R}$, respectively. Hence by independence, $U + V$ has PGF $P_m P_n = P_{m+n}$.

Sampling and the Hypergeometric Distribution

Suppose that we have a *dichotomous population*, that is a population of two types of objects. Specifically, suppose that we have m objects, and that r of the objects are type 1 and the remaining $m - r$ objects are type 0. Thus $m \in \mathbb{N}_+$ and $r \in \{0, 1, \dots, m\}$. We select n objects at random from the population, so that all samples of size n are equally likely. If the sampling is *with replacement*, the sample size n can be any positive integer. If the sampling is *without replacement*, then we must have $n \in \{1, 2, \dots, m\}$.

In either case, let X_i denote the type of the i 'th object selected for $i \in \{1, 2, \dots, n\}$ so that $Y = \sum_{i=1}^n X_i$ is the number of type 1 objects in the sample. As noted in the Introduction, if the sampling is with replacement, (X_1, X_2, \dots, X_n) is a sequence of Bernoulli trials, and hence Y has the binomial distribution parameters n and $p = r/m$. If the sampling is without replacement, then Y has the *hypergeometric distribution* with parameters m , r , and n . The hypergeometric distribution is studied in detail in the chapter on Finite Sampling Models. For reference, the probability density function of Y is given by

$$\mathbb{P}(Y = y) = \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}} = \binom{n}{y} \frac{r^{(y)} (m-r)^{(n-y)}}{m^{(n)}}, \quad y \in \{0, 1, \dots, n\} \quad (11.2.22)$$

and the mean and variance of Y are

$$\mathbb{E}(Y) = n \frac{r}{m}, \quad \text{var}(Y) = n \frac{r}{m} \left(1 - \frac{r}{m}\right) \frac{m-n}{m-1} \quad (11.2.23)$$

If the population size m is large compared to the sample size n , then the dependence between the indicator variables is slight, and so the hypergeometric distribution should be close to the binomial distribution. The following theorem makes this precise.

Suppose that $r_m \in \{0, 1, \dots, m\}$ for each $m \in \mathbb{N}_+$ and that $r_m/m \rightarrow p \in [0, 1]$ as $m \rightarrow \infty$. Then for fixed $n \in \mathbb{N}_+$ the hypergeometric distribution with parameters m , r_m and n converges to the binomial distribution with parameters n and p as $m \rightarrow \infty$.

Proof

The hypergeometric PDF has the form

$$g_m(y) = \binom{n}{y} \frac{r_m^{(y)} (m - r_m)^{(n-y)}}{m^{(n)}}, \quad y \in \{0, 1, \dots, n\} \quad (11.2.24)$$

Note that the fraction above has n factors in the numerator and n factors in the denominator. We can group these, in order, to form a product of n fractions. The first y fractions have the form

$$\frac{r_m - i}{m - i} \quad (11.2.25)$$

where $i \in \{0, 1, \dots, y-1\}$. Each of these converges to p as $m \rightarrow \infty$. The remaining $n - y$ fractions have the form

$$\frac{m - r_m - j}{m - y - j} \quad (11.2.26)$$

where $j \in \{0, 1, \dots, n - y - 1\}$. For fixed y and n , each of these converges to $1 - p$ as $m \rightarrow \infty$. Hence $g_m(y) \rightarrow \binom{n}{y} p^y (1-p)^{n-y}$ as $m \rightarrow \infty$ for each $y \in \{0, 1, \dots, n\}$

Under the conditions in the previous theorem, the mean and variance of the hypergeometric distribution converge to the mean and variance of the limiting binomial distribution:

1. $n \frac{r_m}{m} \rightarrow np$ as $m \rightarrow \infty$
2. $n \frac{r_m}{m} \left(1 - \frac{r_m}{m}\right) \frac{m-n}{m-1} \rightarrow np(1-p)$ as $m \rightarrow \infty$

Proof

By assumption $r_m/m \rightarrow p$ as $m \rightarrow \infty$ and n is fixed, so also $(m-n)/(m-1) \rightarrow 1$ as $n \rightarrow \infty$

In the ball and urn experiment, vary the parameters and switch between sampling without replacement and sampling with replacement. Note the difference between the graphs of the hypergeometric probability density function and the binomial probability density function. In particular, note the similarity when m is large and n small. For selected values of the parameters, and for both sampling modes, run the experiment 1000 times.

From a practical point of view, the convergence of the hypergeometric distribution to the binomial means that if the population size m is “large” compared to the sample size, then the hypergeometric distribution with parameters m , r and n is well approximated by the binomial distribution with parameters n and $p = r/m$. This is often a useful result, because the binomial distribution has fewer parameters than the hypergeometric distribution (and often in real problems, the parameters may only be known approximately). Specifically, in the approximating binomial distribution, we do not need to know the population size m and the number of type 1 objects r individually, but only in the ratio r/m . Generally, the approximation works well if m is large compared to n that $\frac{m-n}{m-1}$ is close to 1. This ensures that the variance of the hypergeometric distribution is close to the variance of the approximating binomial distribution.

Now let's return to our usual sequence of Bernoulli trials $\mathbf{X} = (X_1, X_2, \dots)$, with success parameter p , and to the binomial variables $Y_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}$. Our next result shows that given k successes in the first n trials, the trials on which the successes occur is simply a random sample of size k chosen without replacement from $\{1, 2, \dots, n\}$.

Suppose that $n \in \mathbb{N}_+$ and $k \in \{0, 1, \dots, n\}$. Then for $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ with $\sum_{i=1}^n x_i = k$,

$$\mathbb{P}[(X_1, X_2, \dots, X_n) = (x_1, x_2, \dots, x_n) \mid Y_n = k] = \frac{1}{\binom{n}{k}} \quad (11.2.27)$$

Proof

From the definition of conditional probability,

$$\mathbb{P}[(X_1, X_2, \dots, X_n) = (x_1, x_2, \dots, x_n) \mid Y_n = k] = \frac{\mathbb{P}[(X_1, X_2, \dots, X_n) = (x_1, x_2, \dots, x_n)]}{\mathbb{P}(Y_n = k)} = \frac{p^k(1-p)^{n-k}}{\binom{n}{k}p^k(1-p)^{n-k}} = \frac{1}{\binom{n}{k}} \quad (11.2.28)$$

Note in particular that the conditional distribution above does not depend on p . In statistical terms, this means that relative to (X_1, X_2, \dots, X_n) , random variable Y_n is a *sufficient statistic* for p . Roughly, Y_n contains all of the information about p that is available in the entire sample (X_1, X_2, \dots, X_n) . Sufficiency is discussed in more detail in the chapter on Point Estimation. Next, if $m \leq n$ then the conditional distribution of Y_m given $Y_n = k$ is hypergeometric, with population size n , type 1 size m , and sample size k .

Suppose that $p \in (0, 1)$ and that $m, n, k \in \mathbb{N}_+$ with $m \leq n$ and $k \leq n$. Then

$$\mathbb{P}(Y_m = j \mid Y_n = k) = \frac{\binom{m}{j} \binom{m-n}{k-j}}{\binom{n}{k}}, \quad j \in \{0, 1, \dots, k\} \quad (11.2.29)$$

Proof from the previous result

Given $Y_n = k$, the trial numbers of the successes form a random sample of size k chosen without replacement from $\{1, 2, \dots, n\}$. Designate trials $\{1, 2, \dots, m\}$ as type 1 and trials $\{m+1, m+2, \dots, n\}$ as type 0. Then Y_m is the number of type 1 trials in the sample, and hence (given $Y_n = k$) has the hypergeometric distribution with population size n , type 1 size m , and sample size k .

Direct proof

From the definition of conditional probability,

$$\mathbb{P}(Y_m = j \mid Y_n = k) = \frac{\mathbb{P}(Y_m = j, Y_n = k)}{\mathbb{P}(Y_n = k)} = \frac{\mathbb{P}(Y_m = j, Y_n - Y_m = k - j)}{\mathbb{P}(Y_n = k)} \quad (11.2.30)$$

But Y_m and $Y_n - Y_m$ are independent. Both variables have binomial distributions; the first with parameters m and p , and the second with parameters $n - m$ and p . Hence

$$\mathbb{P}(Y_m = j \mid Y_n = k) = \frac{\binom{m}{j} p^j (1-p)^{m-j} \binom{n-m}{k-j} p^{k-j} (1-p)^{n-m-(k-j)}}{\binom{n}{k} p^k (1-p)^{n-k}} = \frac{\binom{m}{j} \binom{m-n}{k-j}}{\binom{n}{k}} \quad (11.2.31)$$

Once again, note that the conditional distribution is independent of the success parameter p .

The Poisson Approximation

The *Poisson process* on $[0, \infty)$, named for Simeon Poisson, is a model for random points in continuous time. There are many deep and interesting connections between the Bernoulli trials process (which can be thought of as a model for random points in *discrete* time) and the Poisson process. These connections are explored in detail in the chapter on the Poisson process. In this section we just give the most famous and important result—the convergence of the binomial distribution to the Poisson distribution.

For reference, the Poisson distribution with rate parameter $r \in (0, \infty)$ has probability density function

$$g(n) = e^{-r} \frac{r^n}{n!}, \quad n \in \mathbb{N} \quad (11.2.32)$$

The parameter r is both the mean and the variance of the distribution. In addition, the probability generating function is $t \mapsto e^{t(r-1)}$.

Suppose that $p_n \in (0, 1)$ for $n \in \mathbb{N}_+$ and that $np_n \rightarrow r \in (0, \infty)$ as $n \rightarrow \infty$. Then the binomial distribution with parameters n and p_n converges to the Poisson distribution with parameter r as $n \rightarrow \infty$.

Proof from density functions

Let f_n denote the binomial PDF with parameters n and p_n . Then for $k \in \{0, 1, \dots, n\}$

$$f_n(k) = \binom{n}{k} p_n^k (1 - p_n)^{n-k} = \frac{1}{k!} [np_n] [(n-1)p_n] \cdots [(n-k+1)p_n] \left(1 - \frac{p_n}{n}\right)^{n-k} \quad (11.2.33)$$

But $(n-j)p_n \rightarrow r$ as $n \rightarrow \infty$ for fixed j . Also, using a basic theorem from calculus, $(1 - p_n/n)^{n-k} \rightarrow e^{-r}$ as $n \rightarrow \infty$. Hence $f_n(k) \rightarrow e^{-r} \frac{r^k}{k!}$ as $n \rightarrow \infty$.

Proof from generating functions

For $t \in \mathbb{R}$, using the same basic limit from calculus,

$$[(1 - p_n) + p_n t]^n = \left[1 + n \frac{p_n}{n} (t - 1)\right]^n \rightarrow e^{r(t-1)} \text{ as } n \rightarrow \infty \quad (11.2.34)$$

The left side is the PGF of the binomial distribution with parameters n and p_n , while the right side is the PGF of the Poisson distribution with parameter r .

Under the same conditions as the previous theorem, the mean and variance of the binomial distribution converge to the mean and variance of the limiting Poisson distribution, respectively.

1. $np_n \rightarrow r$ as $n \rightarrow \infty$
2. $np_n(1 - p_n) \rightarrow r$ as $n \rightarrow \infty$

Proof

By assumption, $np_n \rightarrow r$ as $n \rightarrow \infty$, and so it also follows that $p_n \rightarrow 0$ as $n \rightarrow \infty$.

Compare the Poisson experiment and the binomial timeline experiment.

1. Open the Poisson experiment and set $r = 1$ and $t = 5$. Run the experiment a few times and note the general behavior of the random points in time. Note also the shape and location of the probability density function and the mean \pm standard deviation bar.
2. Now open the binomial timeline experiment and set $n = 100$ and $p = 0.05$. Run the experiment a few times and note the general behavior of the random points in time. Note also the shape and location of the probability density function and the mean \pm standard deviation bar.

From a practical point of view, the convergence of the binomial distribution to the Poisson means that if the number of trials n is “large” and the probability of success p “small”, so that np^2 is small, then the binomial distribution with parameters n and p is well approximated by the Poisson distribution with parameter $r = np$. This is often a useful result, because the Poisson distribution has fewer parameters than the binomial distribution (and often in real problems, the parameters may only be known approximately). Specifically, in the approximating Poisson distribution, we do not need to know the number of trials n and the probability of success p *individually*, but only in the *product* np . The condition that np^2 be small means that the variance of the binomial distribution, namely $np(1 - p) = np - np^2$ is approximately r , the variance of the approximating Poisson distribution.

The Normal Approximation

Open the binomial timeline experiment. For selected values of $p \in (0, 1)$, start with $n = 1$ and successively increase n by 1. For each value of n , Note the shape of the probability density function of the number of successes and the proportion of successes. With $n = 100$, run the experiment 1000 times and compare the empirical density function to the probability density function for the number of successes and the proportion of successes

The characteristic bell shape that you should observe in the previous exercise is an example of the central limit theorem, because the binomial variable can be written as a sum of n independent, identically distributed random variables (the indicator variables).

The standard score Z_n of Y_n is the same as the standard score of M_n :

$$Z_n = \frac{Y_n - np}{\sqrt{np(1-p)}} = \frac{M_n - p}{\sqrt{p(1-p)/n}} \quad (11.2.35)$$

The distribution of Z_n converges to the standard normal distribution as $n \rightarrow \infty$.

This version of the central limit theorem is known as the *DeMoivre-Laplace theorem*, and is named after Abraham DeMoivre and Simeon Laplace. From a practical point of view, this result means that, for large n , the distribution of Y_n is approximately normal, with mean np and standard deviation $\sqrt{np(1-p)}$ and the distribution of M_n is approximately normal, with mean p and standard deviation $\sqrt{p(1-p)/n}$. Just how large n needs to be for the normal approximation to work well depends on the value of p . The rule of thumb is that we need $np \geq 5$ and $n(1-p) \geq 5$ (the first condition is the significant one when $p \leq \frac{1}{2}$ and the second condition is the significant one when $p \geq \frac{1}{2}$). Finally, when using the normal approximation, we should remember to use the continuity correction, since the binomial is a discrete distribution.

General Families

For a fixed number of trials n , the binomial distribution is a member of two general families of distributions. First, it is a general exponential distribution.

Suppose that Y has the binomial distribution with parameters n and p , where $n \in \mathbb{N}_+$ is fixed and $p \in (0, 1)$. The distribution of Y is a one-parameter exponential family with natural parameter $\ln\left(\frac{p}{1-p}\right)$ and natural statistic Y .

Proof

This follows from the definition of the general exponential family. The support set $\{0, 1, \dots, n\}$ does not depend on p , and for y in this set,

$$f_n(y) = \binom{n}{y} p^y (1-p)^{n-y} = \binom{n}{y} (1-p)^n \left(\frac{p}{1-p}\right)^y = \binom{n}{y} (1-p)^n \exp\left[y \ln\left(\frac{p}{1-p}\right)\right] \quad (11.2.36)$$

Note that the natural parameter is the logarithm of the *odds ratio* corresponding to p . This function is sometimes called the *logit* function. The binomial distribution is also a power series distribution

Suppose again that Y has the binomial distribution with parameters n and p , where $n \in \mathbb{N}_+$ is fixed and $p \in (0, 1)$. The distribution of Y is a power series distribution in the parameter $\theta = \frac{p}{1-p}$, corresponding to the function $\theta \mapsto (1+\theta)^n$.

Proof

This follows from the definition of the power series distribution. As before, for $y \in \{0, 1, \dots, n\}$

$$f_n(y) = \binom{n}{y} p^y (1-p)^{n-y} = \binom{n}{y} (1-p)^n \left(\frac{p}{1-p}\right)^y = \frac{1}{(1+\theta)^n} \binom{n}{y} \theta^y \quad (11.2.37)$$

where $\theta = \frac{p}{1-p}$. This is the power series distribution in θ , with coefficients $\binom{n}{y}$, corresponding to the function $\theta \mapsto (1+\theta)^n$.

The Proportion of Successes

Suppose again that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter p , and that as usual, $Y_n = \sum_{i=1}^n X_i$ is the number of successes in the first n trials for $n \in \mathbb{N}$. The *proportion of successes* in the first n trials is the random variable

$$M_n = \frac{Y_n}{n} = \frac{1}{n} \sum_{i=1}^n X_i \quad (11.2.38)$$

In statistical terms, M_n is the sample mean of the random sample (X_1, X_2, \dots, X_n) . The proportion of successes M_n is typically used to estimate the probability of success p when this probability is unknown.

It is easy to express the probability density function of the proportion of successes M_n in terms of the probability density function of the number of successes Y_n . First, note that M_n takes the values k/n where $k \in \{0, 1, \dots, n\}$.

The probability density function of M_n is given by

$$\mathbb{P}\left(M_n = \frac{k}{n}\right) = \binom{n}{k} p^k (1-p)^{n-k}, \quad k \in \{0, 1, \dots, n\} \quad (11.2.39)$$

Proof

Trivially, $M_n = k/n$ if and only if $Y_n = k$ for $k \in \{0, 1, \dots, n\}$.

In the binomial coin experiment, select the proportion of heads. Vary n and p with the scroll bars and note the shape of the probability density function. For selected values of the parameters, run the experiment 1000 times and compare the relative frequency function to the probability density function.

The mean and variance of the proportion of successes M_n are easy to compute from the mean and variance of the number of successes Y_n .

The mean and variance of M_n are

1. $\mathbb{E}(M_n) = p$
2. $\text{var}(M_n) = \frac{1}{n}p(1-p)$.

Proof

From the scaling properties of expected value and variance,

1. $\mathbb{E}(M_n) = \frac{1}{n}\mathbb{E}(Y_n) = \frac{1}{n}np = p$
2. $\text{var}(M_n) = \frac{1}{n^2}\text{var}(Y_n) = \frac{1}{n^2}np(1-p) = \frac{1}{n}p(1-p)$

In the binomial coin experiment, select the proportion of heads. Vary n and p and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the empirical moments to the distribution moments.

Recall that skewness and kurtosis are standardized measures. Since M_n and Y_n have the same standard score, the skewness and kurtosis of M_n are the same as the [skewness](#) and [kurtosis](#) of Y_n given above.

In statistical terms, part (a) of the moment result above means that M_n is an unbiased estimator of p . From part (b) note that $\text{var}(M_n) \leq \frac{1}{4n}$ for any $p \in [0, 1]$. In particular, $\text{var}(M_n) \rightarrow 0$ as $n \rightarrow \infty$ and the convergence is uniform in $p \in [0, 1]$. Thus, the estimate improves as n increases; in statistical terms, this is known as *consistency*.

For every $\epsilon > 0$, $\mathbb{P}(|M_n - p| \geq \epsilon) \rightarrow 0$ as $n \rightarrow \infty$ and the convergence is uniform in $p \in [0, 1]$.

Proof

This follows from the last result and Chebyshev's inequality.

The last result is a special case of the weak law of large numbers and means that $M_n \rightarrow p$ as $n \rightarrow \infty$ in probability. The strong law of large numbers states that the convergence actually holds with probability 1.

The *proportion* of successes M_n has a number of nice properties as an estimator of the *probability* of success p . As already noted, it is unbiased and consistent. In addition, since Y_n is a [sufficient statistic](#) for p , based on the sample (X_1, X_2, \dots, X_n) , it follows that M_n is sufficient for p as well. Since $\mathbb{E}(X_i) = p$, M_n is trivially the method of moments estimator of p . Assuming that the parameter space for p is $[0, 1]$, it is also the maximum likelihood estimator of p .

The likelihood function for $p \in [0, 1]$, based on the observed sample $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, is $L_y(p) = p^y(1-p)^{n-y}$, where $y = \sum_{i=1}^n x_i$. The likelihood is maximized at $m = y/n$.

Proof

By definition, the likelihood function is simply the joint PDF of (X_1, X_2, \dots, X_n) thought of as a function of the parameter p , for fixed (x_1, x_2, \dots, x_n) . Thus the form of the likelihood function follows from the joint PDF given in the Introduction. If $y = 0$, L_y is decreasing and hence is maximized at $p = 0 = y/n$. If $y = n$, L_y is increasing and is maximized at $p = 1 = y/n$. If $0 < y < n$, the log-likelihood function is

$$\ln[L_y(p)] = y \ln(p) + (n-y) \ln(1-p) \quad (11.2.40)$$

and the derivative is

$$\frac{d}{dp} \ln[L_y(p)] = \frac{y}{p} - \frac{n-y}{1-p} \quad (11.2.41)$$

There is a single critical point at y/n . The second derivative of the log-likelihood function is negative, so the maximum on $(0, 1)$ occurs at the critical point.

See Estimation in the Bernoulli Model in the chapter on Set Estimation for a different approach to the problem of estimating p .

Examples and Applications

Simple Exercises

A student takes a multiple choice test with 20 questions, each with 5 choices (only one of which is correct). Suppose that the student blindly guesses. Let X denote the number of questions that the student answers correctly. Find each of the following:

1. The probability density function of X .
2. The mean of X .
3. The variance of X .
4. The probability that the student answers at least 12 questions correctly (the score that she needs to pass).

Answer

1. $\mathbb{P}(X = x) = \binom{20}{x} \left(\frac{1}{5}\right)^x \left(\frac{4}{5}\right)^{20-x}$ for $x \in \{0, 1, \dots, 20\}$
2. $\mathbb{E}(X) = 4$
3. $\text{var}(X) = \frac{16}{5}$

4. $\mathbb{P}(X \geq 12) \approx 0.000102$ She has no hope of passing.

A certain type of missile has failure probability 0.02. Let Y denote the number of failures in 50 tests. Find each of the following:

1. The probability density function of Y .
2. The mean of Y .
3. The variance of Y .
4. The probability of at least 47 successful tests.

Answer

1. $\mathbb{P}(Y = y) = \binom{50}{y} \left(\frac{1}{50}\right)^y \left(\frac{49}{50}\right)^{50-y}$ for $y \in \{0, 1, \dots, 50\}$
2. $\mathbb{E}(Y) = 1$
3. $\text{var}(Y) = \frac{49}{50}$
4. $\mathbb{P}(Y \leq 3) \approx 0.9822$

Suppose that in a certain district, 40% of the registered voters prefer candidate A . A random sample of 50 registered voters is selected. Let Z denote the number in the sample who prefer A . Find each of the following:

1. The probability density function of Z .
2. The mean of Z .
3. The variance of Z .
4. The probability that Z is less than 19.
5. The normal approximation to the probability in (d).

Answer

1. $\mathbb{P}(Z = z) = \binom{50}{z} \left(\frac{2}{5}\right)^z \left(\frac{3}{5}\right)^{50-z}$ for $z \in \{0, 1, \dots, 50\}$
2. $\mathbb{E}(Z) = 20$
3. $\text{var}(Z) = 12$
4. $\mathbb{P}(Z < 19) = 0.3356$
5. $\mathbb{P}(Z < 19) \approx 0.3330$

Coins and Dice

Recall that a *standard* die is a six-sided die. A *fair* die is one in which the faces are equally likely. An *ace-six flat* die is a standard die in which faces 1 and 6 have probability $\frac{1}{4}$ each, and faces 2, 3, 4, and 5 have probability $\frac{1}{8}$.

A standard, fair die is tossed 10 times. Let N denote the number of aces. Find each of the following:

1. The probability density function of N .
2. The mean of N .
3. The variance of N .

Answer

1. $\mathbb{P}(N = k) = \binom{10}{k} \left(\frac{1}{6}\right)^k \left(\frac{5}{6}\right)^{10-k}$ for $k \in \{0, 1, \dots, 10\}$
2. $\mathbb{E}(N) = \frac{5}{3}$
3. $\text{var}(N) = \frac{25}{18}$

A coin is tossed 100 times and results in 30 heads. Find the probability density function of the number of heads in the first 20 tosses.

Answer

Let Y_n denote the number of heads in the first n tosses.

$$\mathbb{P}(Y_{20} = y \mid Y_{100} = 30) = \frac{\binom{20}{y} \binom{80}{30-y}}{\binom{100}{30}}, \quad y \in \{0, 1, \dots, 20\} \quad (11.2.42)$$

An ace-six flat die is rolled 1000 times. Let Z denote the number of times that a score of 1 or 2 occurred. Find each of the following:

1. The probability density function of Z .
2. The mean of Z .
3. The variance of Z .
4. The probability that Z is at least 400.
5. The normal approximation of the probability in (d)

Answer

1. $\mathbb{P}(Z = z) = \binom{1000}{z} \left(\frac{3}{8}\right)^z \left(\frac{5}{8}\right)^{1000-z}$ for $z \in \{0, 1, \dots, 1000\}$
2. $\mathbb{E}(Z) = 375$
3. $\text{var}(Z) = 1875/8$
4. $\mathbb{P}(Z \geq 400) \approx 0.0552$
5. $\mathbb{P}(Z \geq 400) \approx 0.0550$

In the binomial coin experiment, select the proportion of heads. Set $n = 10$ and $p = 0.4$. Run the experiment 100 times. Over all 100 runs, compute the square root of the average of the squares of the errors, when M used to estimate p . This number is a measure of the quality of the estimate.

In the binomial coin experiment, select the number of heads Y , and set $p = 0.5$ and $n = 15$. Run the experiment 1000 times with and compute the following:

1. $\mathbb{P}(5 \leq Y \leq 10)$
2. The relative frequency of the event $\{5 \leq Y \leq 10\}$
3. The normal approximation to $\mathbb{P}(5 \leq Y \leq 10)$

Answer

1. $\mathbb{P}(5 \leq Y \leq 10) = 0.8815$
3. $\mathbb{P}(5 \leq Y \leq 10) \approx 0.878$

In the binomial coin experiment, select the proportion of heads M and set $n = 30$, $p = 0.6$. Run the experiment 1000 times and compute each of the following:

1. $\mathbb{P}(0.5 \leq M \leq 0.7)$
2. The relative frequency of the event $\{0.5 \leq M \leq 0.7\}$
3. The normal approximation to $\mathbb{P}(0.5 \leq M \leq 0.7)$

Answer

1. $\mathbb{P}(0.5 \leq M \leq 0.7) = 0.8089$
3. $\mathbb{P}(0.5 \leq M \leq 0.7) \approx 0.808$

Famous Problems

In 1693, Samuel Pepys asked Isaac Newton whether it is more likely to get at least one 6 in 6 rolls of a die, or at least two 6's in 12 rolls, or at least three 6's in 18 rolls. This problem is known as a *Pepys' problem*; naturally, Pepys had fair dice in mind.

Solve Pepys' problem using the binomial distribution.

Answer

Let Y_n denote the number of 6's in n rolls of a fair die, so that Y_n has the binomial distribution with parameters n and $p = \frac{1}{6}$. Using the binomial PDF (and complements to simplify the computations),

1. $\mathbb{P}(Y_6 \geq 1) = 1 - \mathbb{P}(Y_6 = 0) = 0.6651$
2. $\mathbb{P}(Y_{12} \geq 2) = 1 - \mathbb{P}(Y_{12} \leq 1) = 0.6187$
3. $\mathbb{P}(Y_{18} \geq 3) = 1 - \mathbb{P}(Y_{18} \leq 2) = 0.5973$

So the first event, at least one 6 in 6 rolls, is the most likely, and in fact the three events, in the order given, decrease in probability.

With fair dice run the simulation of the dice experiment 500 times and compute the relative frequency of the event of interest. Compare the results with the theoretical probabilities above.

1. At least one 6 with $n = 6$.
2. At least two 6's with $n = 12$.
3. At least three 6's with $n = 18$.

It appears that Pepys had some sort of misguided linearity in mind, given the comparisons that he wanted. Let's solve the general problem.

For $k \in \mathbb{N}_+$, let Y_k denote the number of 6's in $6k$ rolls of a fair die.

1. Identify the distribution of Y_k .
2. Find the mean and variance of Y_k .
3. Give an exact formula for $\mathbb{P}(Y_k \geq k)$, the probability of at least k 6's in $6k$ rolls of a fair die.
4. Give the normal approximation of the probability in (c).
5. Find the limit of the probability in (d) as $k \rightarrow \infty$.

Proof

1. Y_k has the binomial distribution with parameters $n = 6k$ and $p = \frac{1}{6}$.

2. $\mathbb{E}(Y_k) = (6k)\frac{1}{6} = k$, $\text{var}(Y_k) = (6k)\frac{1}{6}\frac{5}{6} = \frac{5}{6}k$
3. $\mathbb{P}(Y_k \geq k) = 1 - \mathbb{P}(Y \leq k-1) = 1 - \sum_{j=0}^{k-1} \binom{6k}{j} \left(\frac{1}{6}\right)^j \left(\frac{5}{6}\right)^{6k-j}$
4. Using the continuity correction,

$$\mathbb{P}(Y_k \geq k) = \mathbb{P}\left(Y_k \geq k - \frac{1}{2}\right) = \mathbb{P}\left(\frac{Y_k - k}{\sqrt{(5/6)k}} \geq \frac{-1/2}{\sqrt{(5/6)k}}\right) = 1 - \Phi\left(\frac{-1/2}{\sqrt{(5/6)k}}\right) \quad (11.2.43)$$

5. $1 - \Phi\left(\frac{-1/2}{\sqrt{(5/6)k}}\right) \rightarrow 1 - \Phi(0) = \frac{1}{2}$ as $k \rightarrow \infty$

So on average, the number of 6's in $6k$ rolls of a fair die is k , and that fact might have influenced Pepys thinking. The next problem is known as *DeMere's problem*, named after Chevalier De Mere

Which is more likely: at least one 6 with 4 throws of a fair die or at least one double 6 in 24 throws of two fair dice? .

Answer

Let Y_n denote the number of 6's in n rolls of a fair die, and let Z_n denote the number of double 6's in n rolls of a pair of fair dice. Then Y_n as the binomial distribution with parameters n and $p = \frac{1}{6}$, and Z_n has the binomial distribution with parameters n and $p = \frac{1}{36}$. Using the binomial PDF,

1. $\mathbb{P}(Y_4 \geq 1) = 0.5177$
2. $\mathbb{P}(Z_{24} \geq 1) = 0.4914$

Data Analysis Exercises

In the cicada data, compute the proportion of males in the entire sample, and the proportion of males of each species in the sample.

Answer

1. $m = 0.433$
2. $m_0 = 0.636$
3. $m_1 = 0.259$
4. $m_2 = 0.5$

In the M&M data, pool the bags to create a large sample of M&Ms. Now compute the sample proportion of red M&Ms.

Answer

$$m_{\text{red}} = 0.168$$

The Galton Board

The *Galton board* is a triangular array of pegs. The rows are numbered by the natural numbers $\mathbb{N} = \{0, 1, \dots\}$ from top downward. Row n has $n+1$ pegs numbered from left to right by the integers $\{0, 1, \dots, n\}$. Thus a peg can be uniquely identified by the ordered pair (n, k) where n is the row number and k is the peg number in that row. The Galton board is named after Francis Galton.

Now suppose that a ball is dropped from above the top peg $(0, 0)$. Each time the ball hits a peg, it bounces to the right with probability p and to the left with probability $1-p$, independently from bounce to bounce.

The number of the peg that the ball hits in row n has the binomial distribution with parameters n and p .

In the Galton board experiment, select random variable Y (the number of moves right). Vary the parameters n and p and note the shape and location of the probability density function and the mean \pm standard deviation bar. For selected values of the parameters, click single step several times and watch the ball fall through the pegs. Then run the experiment 1000 times and watch the path of the ball. Compare the relative frequency function and empirical moments to the probability density function and distribution moments, respectively.

Structural Reliability

Recall the discussion of *structural reliability* given in the last section on Bernoulli trials. In particular, we have a system of n similar components that function independently, each with reliability p . Suppose now that the system as a whole functions properly if and only if at least k of the n components are good. Such a systems is called, appropriately enough, a *k out of n system*. Note that the *series* and *parallel* systems considered in the previous section are n out of n and 1 out of n systems, respectively.

Consider the k out of n system.

1. The state of the system is $\mathbf{1}(Y_n \geq k)$ where Y_n is the number of working components.
2. The reliability function is $r_{n,k}(p) = \sum_{i=k}^n \binom{n}{i} p^i (1-p)^{n-i}$.

In the binomial coin experiment, set $n = 10$ and $p = 0.9$ and run the simulation 1000 times. Compute the empirical reliability and compare with the true reliability in each of the following cases:

1. 10 out of 10 (series) system.
2. 1 out of 10 (parallel) system.
3. 4 out of 10 system.

Consider a system with $n = 4$ components. Sketch the graphs of $r_{4,1}$, $r_{4,2}$, $r_{4,3}$, and $r_{4,4}$ on the same set of axes.

An n out of $2n - 1$ system is a *majority rules* system.

1. Compute the reliability of a 2 out of 3 system.
2. Compute the reliability of a 3 out of 5 system
3. For what values of p is a 3 out of 5 system more reliable than a 2 out of 3 system?
4. Sketch the graphs of $r_{3,2}$ and $r_{5,3}$ on the same set of axes.
5. Show that $r_{2n-1,n}(\frac{1}{2}) = \frac{1}{2}$.

Answer

1. $r_{3,2}(p) = 3p^2 - 2p^3$
2. $r_{5,3}(p) = 10p^3 - 15p^4 + 6p^5$
3. 3 out of 5 is better for $p \geq \frac{1}{2}$

In the binomial coin experiment, compute the empirical reliability, based on 100 runs, in each of the following cases. Compare your results to the true probabilities.

1. A 2 out of 3 system with $p = 0.3$
2. A 3 out of 5 system with $p = 0.3$
3. A 2 out of 3 system with $p = 0.8$
4. A 3 out of 5 system with $p = 0.8$

Reliable Communications

Consider the transmission of bits (0s and 1s) through a noisy channel. Specifically, suppose that when bit $i \in \{0, 1\}$ is transmitted, bit i is received with probability $p_i \in (0, 1)$ and the complementary bit $1 - i$ is received with probability $1 - p_i$. Given the bits transmitted, bits are received correctly or incorrectly independently of one-another. Suppose now, that to increase reliability, a given bit I is repeated $n \in \mathbb{N}_+$ times in the transmission. *A priori*, we believe that $\mathbb{P}(I = 1) = \alpha \in (0, 1)$ and $\mathbb{P}(I = 0) = 1 - \alpha$. Let X denote the number of 1s received when bit I is transmitted n times.

Find each of the following:

1. The conditional distribution of X given $I = i \in \{0, 1\}$
2. the probability density function of X
3. $\mathbb{E}(X)$
4. $\text{var}(X)$

Answer

1. Give $I = 1$, X has the binomial distribution with parameters n and p_1 . Given $I = 0$, X has the binomial distribution with parameters n and $1 - p_0$.
2. $\mathbb{P}(X = k) = \binom{n}{k} [\alpha p_1^k (1 - p_1)^{n-k} + (1 - \alpha)(1 - p_0)^k p_0^{n-k}]$ for $k \in \{0, 1, \dots, n\}$.
3. $\mathbb{E}(X) = n [\alpha p_1 + (1 - \alpha)(1 - p_0)]$
4. $\text{var}(X) = \alpha [np_1(1 - p_1) + n^2 p_1^2] + (1 - \alpha) [np_0(1 - p_0) + n^2 (1 - p_0)^2] - n^2 [\alpha p_1 + (1 - \alpha)(1 - p_0)]^2$

Simplify the results in the last exercise in the symmetric case where $p_1 = p_0 =: p$ (so that the bits are equally reliable) and with $\alpha = \frac{1}{2}$ (so that we have no prior information).

Answer

1. Give $I = 1$, X has the binomial distribution with parameters n and p . Given $I = 0$, X has the binomial distribution with parameters n and $1 - p$.
2. $\mathbb{P}(X = k) = \frac{1}{2} \binom{n}{k} [p^k (1 - p)^{n-k} + (1 - p)^k p^{n-k}]$ for $k \in \{0, 1, \dots, n\}$.
3. $\mathbb{E}(X) = \frac{1}{2} n$
4. $\text{var}(X) = np(1 - p) + \frac{1}{2} n^2 [p^2 + (1 - p)^2] - \frac{1}{4} n^2$

Our interest, of course, is predicting the bit transmitted given the bits received.

Find the posterior probability that $I = 1$ given $X = k \in \{0, 1, \dots, n\}$.

Answer

$$\mathbb{P}(I = 1 \mid X = k) = \frac{\alpha p_1^k (1 - p_1)^{n-k}}{\alpha p_1^k (1 - p_1)^{n-k} + (1 - \alpha)(1 - p_0)^k p_0^{n-k}} \quad (11.2.44)$$

Presumably, our decision rule would be to conclude that 1 was transmitted if the posterior probability in the previous exercise is greater than $\frac{1}{2}$ and to conclude that 0 was transmitted if this probability is less than $\frac{1}{2}$. If the probability equals $\frac{1}{2}$, we have no basis to prefer one bit over the other.

Give the decision rule in the symmetric case where $p_1 = p_0 =: p$, so that the bits are equally reliable. Assume that $p > \frac{1}{2}$, so that we at least have a better than even chance of receiving the bit transmitted.

Answer

Give $X = k$, we conclude that bit 1 was transmitted if

$$k > \frac{n}{2} - \frac{1}{2} \frac{\ln(\alpha) - \ln(1 - \alpha)}{\ln(p) - \ln(1 - p)} \quad (11.2.45)$$

and we conclude that bit 0 was transmitted if the reverse inequality holds.

Not surprisingly, in the symmetric case with no prior information, so that $\alpha = \frac{1}{2}$, we conclude that bit i was transmitted if a majority of bits received are i .

Bernstein Polynomials

The *Weierstrass Approximation Theorem*, named after Karl Weierstrass, states that any real-valued function that is continuous on a closed, bounded interval can be uniformly approximated on that interval, to any degree of accuracy, with a polynomial. The theorem is important, since polynomials are simple and basic functions, and a bit surprising, since continuous functions can be quite strange.

In 1911, Sergei Bernstein gave an explicit construction of polynomials that uniformly approximate a given continuous function, using Bernoulli trials. Bernstein's result is a beautiful example of the *probabilistic method*, the use of probability theory to obtain results in other areas of mathematics that are seemingly unrelated to probability.

Suppose that f is a real-valued function that is continuous on the interval $[0, 1]$. The *Bernstein polynomial* of degree n for f is defined by

$$b_n(p) = \mathbb{E}_p[f(M_n)], \quad p \in [0, 1] \quad (11.2.46)$$

where M_n is the proportion of successes in the first n Bernoulli trials with success parameter p , as defined earlier. Note that we are emphasizing the dependence on p in the expected value operator. The next exercise gives a more explicit representation, and shows that the Bernstein polynomial is, in fact, a polynomial

The Bernstein polynomial of degree n can be written as follows:

$$b_n(p) = \sum_{k=0}^n f\left(\frac{k}{n}\right) \binom{n}{k} p^k (1-p)^{n-k}, \quad p \in [0, 1] \quad (11.2.47)$$

Proof

This follows from the change of variables theorem for expected value.

The Bernstein polynomials satisfy the following properties:

1. $b_n(0) = f(0)$ and $b_n(1) = f(1)$
2. $b_1(p) = f(0) + [f(1) - f(0)]p$ for $p \in [0, 1]$.
3. $b_2(p) = f(0) + 2[f(\frac{1}{2}) - f(0)]p + [f(1) - 2f(\frac{1}{2}) + f(0)]p^2$ for $p \in [0, 1]$

From part (a), the graph of b_n passes through the endpoints $(0, f(0))$ and $(1, f(1))$. From part (b), the graph of b_1 is a line connecting the endpoints. From (c), the graph of b_2 is parabola passing through the endpoints and the point $(\frac{1}{2}, \frac{1}{4}f(0) + \frac{1}{2}f(\frac{1}{2}) + \frac{1}{4}f(1))$.

The next result gives *Bernstein's theorem* explicitly.

$b_n \rightarrow f$ as $n \rightarrow \infty$ uniformly on $[0, 1]$.

Proof

Since f is continuous on the closed, bounded interval $[0, 1]$ it is bounded on this interval. Thus, there exists a constant C such that $|f(p)| \leq C$ for all $p \in [0, 1]$. Also, f is uniformly continuous on $[0, 1]$. Thus, for any $\epsilon > 0$ there exists $\delta > 0$ such that if $p, q \in [0, 1]$ and $|p - q| < \delta$ then $|f(p) - f(q)| < \epsilon$. From basic properties of expected value,

$$|b_n(p) - f(p)| \leq \mathbb{E}_p[|f(M_n) - f(p)|, |M_n - p| < \delta] + \mathbb{E}_p[|f(M_n) - f(p)|, |M_n - p| \geq \delta] \quad (11.2.48)$$

Hence $|b_n(p) - f(p)| \leq \epsilon + 2C\mathbb{P}_p(|M_n - p| \geq \delta)$ for any $p \in [0, 1]$. But by [weak law of large numbers](#) above, $\mathbb{P}_p(|M_n - p| \geq \delta) \rightarrow 0$ as $n \rightarrow \infty$ uniformly in $p \in [0, 1]$.

Compute the Bernstein polynomials of orders 1, 2, and 3 for the function f defined by $f(x) = \cos(\pi x)$ for $x \in [0, 1]$. Graph f and the three polynomials on the same set of axes.

Answer

1. $b_1(p) = 1 - 2p$
2. $b_2(p) = 1 - 2p$
3. $b_3(p) = 1 - \frac{3}{2}p - \frac{3}{2}p^2 + p^3$

Use a computer algebra system to compute the Bernstein polynomials of orders 10, 20, and 30 for the function f defined below. Use the CAS to graph the function and the three polynomials on the same axes.

$$f(x) = \begin{cases} 0, & x = 0 \\ x \sin\left(\frac{\pi}{x}\right), & x \in (0, 1] \end{cases} \quad (11.2.49)$$

This page titled [11.2: The Binomial Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.3: The Geometric Distribution

Basic Theory

Definitions

Suppose again that our random experiment is to perform a sequence of Bernoulli trials $\mathbf{X} = (X_1, X_2, \dots)$ with success parameter $p \in (0, 1]$. In this section we will study the random variable N that gives the *trial number of the first success* and the random variable M that gives the number of failures before the first success.

Let $N = \min\{n \in \mathbb{N}_+ : X_n = 1\}$, the trial number of the first success, and let $M = N - 1$, the number of failures before the first success. The distribution of N is the *geometric distribution on \mathbb{N}_+* and the distribution of M is the *geometric distribution on \mathbb{N}* . In both cases, p is the *success parameter* of the distribution.

Since N and M differ by a constant, the properties of their distributions are very similar. Nonetheless, there are applications where it is more natural to use one rather than the other, and in the literature, the term *geometric distribution* can refer to either. In this section, we will concentrate on the distribution of N , pausing occasionally to summarize the corresponding results for M .

The Probability Density Function

N has probability density function f given by $f(n) = p(1-p)^{n-1}$ for $n \in \mathbb{N}_+$.

Proof

Note first that $\{N = n\} = \{X_1 = 0, \dots, X_{n-1} = 0, X_n = 1\}$. By independence, the probability of this event is $(1-p)^{n-1}p$.

Check that f is a valid PDF

By standard results for geometric series

$$\sum_{n=1}^{\infty} \mathbb{P}(N = n) = \sum_{n=1}^{\infty} (1-p)^{n-1}p = \frac{p}{1-(1-p)} = 1 \quad (11.3.1)$$

A priori, we might have thought it possible to have $N = \infty$ with positive probability; that is, we might have thought that we could run Bernoulli trials forever without ever seeing a success. However, we now know this cannot happen when the success parameter p is positive.

The probability density function of M is given by $\mathbb{P}(M = n) = p(1-p)^n$ for $n \in \mathbb{N}$.

In the negative binomial experiment, set $k = 1$ to get the geometric distribution on \mathbb{N}_+ . Vary p with the scroll bar and note the shape and location of the probability density function. For selected values of p , run the simulation 1000 times and compare the relative frequency function to the probability density function.

Note that the probability density functions of N and M are decreasing, and hence have modes at 1 and 0, respectively. The geometric form of the probability density functions also explains the term *geometric distribution*.

Distribution Functions and the Memoryless Property

Suppose that T is a random variable taking values in \mathbb{N}_+ . Recall that the ordinary distribution function of T is the function $n \mapsto \mathbb{P}(T \leq n)$. In this section, the complementary function $n \mapsto \mathbb{P}(T > n)$ will play a fundamental role. We will refer to this function as the *right distribution function* of T . Of course both functions completely determine the distribution of T . Suppose again that N has the geometric distribution on \mathbb{N}_+ with success parameter $p \in (0, 1]$.

N has right distribution function G given by $G(n) = (1-p)^n$ for $n \in \mathbb{N}$.

Proof from Bernoulli trials

Note that $\{N > n\} = \{X_1 = 0, \dots, X_n = 0\}$. By independence, the probability of this event is $(1-p)^n$.

Direct proof

Using geometric series,

$$\mathbb{P}(N > n) = \sum_{k=n+1}^{\infty} \mathbb{P}(N = k) = \sum_{k=n+1}^{\infty} (1-p)^{k-1} p = \frac{p(1-p)^n}{1-(1-p)} = (1-p)^n \quad (11.3.2)$$

From the last result, it follows that the ordinary (left) distribution function of N is given by

$$F(n) = 1 - (1-p)^n, \quad n \in \mathbb{N} \quad (11.3.3)$$

We will now explore another characterization known as the *memoryless property*.

For $m \in \mathbb{N}$, the conditional distribution of $N - m$ given $N > m$ is the same as the distribution of N . That is,

$$\mathbb{P}(N > n + m \mid N > m) = \mathbb{P}(N > n); \quad m, n \in \mathbb{N} \quad (11.3.4)$$

Proof

From the [result above](#) and the definition of conditional probability,

$$\mathbb{P}(N > n + m \mid N > m) = \frac{\mathbb{P}(N > n + m)}{\mathbb{P}(N > m)} = \frac{(1-p)^{n+m}}{(1-p)^m} = (1-p)^n = \mathbb{P}(N > n) \quad (11.3.5)$$

Thus, if the first success has not occurred by trial number m , then the remaining number of trials needed to achieve the first success has the same distribution as the trial number of the first success in a fresh sequence of Bernoulli trials. In short, Bernoulli trials have no memory. This fact has implications for a gambler betting on Bernoulli trials (such as in the casino games roulette or craps). No betting strategy based on observations of past outcomes of the trials can possibly help the gambler.

Conversely, if T is a random variable taking values in \mathbb{N}_+ that satisfies the [memoryless property](#), then T has a geometric distribution.

Proof

Let $G(n) = \mathbb{P}(T > n)$ for $n \in \mathbb{N}$. The memoryless property and the definition of conditional probability imply that $G(m+n) = G(m)G(n)$ for $m, n \in \mathbb{N}$. Note that this is the *law of exponents* for G . It follows that $G(n) = G^n(1)$ for $n \in \mathbb{N}$. Hence T has the geometric distribution with parameter $p = 1 - G(1)$.

Moments

Suppose again that N is the trial number of the first success in a sequence of Bernoulli trials, so that N has the geometric distribution on \mathbb{N}_+ with parameter $p \in (0, 1]$. The mean and variance of N can be computed in several different ways.

$$\mathbb{E}(N) = \frac{1}{p}$$

Proof from the density function

Using the derivative of the geometric series,

$$\mathbb{E}(N) = \sum_{n=1}^{\infty} np(1-p)^{n-1} = p \sum_{n=1}^{\infty} n(1-p)^{n-1} \quad (11.3.6)$$

$$= p \sum_{n=1}^{\infty} -\frac{d}{dp}(1-p)^n = -p \frac{d}{dp} \sum_{n=0}^{\infty} (1-p)^n \quad (11.3.7)$$

$$= -p \frac{d}{dp} \frac{1}{p} = -p \left(-\frac{1}{p^2} \right) = \frac{1}{p} \quad (11.3.8)$$

Proof from the right distribution function

Recall that since N takes positive integer values, its expected value can be computed as the sum of the right distribution function. Hence

$$\mathbb{E}(N) = \sum_{n=0}^{\infty} \mathbb{P}(N > n) = \sum_{n=0}^{\infty} (1-p)^n = \frac{1}{p} \quad (11.3.9)$$

Proof from Bernoulli trials

We condition on the first trial X_1 : If $X_1 = 1$ then $N = 1$ and hence $\mathbb{E}(N | X_1 = 1) = 1$. If $X_1 = 0$ (equivalently $N > 1$) then by the memoryless property, $N - 1$ has the same distribution as N . Hence $\mathbb{E}(N | X_1 = 0) = 1 + \mathbb{E}(N)$. In short

$$\mathbb{E}(N | X_1) = 1 + (1 - X_1)\mathbb{E}(N) \quad (11.3.10)$$

It follows that

$$\mathbb{E}(N) = \mathbb{E}[\mathbb{E}(N | X_1)] = 1 + (1 - p)\mathbb{E}(N) \quad (11.3.11)$$

Solving gives $\mathbb{E}(N) = \frac{1}{p}$.

This result makes intuitive sense. In a sequence of Bernoulli trials with success parameter p we would “expect” to wait $1/p$ trials for the first success.

$$\text{var}(N) = \frac{1-p}{p^2}$$

Direct proof

We first compute $\mathbb{E}[N(N-1)]$. This is an example of a factorial moment, and we will compute the general factorial moments below. Using derivatives of the geometric series again,

$$\mathbb{E}[N(N-1)] = \sum_{n=2}^{\infty} n(n-1)p(1-p)^{n-1} = p(1-p) \sum_{n=2}^{\infty} n(n-1)(1-p)^{n-2} \quad (11.3.12)$$

$$= p(1-p) \frac{d^2}{dp^2} \sum_{n=0}^{\infty} (1-p)^n = p(1-p) \frac{d^2}{dp^2} \frac{1}{p} = p(1-p) \frac{2}{p^3} = 2 \frac{1-p}{p^2} \quad (11.3.13)$$

Since $\mathbb{E}(N) = \frac{1}{p}$, it follows that $\mathbb{E}(N^2) = \frac{2-p}{p^2}$ and hence $\text{var}(N) = \frac{1-p}{p^2}$

Proof from Bernoulli trials

Recall that

$$\mathbb{E}(N | X_1) = 1 + (1 - X_1)\mathbb{E}(N) = 1 + \frac{1}{p}(1 - X_1) \quad (11.3.14)$$

and by the same reasoning, $\text{var}(N | X_1) = (1 - X_1)\text{var}(N)$. Hence

$$\text{var}(N) = \text{var}[\mathbb{E}(N | X_1)] + \mathbb{E}[\text{var}(N | X_1)] = \frac{1}{p^2}p(1-p) + (1-p)\text{var}(N) \quad (11.3.15)$$

Solving gives $\text{var}(N) = \frac{1-p}{p^2}$.

Note that $\text{var}(N) = 0$ if $p = 1$, hardly surprising since N is deterministic (taking just the value 1) in this case. At the other extreme, $\text{var}(N) \uparrow \infty$ as $p \downarrow 0$.

In the negative binomial experiment, set $k = 1$ to get the geometric distribution. Vary p with the scroll bar and note the location and size of the mean \pm standard deviation bar. For selected values of p , run the simulation 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

the probability generating function P of N is given by

$$P(t) = \mathbb{E}(t^N) = \frac{pt}{1 - (1-p)t}, \quad |t| < \frac{1}{1-p} \quad (11.3.16)$$

Proof

This result follows from yet another application of geometric series:

$$\mathbb{E}(t^N) = \sum_{n=1}^{\infty} t^n p(1-p)^{n-1} = pt \sum_{n=1}^{\infty} [t(1-p)]^{n-1} = \frac{pt}{1-(1-p)t}, \quad |(1-p)t| < 1 \quad (11.3.17)$$

Recall again that for $x \in \mathbb{R}$ and $k \in \mathbb{N}$, the *falling power* of x of order k is $x^{(k)} = x(x-1)\cdots(x-k+1)$. If X is a random variable, then $\mathbb{E}[X^{(k)}]$ is the *factorial moment* of X of order k .

The factorial moments of N are given by

$$\mathbb{E}[N^{(k)}] = k! \frac{(1-p)^{k-1}}{p^k}, \quad k \in \mathbb{N}_+ \quad (11.3.18)$$

Proof from geometric series

Using derivatives of geometric series again,

$$\mathbb{E}[N^{(k)}] = \sum_{n=k}^{\infty} n^{(k)} p(1-p)^{n-1} = p(1-p)^{k-1} \sum_{n=k}^{\infty} n^{(k)} (1-p)^{n-k} \quad (11.3.19)$$

$$= p(1-p)^{k-1} (-1)^k \frac{d^k}{dp^k} \sum_{n=0}^{\infty} (1-p)^n = p(1-p)^{k-1} (-1)^k \frac{d^k}{dp^k} \frac{1}{p} = k! \frac{(1-p)^{k-1}}{p^k} \quad (11.3.20)$$

Proof from the generating function

Recall that $\mathbb{E}[N^{(k)}] = P^{(k)}(1)$ where P is the probability generating function of N . So the result follows from standard calculus.

Suppose that $p \in (0, 1)$. The skewness and kurtosis of N are

1. $\text{skew}(N) = \frac{2-p}{\sqrt{1-p}}$
2. $\text{kurt}(N) = \frac{p^2}{1-p}$

Proof

The factorial moments can be used to find the moments of N about 0. The results then follow from the standard computational formulas for skewness and kurtosis.

Note that the geometric distribution is always positively skewed. Moreover, $\text{skew}(N) \rightarrow \infty$ and $\text{kurt}(N) \rightarrow \infty$ as $p \uparrow 1$.

Suppose now that $M = N - 1$, so that M (the number of failures before the first success) has the geometric distribution on \mathbb{N} . Then

1. $\mathbb{E}(M) = \frac{1-p}{p}$
2. $\text{var}(M) = \frac{1-p}{p^2}$
3. $\text{skew}(M) = \frac{2-p}{\sqrt{1-p}}$
4. $\text{kurt}(M) = \frac{p^2}{1-p}$
5. $\mathbb{E}(t^M) = \frac{p}{1-(1-p)t}$ for $|t| < \frac{1}{1-p}$

Of course, the fact that the variance, skewness, and kurtosis are unchanged follows easily, since N and M differ by a constant.

The Quantile Function

Let F denote the distribution function of N , so that $F(n) = 1 - (1-p)^n$ for $n \in \mathbb{N}$. Recall that $F^{-1}(r) = \min\{n \in \mathbb{N}_+ : F(n) \geq r\}$ for $r \in (0, 1)$ is the quantile function of N .

The quantile function of N is

$$F^{-1}(r) = \left\lceil \frac{\ln(1-r)}{\ln(1-p)} \right\rceil, \quad r \in (0, 1) \quad (11.3.21)$$

Of course, the quantile function, like the probability density function and the distribution function, completely determines the distribution of N . Moreover, we can compute the median and quartiles to get measures of center and spread.

The first quartile, the median (or second quartile), and the third quartile are

$$\begin{aligned} 1. F^{-1}\left(\frac{1}{4}\right) &= \lceil \ln(3/4) / \ln(1-p) \rceil \approx \lceil -0.2877 / \ln(1-p) \rceil \\ 2. F^{-1}\left(\frac{1}{2}\right) &= \lceil \ln(1/2) / \ln(1-p) \rceil \approx \lceil -0.6931 / \ln(1-p) \rceil \\ 3. F^{-1}\left(\frac{3}{4}\right) &= \lceil \ln(1/4) / \ln(1-p) \rceil \approx \lceil -1.3863 / \ln(1-p) \rceil \end{aligned}$$

Open the special distribution calculator, and select the geometric distribution and CDF view. Vary p and note the shape and location of the CDF/quantile function. For various values of p , compute the median and the first and third quartiles.

The Constant Rate Property

Suppose that T is a random variable taking values in \mathbb{N}_+ , which we interpret as the first time that some event of interest occurs.

The function h given by

$$h(n) = \mathbb{P}(T = n \mid T \geq n) = \frac{\mathbb{P}(T = n)}{\mathbb{P}(T \geq n)}, \quad n \in \mathbb{N}_+ \quad (11.3.22)$$

is the *rate function* of T .

If T is interpreted as the (discrete) lifetime of a device, then h is a discrete version of the *failure rate function* studied in reliability theory. However, in our usual formulation of Bernoulli trials, the event of interest is success rather than failure (or death), so we will simply use the term *rate function* to avoid confusion. The constant rate property characterizes the geometric distribution. As usual, let N denote the trial number of the first success in a sequence of Bernoulli trials with success parameter $p \in (0, 1)$, so that N has the geometric distribution on \mathbb{N}_+ with parameter p .

N has constant rate p .

Proof

From the results above, $\mathbb{P}(N = n) = p(1-p)^{n-1}$ and $\mathbb{P}(N \geq n) = \mathbb{P}(N > n-1) = (1-p)^{n-1}$, so $\mathbb{P}(N = n) / \mathbb{P}(N \geq n) = p$ for $n \in \mathbb{N}_+$.

Conversely, if T has constant rate $p \in (0, 1)$ then T has the geometric distribution on \mathbb{N}_+ with success parameter p .

Proof

Let $H(n) = \mathbb{P}(T \geq n)$ for $n \in \mathbb{N}_+$. From the constant rate property, $\mathbb{P}(T = n) = p H(n)$ for $n \in \mathbb{N}_+$. Next note that $\mathbb{P}(T = n) = H(n) - H(n+1)$ for $n \in \mathbb{N}_+$. Thus, H satisfies the recurrence relation $H(n+1) = (1-p)H(n)$ for $n \in \mathbb{N}_+$. Also H satisfies the initial condition $H(1) = 1$. Solving the recurrence relation gives $H(n) = (1-p)^{n-1}$ for $n \in \mathbb{N}_+$.

Relation to the Uniform Distribution

Suppose again that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter $p \in (0, 1)$. For $n \in \mathbb{N}_+$, recall that $Y_n = \sum_{i=1}^n X_i$, the number of successes in the first n trials, has the binomial distribution with parameters n and p . As before, N denotes the trial number of the first success.

Suppose that $n \in \mathbb{N}_+$. The conditional distribution of N given $Y_n = 1$ is uniform on $\{1, 2, \dots, n\}$.

Proof from sampling

We showed in the last section that given $Y_n = k$, the trial numbers of the successes form a random sample of size k chosen without replacement from $\{1, 2, \dots, n\}$. This result is a simple corollary with $k = 1$.

Direct proof

For $j \in \{1, 2, \dots, n\}$

$$\mathbb{P}(N = j \mid Y_n = 1) = \frac{\mathbb{P}(N = j, Y_n = 1)}{\mathbb{P}(Y_n = 1)} = \frac{\mathbb{P}(Y_{j-1} = 0, X_j = 1, Y_n - Y_j = 0)}{\mathbb{P}(Y_n = 1)} \quad (11.3.23)$$

In words, the events in the numerator of the last fraction are that there are no successes in the first $j-1$ trials, a success on trial j , and no successes in trials $j+1$ to n . These events are independent so

$$\mathbb{P}(N = j \mid Y_n = 1) = \frac{(1-p)^{j-1}p(1-p)^{n-j}}{np(1-p)^{n-1}} = \frac{1}{n} \quad (11.3.24)$$

Note that the conditional distribution does not depend on the success parameter p . If we know that there is exactly one success in the first n trials, then the trial number of that success is equally likely to be any of the n possibilities.

Another connection between the geometric distribution and the uniform distribution is given below in the [alternating coin tossing game](#): the conditional distribution of N given $N \leq n$ converges to the uniform distribution on $\{1, 2, \dots, n\}$ as $p \downarrow 0$.

Relation to the Exponential Distribution

The *Poisson process* on $[0, \infty)$, named for Simeon Poisson, is a model for random points in continuous time. There are many deep and interesting connections between the Bernoulli trials process (which can be thought of as a model for random points in *discrete* time) and the Poisson process. These connections are explored in detail in the chapter on the Poisson process. In this section we just give the most famous and important result—the convergence of the geometric distribution to the exponential distribution. The geometric distribution, as we know, governs the time of the first “random point” in the Bernoulli trials process, while the exponential distribution governs the time of the first random point in the Poisson process.

For reference, the exponential distribution with rate parameter $r \in (0, \infty)$ has distribution function $F(x) = 1 - e^{-rx}$ for $x \in [0, \infty)$. The mean of the exponential distribution is $1/r$ and the variance is $1/r^2$. In addition, the moment generating function is $s \mapsto \frac{1}{s-r}$ for $s > r$.

For $n \in \mathbb{N}_+$, suppose that U_n has the geometric distribution on \mathbb{N}_+ with success parameter $p_n \in (0, 1)$, where $np_n \rightarrow r > 0$ as $n \rightarrow \infty$. Then the distribution of U_n/n converges to the exponential distribution with parameter r as $n \rightarrow \infty$.

Proof

Let F_n denote the CDF of U_n/n . Then for $x \in [0, \infty)$

$$F_n(x) = \mathbb{P}\left(\frac{U_n}{n} \leq x\right) = \mathbb{P}(U_n \leq nx) = \mathbb{P}(U_n \leq \lfloor nx \rfloor) = 1 - (1 - p_n)^{\lfloor nx \rfloor} \quad (11.3.25)$$

But by a famous limit from calculus, $(1 - p_n)^n = \left(1 - \frac{np_n}{n}\right)^n \rightarrow e^{-r}$ as $n \rightarrow \infty$, and hence $(1 - p_n)^{nx} \rightarrow e^{-rx}$ as $n \rightarrow \infty$. But by definition, $\lfloor nx \rfloor \leq nx < \lfloor nx \rfloor + 1$ or equivalently, $nx - 1 < \lfloor nx \rfloor \leq nx$ so it follows that $(1 - p_n)^{\lfloor nx \rfloor} \rightarrow e^{-rx}$ as $n \rightarrow \infty$. Hence $F_n(x) \rightarrow 1 - e^{-rx}$ as $n \rightarrow \infty$, which is the CDF of the exponential distribution.

Note that the condition $np_n \rightarrow r$ as $n \rightarrow \infty$ is the same condition required for the convergence of the binomial distribution to the Poisson that we studied in the last section.

Special Families

The geometric distribution on \mathbb{N} is an infinitely divisible distribution and is a compound Poisson distribution. For the details, visit these individual sections and see the next section on the negative binomial distribution.

Examples and Applications

Simple Exercises

A standard, fair die is thrown until an ace occurs. Let N denote the number of throws. Find each of the following:

1. The probability density function of N .
2. The mean of N .
3. The variance of N .
4. The probability that the die will have to be thrown at least 5 times.

5. The quantile function of N .
6. The median and the first and third quartiles.

Answer

1. $\mathbb{P}(N = n) = \left(\frac{5}{6}\right)^{n-1} \frac{1}{6}$ for $n \in \mathbb{N}_+$
2. $\mathbb{E}(N) = 6$
3. $\text{var}(N) = 30$
4. $\mathbb{P}(N \geq 5) = 525/1296$
5. $F^{-1}(r) = \lceil \ln(1-r)/\ln(5/6) \rceil$ for $r \in (0, 1)$
6. Quartiles $q_1 = 2$, $q_2 = 4$, $q_3 = 8$

A type of missile has failure probability 0.02. Let N denote the number of launches before the first failure. Find each of the following:

1. The probability density function of N .
2. The mean of N .
3. The variance of N .
4. The probability of 20 consecutive successful launches.
5. The quantile function of N .
6. The median and the first and third quartiles.

Answer

1. $\mathbb{P}(N = n) = \left(\frac{49}{50}\right)^{n-1} \frac{1}{50}$ for $n \in \mathbb{N}_+$
2. $\mathbb{E}(N) = 50$
3. $\text{var}(N) = 2450$
4. $\mathbb{P}(N > 20) = 0.6676$
5. $F^{-1}(r) = \lceil \ln(1-r)/\ln(0.98) \rceil$ for $r \in (0, 1)$
6. Quartiles $q_1 = 15$, $q_2 = 35$, $q_3 = 69$

A student takes a multiple choice test with 10 questions, each with 5 choices (only one correct). The student blindly guesses and gets one question correct. Find the probability that the correct question was one of the first 4.

Answer

0.4

Recall that an American roulette wheel has 38 slots: 18 are red, 18 are black, and 2 are green. Suppose that you observe red or green on 10 consecutive spins. Give the conditional distribution of the number of additional spins needed for black to occur.

Answer

Geometric with $p = \frac{18}{38}$

The game of roulette is studied in more detail in the chapter on Games of Chance.

In the negative binomial experiment, set $k = 1$ to get the geometric distribution and set $p = 0.3$. Run the experiment 1000 times. Compute the appropriate relative frequencies and empirically investigate the memoryless property

$$\mathbb{P}(V > 5 \mid V > 2) = \mathbb{P}(V > 3) \quad (11.3.26)$$

The Petersburg Problem

We will now explore a gambling situation, known as the *Petersburg problem*, which leads to some famous and surprising results. Suppose that we are betting on a sequence of Bernoulli trials with success parameter $p \in (0, 1)$. We can bet any amount of money on a trial at *even stakes*: if the trial results in success, we receive that amount, and if the trial results in failure, we must pay that amount. We will use the following strategy, known as a *martingale strategy*:

1. We bet c units on the first trial.
2. Whenever we lose a trial, we double the bet for the next trial.
3. We stop as soon as we win a trial.

Let N denote the number of trials played, so that N has the geometric distribution with parameter p , and let W denote our net winnings when we stop.

$$W = c$$

Proof

The first win occurs on trial N , so the initial bet was doubled $N - 1$ times. The net winnings are

$$W = -c \sum_{i=0}^{N-2} 2^i + c2^{N-1} = c(1 - 2^{N-1} + 2^{N-1}) = c \quad (11.3.27)$$

Thus, W is not random and W is independent of p ! Since c is an arbitrary constant, it would appear that we have an ideal strategy. However, let us study the amount of money Z needed to play the strategy.

$$Z = c(2^N - 1)$$

The expected amount of money needed for the martingale strategy is

$$\mathbb{E}(Z) = \begin{cases} \frac{c}{2p-1}, & p > \frac{1}{2} \\ \infty, & p \leq \frac{1}{2} \end{cases} \quad (11.3.28)$$

Thus, the strategy is fatally flawed when the trials are unfavorable and even when they are fair, since we need infinite expected capital to make the strategy work in these cases.

Compute $\mathbb{E}(Z)$ explicitly if $c = 100$ and $p = 0.55$.

Answer

\$1000

In the negative binomial experiment, set $k = 1$. For each of the following values of p , run the experiment 100 times. For each run compute Z (with $c = 1$). Find the average value of Z over the 100 runs:

1. $p = 0.2$
2. $p = 0.5$
3. $p = 0.8$

For more information about gambling strategies, see the section on Red and Black. Martingales are studied in detail in a separate chapter.

The Alternating Coin-Tossing Game

A coin has probability of heads $p \in (0, 1]$. There are n players who take turns tossing the coin in round-robin style: player 1 first, then player 2, continuing until player n , then player 1 again, and so forth. The first player to toss heads wins the game.

Let N denote the number of the first toss that results in heads. Of course, N has the geometric distribution on \mathbb{N}_+ with parameter p . Additionally, let W denote the winner of the game; W takes values in the set $\{1, 2, \dots, n\}$. We are interested in the probability distribution of W .

For $i \in \{1, 2, \dots, n\}$, $W = i$ if and only if $N = i + kn$ for some $k \in \mathbb{N}$. That is, using modular arithmetic,

$$W = [(N - 1) \bmod n] + 1 \quad (11.3.29)$$

The winning player W has probability density function

$$\mathbb{P}(W = i) = \frac{p(1-p)^{i-1}}{1 - (1-p)^n}, \quad i \in \{1, 2, \dots, n\} \quad (11.3.30)$$

Proof

This follows from the previous exercise and the geometric distribution of N .

$$\mathbb{P}(W = i) = (1-p)^{i-1} \mathbb{P}(W = 1) \text{ for } i \in \{1, 2, \dots, n\}.$$

Proof

This result can be argued directly, using the memoryless property of the geometric distribution. In order for player i to win, the previous $i-1$ players must first all toss tails. Then, player i effectively becomes the first player in a new sequence of tosses. This result can be used to give another derivation of the probability density function in the previous exercise.

Note that $\mathbb{P}(W = i)$ is a decreasing function of $i \in \{1, 2, \dots, n\}$. Not surprisingly, the lower the toss order the better for the player.

Explicitly compute the probability density function of W when the coin is fair ($p = 1/2$).

Answer

$$\mathbb{P}(W = i) = 2^{n-1} / (2^n - 1), \quad i \in \{1, 2, \dots, n\}$$

Note from the [result above](#) that W itself has a *truncated geometric distribution*.

The distribution of W is the same as the conditional distribution of N given $N \leq n$:

$$\mathbb{P}(W = i) = \mathbb{P}(N = i \mid N \leq n), \quad i \in \{1, 2, \dots, n\} \quad (11.3.31)$$

The following problems explore some limiting distributions related to the alternating coin-tossing game.

For fixed $p \in (0, 1]$, the distribution of W converges to the geometric distribution with parameter p as $n \uparrow \infty$.

For fixed n , the distribution of W converges to the uniform distribution on $\{1, 2, \dots, n\}$ as $p \downarrow 0$.

Players at the end of the tossing order should hope for a coin biased towards tails.

Odd Man Out

In the game of *odd man out*, we start with a specified number of players, each with a coin that has the same probability of heads. The players toss their coins at the same time. If there is an *odd man*, that is a player with an outcome different than all of the other players, then the odd player is eliminated; otherwise no player is eliminated. In any event, the remaining players continue the game in the same manner. A slight technical problem arises with just two players, since different outcomes would make both players “odd”. So in this case, we might (arbitrarily) make the player with tails the odd man.

Suppose there are $k \in \{2, 3, \dots\}$ players and $p \in [0, 1]$. In a single round, the probability of an odd man is

$$r_k(p) = \begin{cases} 2p(1-p), & k = 2 \\ kp(1-p)^{k-1} + kp^{k-1}(1-p), & k \in \{3, 4, \dots\} \end{cases} \quad (11.3.32)$$

Proof

Let Y denote the number of heads. If $k = 2$, the event that there is an odd man is $\{Y = 1\}$. If $k \geq 3$, the event that there is an odd man is $\{Y \in \{1, k-1\}\}$. The result now follows since Y has a binomial distribution with parameters k and p .

The graph of r_k is more interesting than you might think.




Figure 11.3.1: The graphs of r_k for $k \in \{3, 4, 5, 6\}$

For $k \in \{2, 3, \dots\}$, r_k has the following properties:

1. $r_k(0) = r_k(1) = 0$
2. r_k is symmetric about $p = \frac{1}{2}$
3. For fixed $p \in [0, 1]$, $r_k(p) \rightarrow 0$ as $k \rightarrow \infty$.

Proof

These properties are clear from the functional form of $r_k(p)$. Note that $r_k(p) = r_k(1-p)$.

For $k \in \{2, 3, 4\}$, r_k has the following properties:

1. r_k increases and then decreases, with maximum at $p = \frac{1}{2}$.
2. r_k is concave downward

Proof

This follows by computing the first derivatives: $r'_2(p) = 2(1-2p)$, $r'_3(p) = 3(1-2p)$, $r'_4(p) = 4(1-2p)^3$, and the second derivatives: $r''_2(p) = -4$, $r''_3(p) = -6$, $r''_4(p) = -24(1-2p)^2$.

For $k \in \{5, 6, \dots\}$, r_k has the following properties:

1. The maximum occurs at two points of the form p_k and $1-p_k$ where $p_k \in (0, \frac{1}{2})$ and $p_k \rightarrow 0$ as $k \rightarrow \infty$.
2. The maximum value $r_k(p_k) \rightarrow 1/e \approx 0.3679$ as $k \rightarrow \infty$.
3. The graph has a local minimum at $p = \frac{1}{2}$.

Proof sketch

Note that $r_k(p) = s_k(p) + s_k(1-p)$ where $s_k(t) = kt^{k-1}(1-t)$ for $t \in [0, 1]$. Also, $p \mapsto s_k(p)$ is the dominant term when $p > \frac{1}{2}$ while $p \mapsto s_k(1-p)$ is the dominant term when $p < \frac{1}{2}$. A simple analysis of the derivative shows that s_k increases and then decreases, reaching its maximum at $(k-1)/k$. Moreover, the maximum value is $s_k[(k-1)/k] = (1-1/k)^{k-1} \rightarrow e^{-1}$ as $k \rightarrow \infty$. Also, s_k is concave upward and then downward, with inflection point at $(k-2)/k$.

Suppose $p \in (0, 1)$, and let N_k denote the number of rounds until an odd man is eliminated, starting with k players. Then N_k has the geometric distribution on \mathbb{N}_+ with parameter $r_k(p)$. The mean and variance are

1. $\mu_k(p) = 1/r_k(p)$
2. $\sigma_k^2(p) = [1 - r_k(p)] / r_k^2(p)$

As we might expect, $\mu_k(p) \rightarrow \infty$ and $\sigma_k^2(p) \rightarrow \infty$ as $k \rightarrow \infty$ for fixed $p \in (0, 1)$. On the other hand, from the result [above](#), $\mu_k(p_k) \rightarrow e$ and $\sigma_k^2(p_k) \rightarrow e^2 - e$ as $k \rightarrow \infty$.

Suppose we start with $k \in \{2, 3, \dots\}$ players and $p \in (0, 1)$. The number of rounds until a single player remains is $M_k = \sum_{j=2}^k N_j$ where (N_2, N_3, \dots, N_k) are independent and N_j has the geometric distribution on \mathbb{N}_+ with parameter $r_j(p)$. The mean and variance are

1. $\mathbb{E}(M_k) = \sum_{j=2}^k 1/r_j(p)$
2. $\text{var}(M_k) = \sum_{j=2}^k [1 - r_j(p)] / r_j^2(p)$

Proof

The form of M_k follows from the previous result: N_k is the number of rounds until the first player is eliminated. Then the game continues independently with $k-1$ players, so N_{k-1} is the number of additional rounds until the second player is eliminated, and so forth. Parts (a) and (b) follow from the previous result and standard properties of expected value and variance.

Starting with k players and probability of heads $p \in (0, 1)$, the total number of coin tosses is $T_k = \sum_{j=2}^k jN_j$. The mean and variance are

1. $\mathbb{E}(T_k) = \sum_{j=2}^k j/r_j(p)$
2. $\text{var}(T_k) = \sum_{j=2}^k j^2 [1 - r_j(p)] / r_j^2(p)$

Proof

As before, the form of M_k follows from [result above](#): N_k is the number of rounds until the first player is eliminated, and each these rounds has k tosses. Then the game continues independently with $k-1$ players, so N_{k-1} is the number of additional rounds until the second player is eliminated with each round having $k-1$ tosses, and so forth. Parts (a) and (b) also follow from the [result above](#) and standard properties of expected value and variance.

Number of Trials Before a Pattern

Consider again a sequence of Bernoulli trials $\mathbf{X} = (X_1, X_2, \dots)$ with success parameter $p \in (0, 1)$. Recall that the number of trials M before the first success (outcome 1) occurs has the geometric distribution on \mathbb{N} with parameter p . A natural generalization is the random variable that gives the number of trials before a specific finite sequence of outcomes occurs for the first time. (Such a sequence is sometimes referred to as a *word* from the *alphabet* $\{0, 1\}$ or simply a *bit string*). In general, finding the distribution of this variable is a difficult problem, with the difficulty depending very much on the nature of the word. The problem of finding just the *expected* number of trials before a word occurs can be solved using powerful tools from the theory of renewal processes and from the theory of martingales.

To set up the notation, let \mathbf{x} denote a finite bit string and let $M_{\mathbf{x}}$ denote the number of trials before \mathbf{x} occurs for the first time. Finally, let $q = 1 - p$. Note that $M_{\mathbf{x}}$ takes values in \mathbb{N} . In the following exercises, we will consider $\mathbf{x} = 10$, a success followed by a failure. As always, try to derive the results yourself before looking at the proofs.

The probability density function f_{10} of M_{10} is given as follows:

1. If $p \neq \frac{1}{2}$ then

$$f_{10}(n) = pq \frac{p^{n+1} - q^{n+1}}{p - q}, \quad n \in \mathbb{N} \quad (11.3.33)$$

2. If $p = \frac{1}{2}$ then $f_{10}(n) = (n+1)\left(\frac{1}{2}\right)^{n+2}$ for $n \in \mathbb{N}$.

Proof

For $n \in \mathbb{N}$, the event $\{M_{10} = n\}$ can only occur if there is an initial string of 0s of length $k \in \{0, 1, \dots, n\}$ followed by a string of 1s of length $n-k$ and then 1 on trial $n+1$ and 0 on trial $n+2$. Hence

$$f_{10}(n) = \mathbb{P}(M_{10} = n) = \sum_{k=0}^n q^k p^{n-k} pq, \quad n \in \mathbb{N} \quad (11.3.34)$$

The stated result then follows from standard results on geometric series.

It's interesting to note that f is symmetric in p and q , that is, symmetric about $p = \frac{1}{2}$. It follows that the distribution function, probability generating function, expected value, and variance, which we consider below, are all also symmetric about $p = \frac{1}{2}$. It's also interesting to note that $f_{10}(0) = f_{10}(1) = pq$, and this is the largest value. So regardless of $p \in (0, 1)$ the distribution is bimodal with modes 0 and 1.

The distribution function F_{10} of M_{10} is given as follows:

1. If $p \neq \frac{1}{2}$ then

$$F_{10}(n) = 1 - \frac{p^{n+3} - q^{n+3}}{p - q}, \quad n \in \mathbb{N} \quad (11.3.35)$$

2. If $p = \frac{1}{2}$ then $F_{10} = 1 - (n+3)\left(\frac{1}{2}\right)^{n+2}$ for $n \in \mathbb{N}$.

Proof

By definition, $F_{10}(n) = \sum_{k=0}^n f_{10}(k)$ for $n \in \mathbb{N}$. The stated result then follows from the previous theorem, standard results on geometric series, and some algebra.

The probability generating function P_{10} of M_{10} is given as follows:

1. If $p \neq \frac{1}{2}$ then

$$P_{10}(t) = \frac{pq}{p-q} \left(\frac{p}{1-tp} - \frac{q}{1-tq} \right), \quad |t| < \min\{1/p, 1/q\} \quad (11.3.36)$$

2. If $p = \frac{1}{2}$ then $P_{10}(t) = 1/(t-2)^2$ for $|t| < 2$

Proof

By definition,

$$P_{10}(t) = \mathbb{E}(t^{M_{10}}) = \sum_{n=0}^{\infty} f_{10}(n)t^n \quad (11.3.37)$$

for all $t \in \mathbb{R}$ for which the series converges absolutely. The stated result then follows from the theorem above, and once again, standard results on geometric series.

The mean of M_{10} is given as follows:

1. If $p \neq \frac{1}{2}$ then

$$\mathbb{E}(M_{10}) = \frac{p^4 - q^4}{pq(p-q)} \quad (11.3.38)$$

2. If $p = \frac{1}{2}$ then $\mathbb{E}(M_{10}) = 2$.

Proof

Recall that $\mathbb{E}(M_{10}) = P'_{10}(1)$ so the stated result follows from calculus, using the previous theorem on the probability generating function. The mean can also be computed from the definition $\mathbb{E}(M_{10}) = \sum_{n=0}^{\infty} n f_{10}(n)$ using standard results from geometric series, but this method is more tedious.

The graph of $\mathbb{E}(M_{10})$ as a function of $p \in (0, 1)$ is given below. It's not surprising that $\mathbb{E}(M_{10}) \rightarrow \infty$ as $p \downarrow 0$ and as $p \uparrow 1$, and that the minimum value occurs when $p = \frac{1}{2}$.




Figure 11.3.2: $\mathbb{E}(M_{10})$ as a function of p

The variance of M_{10} is given as follows:

1. If $p \neq \frac{1}{2}$ then

$$\text{var}(M_{10}) = \frac{2}{p^2 q^2} \left(\frac{p^6 - q^6}{p - q} \right) + \frac{1}{pq} \left(\frac{p^4 - q^4}{p - q} \right) - \frac{1}{p^2 q^2} \left(\frac{p^4 - q^4}{p - q} \right)^2 \quad (11.3.39)$$

2. If $p = \frac{1}{2}$ then $\text{var}(M_{10}) = 4$.

Proof

Recall that $P''_{10}(1) = \mathbb{E}[M_{10}(M_{10} - 1)]$, the second factorial moment, and so

$$\text{var}(M_{10}) = P''_{10}(1) + P'_{10}(1) - [P'_{10}(1)]^2 \quad (11.3.40)$$

The stated result then follows from calculus and the theorem above giving the probability generating function.

This page titled [11.3: The Geometric Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.4: The Negative Binomial Distribution

Basic Theory

Suppose again that our random experiment is to perform a sequence of Bernoulli trials $\mathbf{X} = (X_1, X_2, \dots)$ with success parameter $p \in (0, 1]$. Recall that the number of successes in the first n trials

$$Y_n = \sum_{i=1}^n X_i \quad (11.4.1)$$

has the binomial distribution with parameters n and p . In this section we will study the random variable that gives the *trial number of the k th success*:

$$V_k = \min \{n \in \mathbb{N}_+ : Y_n = k\} \quad (11.4.2)$$

Note that V_1 is the number of trials needed to get the first success, which we now know has the geometric distribution on \mathbb{N}_+ with parameter p .

The Probability Density Function

The probability distribution of V_k is given by

$$\mathbb{P}(V_k = n) = \binom{n-1}{k-1} p^k (1-p)^{n-k}, \quad n \in \{k, k+1, k+2, \dots\} \quad (11.4.3)$$

Proof

Note that $V_k = n$ if and only if $X_n = 1$ and $Y_{n-1} = k-1$. Hence, from independence and the binomial distribution,

$$\mathbb{P}(V_k = n) = \mathbb{P}(Y_{n-1} = k-1) \mathbb{P}(X_n = 1) = \binom{n-1}{k-1} p^{k-1} (1-p)^{(n-1)-(k-1)} p = \binom{n-1}{k-1} p^k (1-p)^{n-k} \quad (11.4.4)$$

The distribution defined by the density function in (1) is known as the *negative binomial distribution*; it has two parameters, the *stopping parameter* k and the *success probability* p .

In the negative binomial experiment, vary k and p with the scroll bars and note the shape of the density function. For selected values of k and p , run the experiment 1000 times and compare the relative frequency function to the probability density function.

The binomial and negative binomial sequences are inverse to each other in a certain sense.

For $n \in \mathbb{N}_+$ and $k \in \{0, 1, \dots, n\}$,

1. $Y_n \geq k \iff V_k \leq n$ and hence $\mathbb{P}(Y_n \geq k) = \mathbb{P}(V_k \leq n)$
2. $k \mathbb{P}(Y_n = k) = n \mathbb{P}(V_k = n)$

Proof

1. The events $\{Y_n \geq k\}$ and $\{V_k \leq n\}$ both mean that there are at least k successes in the first n Bernoulli trials.
2. From the formulas for the binomial and negative binomial PDFs, $k \mathbb{P}(Y_n = k)$ and $n \mathbb{P}(V_k = n)$ both simplify to $\frac{n!}{(k-1)!(n-k)!} p^k (1-p)^{n-k}$.

In particular, it follows from part (a) that any event that can be expressed in terms of the negative binomial variables can also be expressed in terms of the binomial variables.

The negative binomial distribution is unimodal. Let $t = 1 + \frac{k-1}{p}$. Then

1. $\mathbb{P}(V_k = n) > \mathbb{P}(V_k = n-1)$ if and only if $n < t$.
2. The probability density function at first increases and then decreases, reaching its maximum value at $\lfloor t \rfloor$.
3. There is a single mode at $\lfloor t \rfloor$ if t is not an integer, and two consecutive modes at $t-1$ and t if t is an integer.

Times Between Successes

Next we will define the random variables that give the *number of trials between successive successes*. Let $U_1 = V_1$ and $U_k = V_k - V_{k-1}$ for $k \in \{2, 3, \dots\}$

$\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent random variables, each having the geometric distribution on \mathbb{N}_+ with parameter p . Moreover,

$$V_k = \sum_{i=1}^k U_i \quad (11.4.5)$$

In statistical terms, \mathbf{U} corresponds to sampling from the geometric distribution with parameter p , so that for each k , (U_1, U_2, \dots, U_k) is a random sample of size k from this distribution. The sample mean corresponding to this sample is V_k/k ; this random variable gives the average number of trials between the first k successes. In probability terms, the sequence of negative binomial variables \mathbf{V} is the *partial sum process* corresponding to the sequence \mathbf{U} . Partial sum processes are studied in more generality in the chapter on Random Samples.

The random process $\mathbf{V} = (V_1, V_2, \dots)$ has *stationary, independent increments*:

1. If $j < k$ then $V_k - V_j$ has the same distribution as V_{k-j} , namely negative binomial with parameters $k-j$ and p .
2. If $k_1 < k_2 < k_3 < \dots$ then $(V_{k_1}, V_{k_2} - V_{k_1}, V_{k_3} - V_{k_2}, \dots)$ is a sequence of independent random variables.

Actually, any partial sum process corresponding to an independent, identically distributed sequence will have stationary, independent increments.

Basic Properties

The mean, variance and probability generating function of V_k can be computed in several ways. The method using the representation as a sum of independent, identically distributed geometrically distributed variables is the easiest.

V_k has probability generating function P given by

$$P(t) = \left(\frac{pt}{1 - (1-p)t} \right)^k, \quad |t| < \frac{1}{1-p} \quad (11.4.6)$$

Proof

Recall that the probability generating function of a sum of independent variables is the product of the probability generating functions of the variables. Recall also, the probability generating function of the geometric distribution with parameter p is $t \mapsto pt / [1 - (1-p)t]$. Thus, the result follows immediately from the [sum representation](#) above. A derivation can also be given directly from the probability density function.

The mean and variance of V_k are

1. $\mathbb{E}(V_k) = k \frac{1}{p}$.
2. $\text{var}(V_k) = k \frac{1-p}{p^2}$

Proof

The geometric distribution with parameter p has mean $1/p$ and variance $(1-p)/p^2$, so the results follow immediately from the [sum representation](#) above. Recall that the mean of a sum is the sum of the means, and the variance of the sum of *independent* variables is the sum of the variances. These results can also be proven directly from the probability density function or from the probability generating function.

In the negative binomial experiment, vary k and p with the scroll bars and note the location and size of the mean/standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

Suppose that V and W are independent random variables for an experiment, and that V has the negative binomial distribution with parameters j and p , and W has the negative binomial distribution with parameters k and p . Then $V + W$ has the negative binomial distribution with parameters $k+j$ and p .

Proof

Once again, the simplest proof is based on the representation as a sum of independent geometric variables. In the context of the [sum representation](#) above, we can take $V = V_j$ and $W = V_{k+j} - V_j$, so that $V + W = V_{k+j}$. Another simple proof uses probability generating functions. Recall again that the PGF of the sum of independent variables is the product of the PGFs. Finally, a difficult proof can be constructed using probability density functions. Recall that the PDF of a sum of independent variables is the convolution of the PDFs.

Normal Approximation

In the negative binomial experiment, start with various values of p and $k = 1$. Successively increase k by 1, noting the shape of the probability density function each time.

Even though you are limited to $k = 5$ in the app, you can still see the characteristic bell shape. This is a consequence of the central limit theorem because the negative binomial variable can be written as a sum of k independent, identically distributed (geometric) random variables.

The standard score of V_k is

$$Z_k = \frac{p V_k - k}{\sqrt{k(1-p)}} \quad (11.4.7)$$

The distribution of Z_k converges to the standard normal distribution as $k \rightarrow \infty$.

From a practical point of view, this result means that if k is “large”, the distribution of V_k is approximately normal with mean $k\frac{1}{p}$ and variance $k\frac{1-p}{p^2}$. Just how large k needs to be for the approximation to work well depends on p . Also, when using the normal approximation, we should remember to use the continuity correction, since the negative binomial is a discrete distribution.

Relation to Order Statistics

Suppose that $n \in \mathbb{N}_+$ and $k \in \{1, 2, \dots, n\}$, and let $L = \{(n_1, n_2, \dots, n_k) \in \{1, 2, \dots, n\}^k : n_1 < n_2 < \dots < n_k\}$. Then

$$\mathbb{P}(V_1 = n_1, V_2 = n_2, \dots, V_k = n_k \mid Y_n = k) = \frac{1}{\binom{n}{k}}, \quad (n_1, n_2, \dots, n_k) \in L \quad (11.4.8)$$

Proof

$$\mathbb{P}(V_1 = n_1, V_2 = n_2, \dots, V_k = n_k \mid Y_n = k) = \frac{\mathbb{P}(V_1 = n_1, V_2 = n_2, \dots, V_k = n_k, Y_n = k)}{\mathbb{P}(Y_n = k)} = \frac{p^k(1-p)^{n-k}}{\binom{n}{k}p^k(1-p)^{n-k}} = \frac{1}{\binom{n}{k}} \quad (11.4.9)$$

Note that the event in the numerator of the first fraction means that in the first n trials, successes occurred at trials n_1, n_2, \dots, n_k and failures occurred at all other trials.

Thus, given exactly k successes in the first n trials, the vector of success trial numbers is uniformly distributed on the set of possibilities L , regardless of the value of the success parameter p . Equivalently, the vector of success trial numbers is distributed as the vector of order statistics corresponding to a sample of size k chosen at random and without replacement from $\{1, 2, \dots, n\}$.

Suppose that $n \in \mathbb{N}_+$, $k \in \{1, 2, \dots, n\}$, and $j \in \{1, 2, \dots, k\}$. Then

$$\mathbb{P}(V_j = m \mid Y_n = k) = \frac{\binom{m-1}{j-1} \binom{n-m}{k-j}}{\binom{n}{k}}, \quad m \in \{j, j+1, \dots, n+k-j\} \quad (11.4.10)$$

Proof

This follows immediately from the previous result and a theorem in the section on order statistics. However, a direct proof is also easy. Note that the event $\{V_j = m, Y_n = k\}$ means that there were $j-1$ successes in the first $m-1$ trials, a success on trial m and $k-j$ success in trials $m+1$ to n . Hence using the binomial distribution and independence,

$$\mathbb{P}(V_j = m \mid Y_n = k) = \frac{\mathbb{P}(V_j = m, Y_n = k)}{\mathbb{P}(Y_n = k)} = \frac{\binom{m-1}{j-1} p^{j-1} (1-p)^{(m-1)-(j-1)} p \binom{n-m}{k-j} p^{k-j} (1-p)^{(n-m)-(k-j)}}{\binom{n}{k} p^k (1-p)^{n-k}} \quad (11.4.11)$$

$$= \frac{\binom{m-1}{j-1} \binom{n-m}{k-j} p^k (1-p)^{n-k}}{\binom{n}{k} p^k (1-p)^{n-k}} = \frac{\binom{m-1}{j-1} \binom{n-m}{k-j}}{\binom{n}{k}}, \quad (11.4.12)$$

Thus, given exactly k successes in the first n trials, the trial number of the j th success has the same distribution as the j th order statistic when a sample of size k is selected at random and without replacement from the population $\{1, 2, \dots, n\}$. Again, this result does not depend on the value of the success parameter p . The following theorem gives the mean and variance of the conditional distribution.


Suppose again that $n \in \mathbb{N}_+$, $k \in \{1, 2, \dots, n\}$, and $j \in \{1, 2, \dots, k\}$. Then

1. $\mathbb{E}(V_j \mid Y_n = k) = j\frac{n+1}{k+1}$
2. $\text{var}(V_j \mid Y_n = k) = j(k-j+1)\frac{(n+1)(n-k)}{(k+1)^2(k+2)}$

Proof

These moment results follow immediately from the previous theorem and a theorem in the section on order statistics. However, there is also a nice heuristic argument for (a) using indicator variables. Given $Y_n = k$, the k successes divide the set of indices where the failures occur into $k+1$ disjoint sets (some may be empty, of course, if there are adjacent successes).

The red dots are successes and the green dots failures. The 8 successes in the 50 trials divide the set of failures into 9 disjoint sets.



Let I_i take the value 1 if the i th failure occurs before the j th success, and 0 otherwise, for $i \in \{1, 2, \dots, n-k\}$. Then given $Y_n = k$,

$$V_j = j + \sum_{i=1}^{n-k} I_i \quad (11.4.13)$$

Given $Y_n = k$, we know that the k successes and $n-k$ failures are randomly placed in $\{1, 2, \dots, n\}$ with each possible configuration having the same probability. Thus,

$$\mathbb{E}(I_i | Y_n = k) = \mathbb{P}(I_i = 1 | Y_n = k) = \frac{j}{k+1}, \quad i \in \{1, 2, \dots, n-k\} \quad (11.4.14)$$

Hence

$$\mathbb{E}(V_j | Y_n = k) = j + (n-k) \frac{j}{k+1} = j \frac{n+1}{k+1} \quad (11.4.15)$$

Examples and Applications

Coins, Dice and Other Gadgets

A standard, fair die is thrown until 3 aces occur. Let V denote the number of throws. Find each of the following:

1. The probability density function of V .
2. The mean of V .
3. The variance of V .
4. The probability that at least 20 throws will be needed.

Answer

1. $\mathbb{P}(V = n) = \binom{n-1}{2} \left(\frac{1}{6}\right)^3 \left(\frac{5}{6}\right)^{n-3}, \quad n \in \{3, 4, \dots\}$
2. $\mathbb{E}(V) = 18$
3. $\text{var}(V) = 90$
4. $\mathbb{P}(V \geq 20) = 0.3643$

A coin is tossed repeatedly. The 10th head occurs on the 25th toss. Find each of the following:

1. The probability density function of the trial number of the 5th head.
2. The mean of the distribution in (a).
3. The variance of the distribution in (a).

Answer

1. $\mathbb{P}(V_5 = m | V_{10} = 25) = \frac{\binom{m-1}{4} \binom{24-m}{4}}{\binom{24}{9}}, \quad m \in \{5, 6, \dots, 25\}$
2. $\mathbb{E}(V_5 | V_{10} = 25) = \frac{25}{2}$
3. $\text{var}(V_5 | V_{10} = 25) = \frac{375}{44}$

A certain type of missile has failure probability 0.02. Let N denote the launch number of the fourth failure. Find each of the following:

1. The probability density function of N .
2. The mean of N .
3. The variance of N .
4. The probability that there will be at least 4 failures in the first 200 launches.

Answer

1. $\mathbb{P}(N = n) = \binom{n-1}{3} \left(\frac{1}{50}\right)^4 \left(\frac{49}{50}\right)^{n-4}, \quad n \in \{4, 5, \dots\}$
2. $\mathbb{E}(N) = 200$
3. $\text{var}(N) = 9800$
4. $\mathbb{P}(N \leq 200) = 0.5685$

In the negative binomial experiment, set $p = 0.5$ and $k = 5$. Run the experiment 1000 times. Compute and compare each of the following:

1. $\mathbb{P}(8 \leq V_5 \leq 15)$
2. The relative frequency of the event $\{8 \leq V_5 \leq 15\}$ in the simulation
3. The normal approximation to $\mathbb{P}(8 \leq V_5 \leq 15)$

Answer

1. $\mathbb{P}(8 \leq V_5 \leq 15) = 0.7142$
3. $\mathbb{P}(8 \leq V_5 \leq 15) \approx 0.7445$

A coin is tossed until the 50th head occurs.

1. Assuming that the coin is fair, find the normal approximation of the probability that the coin is tossed at least 125 times.
2. Suppose that you perform this experiment, and 125 tosses are required. Do you believe that the coin is fair?

Answer

1. 0.0072
2. No.

The Banach Match Problem

Suppose that an absent-minded professor (is there any other kind?) has m matches in his right pocket and m matches in his left pocket. When he needs a match to light his pipe, he is equally likely to choose a match from either pocket. We want to compute the probability density function of the random variable W that gives the number of matches remaining when the professor first discovers that one of the pockets is empty. This is known as the *Banach match problem*, named for the mathematician Stefan Banach, who evidently behaved in the manner described.

We can recast the problem in terms of the negative binomial distribution. Clearly the match choices form a sequence of Bernoulli trials with parameter $p = \frac{1}{2}$. Specifically, we can consider a match from the right pocket as a win for player R , and a match from the left pocket as a win for player L . In a hypothetical infinite sequence of trials, let U denote the number of trials necessary for R to win $m + 1$ trials, and V the number of trials necessary for L to win $m + 1$ trials. Note that U and V each have the negative binomial distribution with parameters $m + 1$ and p .

For $k \in \{0, 1, \dots, m\}$,

1. L has $m - k$ wins at the moment when R wins $m + 1$ games if and only if $U = 2m - k + 1$.
2. $\{U = 2m - k + 1\}$ is equivalent to the event that the professor first discovers that the right pocket is empty and that the left pocket has k matches
3. $\mathbb{P}(U = 2m - k + 1) = \binom{2m-k}{m} \left(\frac{1}{2}\right)^{2m-k+1}$

For $k \in \{0, 1, \dots, m\}$,

1. R has $m - k$ wins at the moment when L wins $m + 1$ games if and only if $V = 2m - k + 1$.
2. $\{V = 2m - k + 1\}$ is equivalent to the event that the professor first discovers that the right pocket is empty and that the left pocket has k matches
3. $\mathbb{P}(V = 2m - k + 1) = \binom{2m-k}{m} \left(\frac{1}{2}\right)^{2m-k+1}$.

W has probability density function

$$\mathbb{P}(W = k) = \binom{2m-k}{m} \left(\frac{1}{2}\right)^{2m-k}, \quad k \in \{0, 1, \dots, m\} \quad (11.4.16)$$

Proof

This result follows from the previous two exercises, since $\mathbb{P}(W = k) = \mathbb{P}(U = 2m - k + 1) + \mathbb{P}(V = 2m - k + 1)$.

We can also solve the *non-symmetric Banach match problem*, using the same methods as above. Thus, suppose that the professor reaches for a match in his right pocket with probability p and in his left pocket with probability $1 - p$, where $0 < p < 1$. The essential change in the analysis is that U has the negative binomial distribution with parameters $m + 1$ and p , while V has the negative binomial distribution with parameters $m + 1$ and $1 - p$.

For the Banach match problem with parameter p , W has probability density function

$$\mathbb{P}(W = k) = \binom{2m-k}{m} [p^{m+1}(1-p)^{m-k} + (1-p)^{m+1}p^{m-k}], \quad k \in \{0, 1, \dots, m\} \quad (11.4.17)$$

The Problem of Points

Suppose that two teams (or individuals) A and B play a sequence of Bernoulli trials (which we will also call *points*), where $p \in (0, 1)$ is the probability that player A wins a point. For nonnegative integers n and m , let $A_{n,m}(p)$ denote the probability that A wins n points before B wins m points. Computing $A_{n,m}(p)$ is an historically famous problem, known as the *problem of points*, that was solved by Pierre de Fermat and by Blaise Pascal.

Comment on the validity of the Bernoulli trial assumptions (independence of trials and constant probability of success) for games of sport that have a *skill* component as well as a *random* component.

There is an easy solution to the problem of points using the binomial distribution; this was essentially Pascal's solution. There is also an easy solution to the problem of points using the negative binomial distribution. In a sense, this has to be the case, given the equivalence between the binomial and negative binomial processes in (3). First, let us pretend that the trials go on forever, regardless of the outcomes. Let Y_{n+m-1} denote the number of wins by player A in the first $n + m - 1$ points, and let V_n denote the number of trials needed for A to win n points. By definition, Y_{n+m-1} has the binomial distribution with parameters $n + m - 1$ and p , and V_n has the negative binomial distribution with parameters n and p .

Player A wins n points before B wins m points if and only if $Y_{n+m-1} \geq n$ if and only if $V_n \leq m+n-1$. Hence

$$A_{n,m}(p) = \sum_{k=n}^{n+m-1} \binom{n+m-1}{k} p^k (1-p)^{n+m-1-k} = \sum_{j=n}^{n+m-1} \binom{j-1}{n-1} p^n (1-p)^{j-n} \quad (11.4.18)$$

$A_{n,m}(p)$ satisfies the following properties:

1. $A_{n,m}(p)$ increases from 0 to 1 as p increases from 0 to 1 for fixed n and m .
2. $A_{n,m}(p)$ decreases as n increases for fixed m and p .
3. $A_{n,m}(p)$ increases as m increases for fixed n and p .

$1 - A_{n,m}(p) = A_{m,n}(1-p)$ for any $m, n \in \mathbb{N}_+$ and $p \in (0, 1)$.

Proof

A simple probabilistic proof is to note that both sides can be interpreted as the probability that a player with point probability $1-p$ wins m points before the opponent wins n points. An analytic proof can also be constructed using the [formulas above](#) for $A_{n,m}(p)$

In the problem of points experiments, vary the parameters n , m , and p , and note how the probability changes. For selected values of the parameters, run the simulation 1000 times and note the apparent convergence of the relative frequency to the probability.

The win probability function for player A satisfies the following recurrence relation and boundary conditions (this was essentially Fermat's solution):

1. $A_{m,n}(p) = p A_{n-1,m}(p) + (1-p) A_{n,m-1}(p)$, $n, m \in \mathbb{N}_+$
2. $A_{n,0}(p) = 0$, $A_{0,m}(p) = 1$

Proof

Condition on the outcome of the first trial.

Next let $N_{n,m}$ denote the number of trials needed until either A wins n points or B wins m points, whichever occurs first—the length of the problem of points experiment. The following result gives the distribution of $N_{n,m}$

For $k \in \{\min\{m, n\}, \dots, n+m-1\}$

$$\mathbb{P}(N_{n,m} = k) = \binom{k-1}{n-1} p^n (1-p)^{k-n} + \binom{k-1}{m-1} (1-p)^m p^{k-m} \quad (11.4.19)$$

Proof

Again, imagine that we continue the trials indefinitely. Let V_n denote the number of trials needed for A to win n points, and let W denote the number of trials needed for B to win m points. Then $\mathbb{P}(N_{n,m} = k) = \mathbb{P}(V_n = k) + \mathbb{P}(W_m = k)$ for k in the indicated range.

Series of Games

The special case of the problem of points experiment with $m = n$ is important, because it corresponds to A and B playing a *best of $2n-1$ game series*. That is, the first player to win n games wins the series. Such series, especially when $n \in \{2, 3, 4\}$, are frequently used in championship tournaments.

Let $A_n(p)$ denote the probability that player A wins the series. Then

$$A_n(p) = \sum_{k=n}^{2n-1} \binom{2n-1}{k} p^k (1-p)^{2n-1-k} = \sum_{j=n}^{2n-1} \binom{j-1}{n-1} p^n (1-p)^{j-n} \quad (11.4.20)$$

Proof

This follows directly from the [problem of points probability](#) above, since $A_n(p) = A_{n,n}(p)$.

Suppose that $p = 0.6$. Explicitly find the probability that team A wins in each of the following cases:

1. A best of 5 game series.
2. A best of 7 game series.

Answer

1. 0.6825.
2. 0.7102

In the problem of points experiments, vary the parameters n , m , and p (keeping $n = m$), and note how the probability changes. Now simulate a best of 5 series by selecting $n = m = 3$, $p = 0.6$. Run the experiment 1000 times and compare the relative frequency to the true probability.

$A_n(1 - p) = 1 - A_n(p)$ for any $n \in \mathbb{N}_+$ and $p \in [0, 1]$. Therefore

1. The graph of A_n is symmetric with respect to $p = \frac{1}{2}$.
2. $A_n\left(\frac{1}{2}\right) = \frac{1}{2}$.

Proof

Again, there is a simple probabilistic argument for the equation: both sides represent the probability that a player with game probability $1 - p$ will win the series.

In the problem of points experiments, vary the parameters n , m , and p (keeping $n = m$), and note how the probability changes. Now simulate a best 7 series by selecting $n = m = 4$, $p = 0.45$. Run the experiment 1000 times and compare the relative frequency to the true probability.

If $n > m$ then

1. $A_n(p) < A_m(p)$ if $0 < p < \frac{1}{2}$
2. $A_n(p) > A_m(p)$ if $\frac{1}{2} < p < 1$

Proof

The greater the number of games in the series, the more the series favors the stronger player (the one with the larger game probability).

Let N_n denote the number of trials in the series. Then N_n has probability density function

$$\mathbb{P}(N_n = k) = \binom{k-1}{n-1} [p^n(1-p)^{k-n} + (1-p)^n p^{k-n}], \quad k \in \{n, n+1, \dots, 2n-1\} \quad (11.4.21)$$

Proof

This result follows directly from the corresponding [problem of points result](#) above with $n = m$.

Explicitly compute the probability density function, expected value, and standard deviation for the number of games in a best of 7 series with the following values of p :

1. 0.5
2. 0.7
3. 0.9

Answer

1. $f(k) = \binom{k-1}{3} \left(\frac{1}{2}\right)^{k-1}$, $k \in \{4, 5, 6, 7\}$, $\mathbb{E}(N) = 5.8125$, $\text{sd}(N) = 1.0136$
2. $f(k) = \binom{k-1}{3} [(0.7)^4(0.3)^{k-4} + (0.3)^4(0.7)^{k-4}]$, $k \in \{4, 5, 6, 7\}$, $\mathbb{E}(N) = 5.3780$, $\text{sd}(N) = 1.0497$
3. $f(k) = \binom{k-1}{3} [(0.9)^4(0.1)^{k-4} + (0.1)^4(0.9)^{k-4}]$, $k \in \{4, 5, 6, 7\}$, $\mathbb{E}(N) = 4.4394$, $\text{sd}(N) = 0.6831$

Division of Stakes

The problem of points originated from a question posed by Chevalier de Mere, who was interested in the fair division of stakes when a game is interrupted. Specifically, suppose that players A and B each put up c monetary units, and then play Bernoulli trials until one of them wins a specified number of trials. The winner then takes the entire $2c$ fortune.

If the game is interrupted when A needs to win n more trials and B needs to win m more trials, then the fortune should be divided between A and B , respectively, as follows:

1. $2cA_{n,m}(p)$ for A
2. $2c[1 - A_{n,m}(p)] = 2cA_{m,n}(1-p)$ for B .

Suppose that players A and B bet \$50 each. The players toss a fair coin until one of them has 10 wins; the winner takes the entire fortune. Suppose that the game is interrupted by the gambling police when A has 5 wins and B has 3 wins. How should the stakes be divided?

Answer

A gets \$72.56, B gets \$27.44

Alternate and General Versions

Let's return to the formulation at the beginning of this section. Thus, suppose that we have a sequence of Bernoulli trials \mathbf{X} with success parameter $p \in (0, 1]$, and for $k \in \mathbb{N}_+$, we let V_k denote the trial number of the k th success. Thus, V_k has the negative binomial distribution with parameters k and p as we studied above. The random variable $W_k = V_k - k$ is the number of failures before the k th success. Let $N_1 = W_1$, the number of failures before the first success, and let $N_k = W_k - W_{k-1}$, the number of failures between the $(k-1)$ st success and the k th success, for $k \in \{2, 3, \dots\}$.

$\mathbf{N} = (N_1, N_2, \dots)$ is a sequence of independent random variables, each having the geometric distribution on \mathbb{N} with parameter p . Moreover,

$$W_k = \sum_{i=1}^k N_i \quad (11.4.22)$$

Thus, $\mathbf{W} = (W_1, W_2, \dots)$ is the partial sum process associated with \mathbf{N} . In particular, \mathbf{W} has stationary, independent increments.

Probability Density Functions

The probability density function of W_k is given by

$$\mathbb{P}(W_k = n) = \binom{n+k-1}{k-1} p^k (1-p)^n = \binom{n+k-1}{n} p^k (1-p)^n, \quad n \in \mathbb{N} \quad (11.4.23)$$

Proof

This result follows directly from the PDF of V_k , since $\mathbb{P}(W_k = n) = \mathbb{P}(V_k = k+n)$ for $n \in \mathbb{N}$.

The distribution of W_k is also referred to as the *negative binomial distribution* with parameters k and p . Thus, the term *negative binomial distribution* can refer either to the distribution of the trial number of the k th success or the distribution of the number of failures before the k th success, depending on the author and the context. The two random variables differ by a constant, so it's not a particularly important issue as long as we know which version is intended. In this text, we will refer to the alternate version as the negative binomial distribution *on* \mathbb{N} , to distinguish it from the original version, which has support set $\{k, k+1, \dots\}$

More interestingly, however, the probability density function in the last result makes sense for any $k \in (0, \infty)$, not just integers. To see this, first recall the definition of the general binomial coefficient: if $a \in \mathbb{R}$ and $n \in \mathbb{N}$, we define

$$\binom{a}{n} = \frac{a^{(n)}}{n!} = \frac{a(a-1) \cdots (a-n+1)}{n!} \quad (11.4.24)$$

The function f given below defines a probability density function for every $p \in (0, 1)$ and $k \in (0, \infty)$:

$$f(n) = \binom{n+k-1}{n} p^k (1-p)^n, \quad n \in \mathbb{N} \quad (11.4.25)$$

Proof

Recall from the section on Combinatorial Structures that $\binom{n+k-1}{n} = (-1)^n \binom{-k}{n}$. From the general binomial theorem,

$$\sum_{n=0}^{\infty} f(n) = p^k \sum_{n=0}^{\infty} \binom{-k}{n} (-1)^n (1-p)^n = p^k [1 - (1-p)]^{-k} = 1 \quad (11.4.26)$$

Once again, the distribution defined by the probability density function in the last theorem is the *negative binomial distribution on* \mathbb{N} , with parameters k and p . The special case when k is a positive integer is sometimes referred to as the *Pascal distribution*, in honor of Blaise Pascal.

The distribution is unimodal. Let $t = |k-1| \frac{1-p}{p}$.

1. $f(n-1) < f(n)$ if and only if $n < t$.
2. The distribution has a single mode at $\lfloor t \rfloor$ if t is not an integer.
3. The distribution has two consecutive modes at $t-1$ and t if t is a positive integer.

Basic Properties

Suppose that W has the negative binomial distribution on \mathbb{N} with parameters $k \in (0, \infty)$ and $p \in (0, 1)$. To establish basic properties, we can no longer use the decomposition of W as a sum of independent geometric variables. Instead, the best approach is to derive the probability generating function and then use the generating function to obtain other basic properties.

W has probability generating function P given by

$$P(t) = \mathbb{E}(t^W) = \left(\frac{p}{1 - (1-p)t} \right)^k, \quad |t| < \frac{1}{1-p} \quad (11.4.27)$$

Proof

This follows from the general binomial theorem: for $|t| < 1/(1-p)$,

$$\mathbb{E}(t^W) = \sum_{n=0}^{\infty} f(n)t^n = p^k \sum_{n=0}^{\infty} \binom{-k}{n} (-1)^n (1-p)^n t^n = p^k [1 - (1-p)t]^{-k} \quad (11.4.28)$$

The moments of W can be obtained from the derivatives of the probability generating function.

W has the following moments:

1. $\mathbb{E}(W) = k \frac{1-p}{p}$
2. $\text{var}(W) = k \frac{1-p}{p^2}$
3. $\text{skew}(W) = \frac{2-p}{\sqrt{k(1-p)}}$
4. $\text{kurt}(W) = \frac{3(k+2)(1-p)+p^2}{k(1-p)}$

Proof

Recall that the factorial moments of W can be obtained from the derivatives of the probability generating function: $\mathbb{E}[W^{(k)}] = P^{(k)}(1)$. Then the various moments above can be obtained from standard formulas.

The negative binomial distribution on \mathbb{N} is preserved under sums of independent variables.

Suppose that V has the negative binomial distribution on \mathbb{N} with parameters $a \in (0, \infty)$ and $p \in (0, 1)$, and that W has the negative binomial distribution on \mathbb{N} with parameters $b \in (0, \infty)$ and $p \in (0, 1)$, and that V and W are independent. Then $V + W$ has the negative binomial on \mathbb{N} distribution with parameters $a + b$ and p .

Proof

This result follows from the probability generating functions. Recall that the PGF of $V + W$ is the product of the PGFs of V and W .

In the last result, note that the success parameter p must be the same for both variables.

Normal Approximation

Because of the decomposition of W when the parameter k is a positive integer, it's not surprising that a central limit theorem holds for the general negative binomial distribution.

Suppose that W has the negative binomial distribution with parameters $k \in (0, \infty)$ and $p \in (0, 1)$. The standard score of W is

$$Z = \frac{pW - k(1-p)}{\sqrt{k(1-p)}} \quad (11.4.29)$$

The distribution of Z converges to the standard normal distribution as $k \rightarrow \infty$.

Thus, if k is large (and not necessarily an integer), then the distribution of W is approximately normal with mean $k \frac{1-p}{p}$ and variance $k \frac{1-p}{p^2}$.

Special Families

The negative binomial distribution on \mathbb{N} belongs to several special families of distributions. First, It follows from the result above on [sums](#) that we can decompose a negative binomial variable on \mathbb{N} into the sum of an arbitrary number of independent, identically distributed variables. This special property is known as infinite divisibility, and is studied in more detail in the chapter on Special Distributions.

The negative binomial distribution on \mathbb{N} is infinitely divisible.

Proof

Suppose that V has the negative binomial distribution on \mathbb{N} with parameters $k \in (0, \infty)$ and $p \in (0, 1)$. It follows from the previous result that for any $n \in \mathbb{N}_+$, V can be represented as $V = \sum_{i=1}^n V_i$ where (V_1, V_2, \dots, V_n) are independent, and each has the negative binomial distribution on \mathbb{N} with parameters k/n and p .

A Poisson-distributed random sum of independent, identically distributed random variables is said to have a compound Poisson distributions; these distributions are studied in more detail in the chapter on the Poisson Process. A theorem of William Feller states that an infinite divisible distribution

on \mathbb{N} must be compound Poisson. Hence it follows from the previous result that the negative binomial distribution on \mathbb{N} belongs to this family. Here is the explicit result:

Let $p, k \in (0, \infty)$. Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent variables, each having the logarithmic series distribution with shape parameter $1 - p$. Suppose also that N is independent of \mathbf{X} and has the Poisson distribution with parameter $-k \ln(p)$. Then $W = \sum_{i=1}^N X_i$ has the negative binomial distribution on \mathbb{N} with parameters k and p .

Proof

From the general theory of compound Poisson distributions, the probability generating function of W is $P(t) = \exp(\lambda[Q(t) - 1])$ where λ is the parameter of the Poisson variable N and $Q(t)$ is the common PGF of the terms in the sum. Using the PGF of the logarithmic series distribution, and the particular values of the parameters, we have

$$P(t) = \exp \left[-k \ln(p) \left(\frac{\ln[1 - (1-p)t]}{\ln(p)} - 1 \right) \right], \quad |t| < \frac{1}{1-p} \quad (11.4.30)$$

Using properties of logarithms and simple algebra, this reduces to

$$P(t) = \left(\frac{p}{1 - (1-p)t} \right)^k, \quad |t| < \frac{1}{1-p} \quad (11.4.31)$$

which is the PGF of the negative binomial distribution with parameters k and p .

As a special case ($k = 1$), it follows that the geometric distribution on \mathbb{N} is infinitely divisible and compound Poisson.

Next, the negative binomial distribution on \mathbb{N} belongs to the general exponential family. This family is important in inferential statistics and is studied in more detail in the chapter on Special Distributions.

Suppose that W has the negative binomial distribution on \mathbb{N} with parameters $k \in (0, \infty)$ and $p \in (0, 1)$. For fixed k , W has a one-parameter exponential distribution with natural statistic W and natural parameter $\ln(1 - p)$.

Proof

The PDF of W can be written as

$$f(n) = \binom{n+k-1}{n} p^k \exp[n \ln(1-p)], \quad n \in \mathbb{N} \quad (11.4.32)$$

so the result follows from the definition of the general exponential family.

Finally, the negative binomial distribution on \mathbb{N} is a power series distribution. Many special discrete distribution belong to this family, which is studied in more detail in the chapter on Special Distributions.

For fixed $k \in (0, \infty)$, the negative binomial distribution on \mathbb{N} with parameters k and $p \in (0, 1)$ is a power series distribution corresponding to the function $g(\theta) = 1/(1 - \theta)^k$ for $\theta \in (0, 1)$, where $\theta = 1 - p$.

Proof

In terms of the new parameter θ , the negative binomial pdf has the form $f(n) = \frac{1}{g(\theta)} \binom{n+k-1}{n} \theta^n$ for $n \in \mathbb{N}$, and $\sum_{n=0}^{\infty} \binom{n+k-1}{n} \theta^n = g(\theta)$.

Computational Exercises

Suppose that W has the negative binomial distribution with parameters $k = \frac{15}{2}$ and $p = \frac{3}{4}$. Compute each of the following:

1. $\mathbb{P}(W = 3)$
2. $\mathbb{E}(W)$
3. $\text{var}(W)$

Answer

1. $\mathbb{P}(W = 3) = 0.1823$
2. $\mathbb{E}(W) = \frac{5}{2}$
3. $\text{var}(W) = \frac{10}{3}$

Suppose that W has the negative binomial distribution with parameters $k = \frac{1}{3}$ and $p = \frac{1}{4}$. Compute each of the following:

1. $\mathbb{P}(W \leq 2)$
2. $\mathbb{E}(W)$
3. $\text{var}(W)$

Answer

1. $\mathbb{P}(W \leq 2) = \frac{11}{8\sqrt[3]{4}}$
2. $\mathbb{E}(W) = 1$
3. $\text{var}(W) = 4$

Suppose that W has the negative binomial distribution with parameters $k = 10\pi$ and $p = \frac{1}{3}$. Compute each of the following:

1. $\mathbb{E}(W)$
2. $\text{var}(W)$
3. The normal approximation to $\mathbb{P}(50 \leq W \leq 70)$

Answer

1. $\mathbb{E}(W) = 20\pi$
2. $\text{var}(W) = 60\pi$
3. $\mathbb{P}(50 \leq W \leq 70) \approx 0.5461$

This page titled [11.4: The Negative Binomial Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.5: The Multinomial Distribution

Basic Theory

Multinomial trials

A *multinomial trials process* is a sequence of independent, identically distributed random variables $\mathbf{X} = (X_1, X_2, \dots)$ each taking k possible values. Thus, the multinomial trials process is a simple generalization of the Bernoulli trials process (which corresponds to $k = 2$). For simplicity, we will denote the set of outcomes by $\{1, 2, \dots, k\}$, and we will denote the common probability density function of the trial variables by

$$p_i = \mathbb{P}(X_j = i), \quad i \in \{1, 2, \dots, k\} \quad (11.5.1)$$

Of course $p_i > 0$ for each i and $\sum_{i=1}^k p_i = 1$. In statistical terms, the sequence \mathbf{X} is formed by sampling from the distribution.

As with our discussion of the binomial distribution, we are interested in the random variables that count the number of times each outcome occurred. Thus, let

$$Y_i = \# \{j \in \{1, 2, \dots, n\} : X_j = i\} = \sum_{j=1}^n \mathbf{1}(X_j = i), \quad i \in \{1, 2, \dots, k\} \quad (11.5.2)$$

Of course, these random variables also depend on the parameter n (the number of trials), but this parameter is fixed in our discussion so we suppress it to keep the notation simple. Note that $\sum_{i=1}^k Y_i = n$ so if we know the values of $k-1$ of the counting variables, we can find the value of the remaining variable.

Basic arguments using independence and combinatorics can be used to derive the joint, marginal, and conditional densities of the counting variables. In particular, recall the definition of the multinomial coefficient: for nonnegative integers (j_1, j_2, \dots, j_k) with $\sum_{i=1}^k j_i = n$,

$$\binom{n}{j_1, j_2, \dots, j_k} = \frac{n!}{j_1! j_2! \cdots j_k!} \quad (11.5.3)$$

Joint Distribution

For nonnegative integers (j_1, j_2, \dots, j_k) with $\sum_{i=1}^k j_i = n$,

$$\mathbb{P}(Y_1 = j_1, Y_2 = j_2, \dots, Y_k = j_k) = \binom{n}{j_1, j_2, \dots, j_k} p_1^{j_1} p_2^{j_2} \cdots p_k^{j_k} \quad (11.5.4)$$

Proof

By independence, any sequence of trials in which outcome i occurs exactly j_i times for $i \in \{1, 2, \dots, k\}$ has probability $p_1^{j_1} p_2^{j_2} \cdots p_k^{j_k}$. The number of such sequences is the multinomial coefficient $\binom{n}{j_1, j_2, \dots, j_k}$. Thus, the result follows from the additive property of probability.

The distribution of $\mathbf{Y} = (Y_1, Y_2, \dots, Y_k)$ is called the *multinomial distribution* with parameters n and $\mathbf{p} = (p_1, p_2, \dots, p_k)$. We also say that $(Y_1, Y_2, \dots, Y_{k-1})$ has this distribution (recall that the values of $k-1$ of the counting variables determine the value of the remaining variable). Usually, it is clear from context which meaning of the term *multinomial distribution* is intended. Again, the ordinary binomial distribution corresponds to $k = 2$.

Marginal Distributions

For each $i \in \{1, 2, \dots, k\}$, Y_i has the binomial distribution with parameters n and p_i :

$$\mathbb{P}(Y_i = j) = \binom{n}{j} p_i^j (1 - p_i)^{n-j}, \quad j \in \{0, 1, \dots, n\} \quad (11.5.5)$$

Proof

There is a simple probabilistic proof. If we think of each trial as resulting in outcome i or not, then clearly we have a sequence of n Bernoulli trials with success parameter p_i . Random variable Y_i is the number of successes in the n trials. The result could also be obtained by summing the joint probability density function in Exercise 1 over all of the other variables, but this would be much harder.

Grouping

The multinomial distribution is preserved when the counting variables are combined. Specifically, suppose that (A_1, A_2, \dots, A_m) is a partition of the index set $\{1, 2, \dots, k\}$ into nonempty subsets. For $j \in \{1, 2, \dots, m\}$ let

$$Z_j = \sum_{i \in A_j} Y_i, \quad q_j = \sum_{i \in A_j} p_i \quad (11.5.6)$$

$Z = (Z_1, Z_2, \dots, Z_m)$ has the multinomial distribution with parameters n and $\mathbf{q} = (q_1, q_2, \dots, q_m)$.

Proof

Again, there is a simple probabilistic proof. Each trial, independently of the others, results in an outcome in A_j with probability q_j . For each j , Z_j counts the number of trials which result in an outcome in A_j . This result could also be derived from the joint probability density function in Exercise 1, but again, this would be a much harder proof.

Conditional Distribution

The multinomial distribution is also preserved when some of the counting variables are observed. Specifically, suppose that (A, B) is a partition of the index set $\{1, 2, \dots, k\}$ into nonempty subsets. Suppose that $(j_i : i \in B)$ is a sequence of nonnegative integers, indexed by B such that $j = \sum_{i \in B} j_i \leq n$. Let $p = \sum_{i \in A} p_i$.

The conditional distribution of $(Y_i : i \in A)$ given $(Y_i = j_i : i \in B)$ is multinomial with parameters $n - j$ and $(p_i/p : i \in A)$.

Proof

Again, there is a simple probabilistic argument and a harder analytic argument. If we know $Y_i = j_i$ for $i \in B$, then there are $n - j$ trials remaining, each of which, independently of the others, must result in an outcome in A . The conditional probability of a trial resulting in $i \in A$ is p_i/p .

Combinations of the basic results involving **grouping** and **conditioning** can be used to compute any marginal or conditional distributions.

Moments

We will compute the mean and variance of each counting variable, and the covariance and correlation of each pair of variables.

For $i \in \{1, 2, \dots, k\}$, the mean and variance of Y_i are

1. $\mathbb{E}(Y_i) = np_i$
2. $\text{var}(Y_i) = np_i(1 - p_i)$

Proof

Recall that Y_i has the binomial distribution with parameters n and p_i .

For distinct $i, j \in \{1, 2, \dots, k\}$,

1. $\text{cov}(Y_i, Y_j) = -np_i p_j$
2. $\text{cor}(Y_i, Y_j) = -\sqrt{p_i p_j / [(1 - p_i)(1 - p_j)]}$

Proof

From the bi-linearity of the covariance operator, we have

$$\text{cov}(Y_i, Y_j) = \sum_{s=1}^n \sum_{t=1}^n \text{cov}[\mathbf{1}(X_s = i), \mathbf{1}(X_t = j)] \quad (11.5.7)$$

If $s = t$, the covariance of the indicator variables is $-p_i p_j$. If $s \neq t$ the covariance is 0 by independence. Part (b) can be obtained from part (a) using the definition of correlation and the variances of Y_i and Y_j given [above](#).

From the [last result](#), note that the number of times outcome i occurs and the number of times outcome j occurs are negatively correlated, but the correlation does not depend on n .

If $k = 2$, then the number of times outcome 1 occurs and the number of times outcome 2 occurs are perfectly correlated.

Proof

This follows immediately from the result above on [covariance](#) since we must have $i = 1$ and $j = 2$, and $p_2 = 1 - p_1$. Of course we can also argue this directly since $Y_2 = n - Y_1$.

Examples and Applications

In the dice experiment, select the number of aces. For each die distribution, start with a single die and add dice one at a time, noting the shape of the probability density function and the size and location of the mean/standard deviation bar. When you get to 10 dice, run the simulation 1000 times and compare the relative frequency function to the probability density function, and the empirical moments to the distribution moments.

Suppose that we throw 10 standard, fair dice. Find the probability of each of the following events:

1. Scores 1 and 6 occur once each and the other scores occur twice each.
2. Scores 2 and 4 occur 3 times each.
3. There are 4 even scores and 6 odd scores.
4. Scores 1 and 3 occur twice each given that score 2 occurs once and score 5 three times.

Answer

1. 0.00375
2. 0.0178
3. 0.205
4. 0.0879

Suppose that we roll 4 ace-six flat dice (faces 1 and 6 have probability $\frac{1}{4}$ each; faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each). Find the joint probability density function of the number of times each score occurs.

Answer

$$f(u, v, w, x, y, z) = \binom{4}{u, v, w, x, y, z} \left(\frac{1}{4}\right)^{u+z} \left(\frac{1}{8}\right)^{v+w+x+y} \text{ for nonnegative integers } u, v, w, x, y, z \text{ that sum to 4}$$

In the dice experiment, select 4 ace-six flats. Run the experiment 500 times and compute the joint relative frequency function of the number times each score occurs. Compare the relative frequency function to the true probability density function.

Suppose that we roll 20 ace-six flat dice. Find the covariance and correlation of the number of 1's and the number of 2's.

Answer

covariance: -0.625 correlation: -0.0386

In the dice experiment, select 20 ace-six flat dice. Run the experiment 500 times, updating after each run. Compute the empirical covariance and correlation of the number of 1's and the number of 2's. Compare the results with the theoretical results computed [previously](#).

This page titled [11.5: The Multinomial Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.6: The Simple Random Walk

The simple random walk process is a minor modification of the Bernoulli trials process. Nonetheless, the process has a number of very interesting properties, and so deserves a section of its own. In some respects, it's a discrete time analogue of the Brownian motion process.

The Basic Process

Suppose that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent random variables, each taking values 1 and -1 with probabilities $p \in [0, 1]$ and $1 - p$ respectively. Let $\mathbf{X} = (X_0, X_1, X_2, \dots)$ be the partial sum process associated with \mathbf{U} , so that

$$X_n = \sum_{i=1}^n U_i, \quad n \in \mathbb{N} \quad (11.6.1)$$

The sequence \mathbf{X} is the *simple random walk* with parameter p .

We imagine a person or a particle on an axis, so that at each discrete time step, the walker moves either one unit to the right (with probability p) or one unit to the left (with probability $1 - p$), independently from step to step. The walker could accomplish this by tossing a coin with probability of heads p at each step, to determine whether to move right or move left. Other types of random walks, and additional properties of this random walk, are studied in the chapter on Markov Chains.

The mean and standard deviation, respectively of a step U are

1. $\mathbb{E}(U) = 2p - 1$
2. $\text{var}(U) = 4p(1 - p)$

Let $I_j = \frac{1}{2}(U_j + 1)$ for $j \in \mathbb{N}_+$. Then $\mathbf{I} = (I_1, I_2, \dots)$ is a Bernoulli trials sequence with success parameter p .

Proof

Note that $I_j = 1$ if $U_j = 1$ and $I_j = 0$ if $U_j = -1$.

In terms of the random walker, I_j is the indicator variable for the event that the j th step is to the right.

Let $R_n = \sum_{i=1}^n I_i$ for $n \in \mathbb{N}$, so that $\mathbf{R} = (R_0, R_1, \dots)$ is the partial sum process associated with \mathbf{I} . Then

1. $X_n = 2R_n - n$ for $n \in \mathbb{N}$.
2. R_n has the binomial distribution with trial parameter n and success parameter p .

In terms of the walker, R_n is the number of steps to the right in the first n steps.

X_n has probability density function

$$\mathbb{P}(X_n = k) = \binom{n}{(n+k)/2} p^{(n+k)/2} (1-p)^{(n-k)/2}, \quad k \in \{-n, -n+2, \dots, n-2, n\} \quad (11.6.2)$$

Proof

Since R_n takes values in $\{0, 1, \dots, n\}$, X_n takes values in $\{-n, -n+2, \dots, n-2, n\}$. For k in this set, $\mathbb{P}(X_n = k) = \mathbb{P}[R_n = (n+k)/2]$, so the result follows from the binomial distribution of R_n .

The mean and variance of X_n are

1. $\mathbb{E}(X_n) = n(2p - 1)$
2. $\text{var}(X_n) = 4np(1 - p)$

The Simple Symmetric Random Walk

Suppose now that $p = \frac{1}{2}$. In this case, $\mathbf{X} = (X_0, X_1, \dots)$ is called the *simple symmetric random walk*. The symmetric random walk can be analyzed using some special and clever combinatorial arguments. But first we give the basic results above for this special case.

For each $n \in \mathbb{N}_+$, the random vector $\mathbf{U}_n = (U_1, U_2, \dots, U_n)$ is uniformly distributed on $\{-1, 1\}^n$, and therefore

$$\mathbb{P}(\mathbf{U}_n \in A) = \frac{\#(A)}{2^n}, \quad A \subseteq S \quad (11.6.3)$$

X_n has probability density function

$$\mathbb{P}(X_n = k) = \binom{n}{(n+k)/2} \frac{1}{2^n}, \quad k \in \{-n, -n+2, \dots, n-2, n\} \quad (11.6.4)$$

The mean and variance of X_n are

1. $\mathbb{E}(X_n) = 0$
2. $\text{var}(X_n) = n$

In the random walk simulation, select the final position. Vary the number of steps and note the shape and location of the probability density function and the mean \pm standard deviation bar. For selected values of the parameter, run the simulation 1000 times and compare the empirical density function and moments to the true probability density function and moments.

In the random walk simulation, select the final position and set the number of steps to 50. Run the simulation 1000 times and compute and compare the following:

1. $\mathbb{P}(-6 \leq X_{50} \leq 10)$
2. The relative frequency of the event $\{-6 \leq X_{50} \leq 10\}$
3. The normal approximation to $\mathbb{P}(-6 \leq X_{50} \leq 10)$

Answer

1. 0.7794
3. 0.7752

The Maximum Position

Consider again the simple, symmetric random walk. Let $Y_n = \max\{X_0, X_1, \dots, X_n\}$, the *maximum position* during the first n steps. Note that Y_n takes values in the set $\{0, 1, \dots, n\}$. The distribution of Y_n can be derived from a simple and wonderful idea known as the *reflection principle*.

For $n \in \mathbb{N}$ and $y \in \{0, 1, \dots, n\}$

$$\mathbb{P}(Y_n = y) = \begin{cases} \mathbb{P}(X_n = y) = \binom{n}{(y+n)/2} \frac{1}{2^n}; & y, n \text{ have the same parity (both even or both odd)} \\ \mathbb{P}(X_n = y+1) = \binom{n}{(y+n+1)/2} \frac{1}{2^n}; & y, n \text{ have opposite parity (one even and one odd)} \end{cases} \quad (11.6.5)$$

Proof

Note first that $Y_n \geq y$ if and only if $X_i = y$ for some $i \in \{0, 1, \dots, n\}$. Suppose that $k \leq y \leq n$. For each path that satisfies $Y_n \geq y$ and $X_n = k$ there is another path that satisfies $X_n = 2y - k$. The second path is obtained from the first path by reflecting in the line $x = y$, after the first path hits y . Since the paths are equally likely,

$$\mathbb{P}(Y_n \geq y, X_n = k) = \mathbb{P}(X_n = 2y - k), \quad k \leq y \leq n \quad (11.6.6)$$

Hence it follows that

$$\mathbb{P}(Y_n = y, X_n = k) = \mathbb{P}(X_n = 2y - k) - \mathbb{P}[X_n = 2(y+1) - k], \quad k \leq y \leq n \quad (11.6.7)$$

In the random walk simulation, select the maximum value variable. Vary the number of steps and note the shape and location of the probability density function and the mean/standard deviation bar. Now set the number of steps to 30 and run the simulation 1000 times. Compare the relative frequency function and empirical moments to the true probability density function and moments.

For every n , the probability density function of Y_n is decreasing.

The last result is a bit surprising; in particular, the single most likely value for the maximum (and hence the mode of the distribution) is 0.

Explicitly compute the probability density function, mean, and standard deviation of Y_5 .

Answer

1. Probability density function of Y_5 : $f(0) = f(1) = \frac{10}{32}$, $f(2) = f(3) = \frac{5}{32}$, $f(4) = f(5) = \frac{1}{32}$
2. $\mathbb{E}(Y_5) = \frac{11}{8}$
3. $\text{var}(Y_5) = \frac{111}{64}$

A fair coin is tossed 10 times. Find the probability that the difference between the number of heads and the number of tails is never greater than 4.

Answer

$$\mathbb{P}(Y_{10} \leq 4) = \frac{57}{64}$$

The Last Visit to 0

Consider again the simple, symmetric random walk. Our next topic is the last visit to 0 during the first $2n$ steps:

$$Z_{2n} = \max \{j \in \{0, 2, \dots, 2n\} : X_j = 0\}, \quad n \in \mathbb{N} \quad (11.6.8)$$

Note that since visits to 0 can only occur at even times, Z_{2n} takes the values in the set $\{0, 2, \dots, 2n\}$. This random variable has a strange and interesting distribution known as the *discrete arcsine distribution*. Along the way to our derivation, we will discover some other interesting results as well.

The probability density function of Z_{2n} is

$$\mathbb{P}(Z_{2n} = 2k) = \binom{2k}{k} \binom{2n-2k}{n-k} \frac{1}{2^{2n}}, \quad k \in \{0, 1, \dots, n\} \quad (11.6.9)$$

Proof

Note that

$$\mathbb{P}(Z_{2n} = 2k) = \mathbb{P}(X_{2k} = 0, X_{2k+1} \neq 0, \dots, X_{2n} \neq 0), \quad k \in \{0, 1, \dots, n\} \quad (11.6.10)$$

From independence and symmetry it follows that

$$\mathbb{P}(Z_{2n} = 2k) = \mathbb{P}(X_{2k} = 0) \mathbb{P}(X_1 \neq 0, X_2 \neq 0, \dots, X_{2n-2k} \neq 0), \quad k \in \{0, 1, \dots, n\} \quad (11.6.11)$$

We know the first factor on the right from the distribution of X_{2k} . Thus, we need to compute the second factor, the probability that our random walk never returns to 0 during a time interval. Using results for the [maximum position](#) we have

$$\mathbb{P}(X_1 \leq 0, X_2 \leq 0, \dots, X_{2j} \leq 0) = \mathbb{P}(Y_{2j} = 0) \binom{2j}{j} \frac{1}{2^{2j}} \quad (11.6.12)$$

From symmetry (which is just the reflection principle at $y = 0$), it follows that

$$\mathbb{P}(X_1 \geq 0, X_2 \geq 0, \dots, X_{2n} \geq 0) = \binom{2n}{n} \frac{1}{2^{2n}} \quad (11.6.13)$$

Next, $\{X_1 > 0, X_2 > 0, \dots, X_{2j} > 0\} = \{X_1 = 1, X_2 \geq 1, \dots, X_{2j} \geq 1\}$. From independence and symmetry,

$$\mathbb{P}(X_1 > 0, X_2 > 0, \dots, X_{2j} > 0) = \mathbb{P}(X_1 = 0) \mathbb{P}(X_1 \geq 0, X_2 \geq 0, \dots, X_{2j-1} \geq 0) \quad (11.6.14)$$

But $X_{2j-1} \geq 0$ implies $X_{2j} \geq 0$. Hence

$$\mathbb{P}(X_1 > 0, X_2 > 0, \dots, X_{2j} > 0) = \binom{2j}{j} \frac{1}{2^{2j+1}} \quad (11.6.15)$$

From symmetry,

$$\mathbb{P}(X_1 \neq 0, X_2 \neq 0, \dots, X_{2j} \neq 0) = \binom{2j}{j} \frac{1}{2^{2j}} \quad (11.6.16)$$

In the random walk simulation, choose the last visit to 0 and then vary the number of steps with the scroll bar. Note the shape and location of the probability density function and the mean/standard deviation bar. For various values of the parameter, run the simulation 1000 times and compare the empirical density function and moments to the true probability density function and moments.

The probability density function of Z_{2n} is symmetric about n and is u -shaped:

1. $\mathbb{P}(Z_{2n} = 2k) = \mathbb{P}(Z_{2n} = 2n - 2k)$
2. $\mathbb{P}(Z_{2n} = 2j) > \mathbb{P}(Z_{2n} = 2k)$ if and only if $j < k$ and $2k \leq n$

In particular, 0 and $2n$ are the most likely values and hence are the modes of the distribution. The discrete arcsine distribution is quite surprising. Since we are tossing a fair coin to determine the steps of the walker, you might easily think that the random walk should be positive half of the time and negative half of the time, and that it should return to 0 frequently. But in fact, the arcsine law implies that with probability $\frac{1}{2}$, there will be *no* return to 0 during the second half of the walk, from time $n + 1$ to $2n$, regardless of n , and it is not uncommon for the walk to stay positive (or negative) during the entire time from 1 to $2n$.

Explicitly compute the probability density function, mean, and variance of Z_{10} .

Answer

1. Probability density function of Z_{10} : $f(0) = f(10) = \frac{63}{256}$, $f(2) = f(8) = \frac{35}{256}$, $f(4) = f(6) = \frac{30}{256}$
2. $\mathbb{E}(Z_{10}) = 5$
3. $\text{var}(Z_{10}) = 15$

The Ballot Problem and the First Return to Zero

The Ballot Problem

Suppose that in an election, candidate A receives a votes and candidate B receives b votes where $a > b$. Assuming a random ordering of the votes, what is the probability that A is always ahead of B in the vote count? This is an historically famous problem known as the *Ballot Problem*, that was solved by Joseph Louis Bertrand in 1887. The ballot problem is intimately related to simple random walks.

Comment on the validity of the assumption that the voters are randomly ordered for a real election.

The ballot problem can be solved by using a simple conditional probability argument to obtain a recurrence relation. Let $f(a, b)$ denote the probability that A is always ahead of B in the vote count.

f satisfies the initial condition $f(1, 0) = 1$ and the following recurrence relation:

$$f(a, b) = \frac{a}{a+b} f(a-1, b) + \frac{b}{a+b} f(a, b-1) \quad (11.6.17)$$

Proof

This follows by conditioning on the candidate that receives the last vote.

The probability that A is always ahead in the vote count is

$$f(a, b) = \frac{a-b}{a+b} \quad (11.6.18)$$

Proof

This follows from the recurrence relation and induction on the total number of votes $n = a + b$

In the ballot experiment, vary the parameters a and b and note the change the ballot probability. For selected values of the parameters, run the experiment 1000 times and compare the relative frequency to the true probability.

In an election for mayor of a small town, Mr. Smith received 4352 votes while Ms. Jones received 7543 votes. Compute the probability that Jones was always ahead of Smith in the vote count.

Answer

$$\frac{3191}{11895} \approx 0.2683$$

Relation to Random Walks

Consider again the [simple random walk](#) X with parameter p .

Given $X_n = k$,

1. There are $\frac{n+k}{2}$ steps to the right and $\frac{n-k}{2}$ steps to the left.
2. All possible orderings of the steps to the right and the steps to the left are equally likely.

For $k > 0$,

$$\mathbb{P}(X_1 > 0, X_2 > 0, \dots, X_{n-1} > 0 \mid X_n = k) = \frac{k}{n} \quad (11.6.19)$$

Proof

This follows from the previous result and the [ballot probability](#).

In the ballot experiment, vary the parameters a and b and note the change the ballot probability. For selected values of the parameters, run the experiment 1000 times and compare the relative frequency to the true probability.

An American roulette wheel has 38 slots; 18 are red, 18 are black, and 2 are green. Fred bet \$1 on red, at even stakes, 50 times, winning 22 times and losing 28 times. Find the probability that Fred's net fortune was always negative.

Answer

$$\frac{3}{25}$$

Roulette is studied in more detail in the chapter on Games of Chance.

The Distribution of the First Zero

Consider again the simple random walk with parameter p , as in the last subsection. Let T denote the time of the first return to 0:

$$T = \min\{n \in \mathbb{N}_+ : X_n = 0\} \quad (11.6.20)$$

Note that returns to 0 can only occur at even times; it may also be possible that the random walk never returns to 0. Thus, T takes values in the set $\{2, 4, \dots\} \cup \{\infty\}$.

The probability density function of T_{2n} is given by

$$\mathbb{P}(T = 2n) = \binom{2n}{n} \frac{1}{2n-1} p^n (1-p)^n, \quad n \in \mathbb{N}_+ \quad (11.6.21)$$

Proof

For $n \in \mathbb{N}_+$

$$\mathbb{P}(T = 2n) = \mathbb{P}(T = 2n, X_{2n} = 0) = \mathbb{P}(T = 2n \mid X_{2n} = 0)\mathbb{P}(X_{2n} = 0) \quad (11.6.22)$$

From the ballot problem,

$$\mathbb{P}(T = 2n \mid X_{2n} = 0) = \frac{1}{2n-1} \quad (11.6.23)$$

Fred and Wilma are tossing a fair coin; Fred gets a point for each head and Wilma gets a point for each tail. Find the probability that their scores are equal for the first time after n tosses, for each $n \in \{2, 4, 6, 8, 10\}$.

Answer

$$f(2) = \frac{1}{2}, f(4) = \frac{1}{8}, f(6) = \frac{1}{16}, f(8) = \frac{5}{128}, f(10) = \frac{7}{512}$$

This page titled [11.6: The Simple Random Walk](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

11.7: The Beta-Bernoulli Process

An interesting thing to do in almost any parametric probability model is to “randomize” one or more of the parameters. Done in a clever way, this often leads to interesting new models and unexpected connections between models. In this section we will randomize the success parameter in the Bernoulli trials process. This leads to interesting and surprising connections with Pólya's urn process.

Basic Theory

Definitions

First, recall that the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$ is a continuous distribution on the interval $(0, 1)$ with probability density function g given by

$$g(p) = \frac{1}{B(a, b)} p^{a-1} (1-p)^{b-1}, \quad p \in (0, 1) \quad (11.7.1)$$

where B is the *beta function*. So $B(a, b)$ is simply the normalizing constant for the function $p \mapsto p^{a-1} (1-p)^{b-1}$ on the interval $(0, 1)$. Here is our main definition:

Suppose that P has the beta distribution with left parameter $a \in (0, \infty)$ and right parameter $b \in (0, \infty)$. Next suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of indicator random variables with the property that given $P = p \in (0, 1)$, \mathbf{X} is a conditionally independent sequence with

$$\mathbb{P}(X_i = 1 \mid P = p) = p, \quad i \in \mathbb{N}_+ \quad (11.7.2)$$

Then \mathbf{X} is the *beta-Bernoulli process* with parameters a and b .

In short, given $P = p$, the sequence \mathbf{X} is a Bernoulli trials sequence with success parameter p . In the usual language of reliability, X_i is the outcome of trial i , where 1 denotes success and 0 denotes failure. For a specific application, suppose that we select a random probability of heads according to the beta distribution with parameters a and b , and then toss a coin with this probability of heads repeatedly.

Outcome Variables

What's our first step? Well, of course we need to compute the finite dimensional distributions of \mathbf{X} . Recall that for $r \in \mathbb{R}$ and $j \in \mathbb{N}$, $r^{[j]}$ denotes the ascending power $r(r+1) \cdots [r+(j-1)]$. By convention, a product over an empty index set is 1, so $r^{[0]} = 1$.

Suppose that $n \in \mathbb{N}_+$ and $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$. Let $k = x_1 + x_2 + \cdots + x_n$. Then

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a^{[k]} b^{[n-k]}}{(a+b)^{[n]}} \quad (11.7.3)$$

Proof

First, note that $\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid P = p) = p^k (1-p)^{n-k}$ by the conditional independence. Thus, conditioning on P gives

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \mathbb{E}[\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid P)] \quad (11.7.4)$$

$$= \int_0^1 p^k (1-p)^{n-k} \frac{1}{B(a, b)} p^{a-1} (1-p)^{b-1} dp \quad (11.7.5)$$

$$= \frac{B[a+k, b+(n-k)]}{B(a, b)} = \frac{a^{[k]} b^{[n-k]}}{(a+b)^{[n]}} \quad (11.7.6)$$

The last step uses a property of the beta function.

From this result, it follows that Pólya's urn process with parameters $a, b, c \in \mathbb{N}_+$ is equivalent to the beta-Bernoulli process with parameters a/c and b/c , quite an interesting result. Note that since the joint distribution above depends only on

$x_1 + x_2 + \dots + x_n$, the sequence \mathbf{X} is exchangeable. Finally, it's interesting to note that the beta-Bernoulli process with parameters a and b could simply be defined as the sequence with the finite-dimensional distributions above, without reference to the beta distribution! It turns out that *every* exchangeable sequence of indicator random variables can be obtained by randomizing the success parameter in a sequence of Bernoulli trials. This is *de Finetti's theorem*, named for Bruno de Finetti, which is studied in the section on backwards martingales.

For each $i \in \mathbb{N}_+$

1. $\mathbb{E}(X_i) = \frac{a}{a+b}$
2. $\text{var}(X_i) = \frac{a}{a+b} \frac{b}{a+b}$

Proof

Since the sequence is exchangeable, X_i has the same distribution as X_1 , so $\mathbb{P}(X_i = 1) = \frac{a}{a+b}$. The mean and variance now follow from standard results for indicator variables.

Thus \mathbf{X} is a sequence of identically distributed variables, quite surprising at first but of course inevitable for any exchangeable sequence. Compare the [joint distribution](#) with the [marginal distributions](#). Clearly the variables are dependent, so let's compute the covariance and correlation of a pair of outcome variables.

Suppose that $i, j \in \mathbb{N}_+$ are distinct. Then

1. $\text{cov}(X_i, X_j) = \frac{ab}{(a+b)^2(a+b+1)}$
2. $\text{cor}(X_i, X_j) = \frac{1}{a+b+1}$

Proof

Since the variables are exchangeable, $\mathbb{P}(X_i = 1, X_j = 1) = \mathbb{P}(X_1 = 1, X_2 = 1) = \frac{a}{a+b} \frac{a+1}{a+b+1}$. The results now follow from standard formulas for covariance and correlation.

Thus, the variables are positively correlated. It turns out that in any *infinite* sequence of exchangeable variables, the the variables must be nonnegatively correlated. Here is another result that explores how the variables are related.

Suppose that $n \in \mathbb{N}_+$ and $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$. Let $k = \sum_{i=1}^n x_i$. Then

$$\mathbb{P}(X_{n+1} = 1 \mid X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a+k}{a+b+n} \quad (11.7.7)$$

Proof

Using the [joint distribution](#),

$$\begin{aligned} \mathbb{P}(X_{n+1} = 1 \mid X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) &= \frac{\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n, X_{n+1} = 1)}{\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)} \\ &= \frac{a^{[k+1]} b^{[n-k]} (a+b)^{[n]}}{(a+b)^{[n+1]} a^{[k]} b^{[n-k]}} = \frac{a+k}{a+b+n} \end{aligned}$$

The beta-Bernoulli model starts with the conditional distribution of \mathbf{X} given P . Let's find the conditional distribution in the other direction.

Suppose that $n \in \mathbb{N}_+$ and $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$. Let $k = \sum_{i=1}^n x_i$. Then the conditional distribution of P given $(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)$ is beta with left parameter $a+k$ and right parameter $b+(n-k)$. Hence

$$\mathbb{E}(P \mid X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a+k}{a+b+k} \quad (11.7.8)$$

Proof

This follows from Bayes' theorem. The conditional PDF $g(\cdot \mid x_1, x_2, \dots, x_n)$ is given by

$$g(p | X_1, x_2, \dots, x_n) = \frac{g(p)\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n | P = p)}{\int_0^1 g(t)\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n | P = t)dt}, \quad p \in (0, 1) \quad (11.7.9)$$

The numerator is

$$\frac{1}{B(a, b)} p^{a-1} (1-p)^{b-1} p^k (1-p)^{n-k} = \frac{1}{B(a, b)} p^{a+k-1} (1-p)^{b+n-k-1} \quad (11.7.10)$$

The denominator is simply the normalizing constant for the expression, as a function of p and is $B(a+k, b+n-k)/B(a, b)$. Hence

$$g(p | k) = \frac{1}{B(a+k, b+n-k)} p^{a+k-1} (1-p)^{b+n-k-1}, \quad p \in (0, 1) \quad (11.7.11)$$

The last result follows since the mean of the beta distribution is the left parameter divided by the sum of the parameters.

Thus, the left parameter increases by the number of successes while the right parameter increases by the number of failures. In the language of Bayesian statistics, the original distribution of P is the *prior distribution*, and the conditional distribution of P given the data (x_1, x_2, \dots, x_n) is the *posterior distribution*. The fact that the posterior distribution is beta whenever the prior distribution is beta means that the beta distributions is *conjugate* to the Bernoulli distribution. The conditional expected value in the last theorem is the *Bayesian estimate* of p when p is modeled by the random variable P . These concepts are studied in more generality in the section on Bayes Estimators in the chapter on Point Estimation. It's also interesting to note that the expected values in the last two theorems are the same: If $n \in \mathbb{N}$, $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ and $k = \sum_{i=1}^n x_i$ then

$$\mathbb{E}(X_{n+1} | X_1 = x_1, \dots, X_n = x_n) = \mathbb{E}(P | X_1 = x_1, \dots, X_n = x_n) = \frac{a+k}{a+b+n} \quad (11.7.12)$$

Run the simulation of the beta coin experiment for various values of the parameter. Note how the posterior probability density function changes from the prior probability density function, given the number of heads.

The Number of Successes

It's already clear that the number of successes in a given number of trials plays an important role, so let's study these variables. For $n \in \mathbb{N}_+$, let

$$Y_n = \sum_{i=1}^n X_i \quad (11.7.13)$$

denote the number of successes in the first n trials. Of course, $\mathbf{Y} = (Y_0, Y_1, \dots)$ is the partial sum process associated with $\mathbf{X} = (X_1, X_2, \dots)$.

Y_n has probability density function given by

$$\mathbb{P}(Y_n = k) = \binom{n}{k} \frac{a^{[k]} b^{[n-k]}}{(a+b)^{[n]}}, \quad k \in \{0, 1, \dots, n\} \quad (11.7.14)$$

Proof

Every bit string of length n with 1 occurring exactly k times has the probability given in the [joint distribution](#) above. There are $\binom{n}{k}$ such bit strings.

The distribution of Y_n is known as the *beta-binomial distribution* with parameters n , a , and b .

In the simulation of the beta-binomial experiment, vary the parameters and note how the shape of the probability density function of Y_n (discrete) parallels the shape of the probability density function of P (continuous). For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The case where the parameters are both 1 is interesting.

If $a = b = 1$, so that P is uniformly distributed on $(0, 1)$, then Y_n is uniformly distributed on $\{0, 1, \dots, n\}$.

Proof

Note that $1^{[j]} = j!$ and $2^{[j]} = (j+1)!$ for $j \in \mathbb{N}$. Hence, from the [general PDF](#) Y_n above

$$\mathbb{P}(Y_n = k) = \frac{n!}{k!(n-k)!} \frac{k!(n-k)!}{(n+1)!} = \frac{1}{n+1}, \quad k \in \{0, 1, \dots, n\} \quad (11.7.15)$$

Next, let's compute the mean and variance of Y_n .

The mean and variance of Y_n are

1. $\mathbb{E}(Y_n) = n \frac{a}{a+b}$
2. $\text{var}(Y_n) = n \frac{ab}{(a+b)^2} \left[1 + (n-1) \frac{1}{a+b+1} \right]$

Proof

These results follow from the [mean](#) and [covariance](#) results given above:

$$\mathbb{E}(Y_n) = \sum_{i=a}^n \mathbb{E}(X_i) = n \frac{a}{a+b} \quad (11.7.16)$$

$$\text{var}(Y_n) = \sum_{i=1}^n \sum_{j=1}^n \text{cov}(X_i, X_j) = n \frac{ab}{(a+b)^2} + n(n-1) \frac{ab}{(a+b)^2(a+b+1)} \quad (11.7.17)$$

In the simulation of the beta-binomial experiment, vary the parameters and note the location and size of the mean-standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical moments to the true moments.

We can restate the conditional distributions in the last subsection more elegantly in terms of Y_n .

Let $n \in \mathbb{N}$.

1. The conditional distribution of X_{n+1} given Y_n is

$$\mathbb{P}(X_{n+1} = 1 \mid Y_n) = \mathbb{E}(X_{n+1} \mid Y_n) = \frac{a + Y_n}{a + b + n} \quad (11.7.18)$$

2. The conditional distribution of P given Y_n is beta with left parameter $a + Y_n$ and right parameter $b + (n - Y_n)$. In particular

$$\mathbb{E}(P \mid Y_n) = \frac{a + Y_n}{a + b + n} \quad (11.7.19)$$

Proof

The proof is easy using the nesting property of conditional expected value and the fact that the conditional distributions given (X_1, X_2, \dots, X_n) depend only on $Y_n = \sum_{i=1}^n X_i$.

1. Note that

$$\mathbb{E}(X_{n+1} \mid Y_n) = \mathbb{E}[\mathbb{E}(X_{n+1} \mid Y_n) \mid X_1, X_2, \dots, X_n] \quad (11.7.20)$$

$$= \mathbb{E}\left[\frac{a + Y_n}{a + b + n} \mid X_1, X_2, \dots, X_n\right] = \mathbb{E}\left(\frac{a + Y_n}{a + b + n} \mid Y_n\right) = \frac{a + Y_n}{a + b + n} \quad (11.7.21)$$

2. Similarly, if $A \subseteq (0, 1)$ is measurable then $\mathbb{P}(P \in A \mid X_1, X_2, \dots, X_n)$ depends only on Y_n and so

$$\mathbb{P}(P \in A \mid Y_n) = \mathbb{E}[\mathbb{P}(P \in A \mid Y_n) \mid X_1, X_2, \dots, X_n] \quad (11.7.22)$$

$$= \mathbb{E}[\mathbb{P}(P \in A \mid X_1, X_2, \dots, X_n) \mid Y_n] = \mathbb{P}(P \in A \mid Y_n) \quad (11.7.23)$$

Once again, the conditional expected value $\mathbb{E}(P \mid Y_n)$ is the *Bayesian estimator* of p . In particular, if $a = b = 1$, so that P has the uniform distribution on $(0, 1)$, then $\mathbb{P}(X_{n+1} = 1 \mid Y_n = n) = \frac{n+1}{n+2}$. This is *Laplace's rule of succession*, another interesting

connection. The rule is named for Pierre Simon Laplace, and is studied from a different point of view in the section on Independence.

The Proportion of Successes

For $n \in \mathbb{N}_+$, let

$$M_n = \frac{Y_n}{n} = \frac{1}{n} \sum_{i=1}^n X_i \quad (11.7.24)$$

so that M_n is the *sample mean* of (X_1, X_2, \dots, X_n) , or equivalently the proportion of successes in the first n trials. Properties of M_n follow easily from the corresponding properties of Y_n . In particular, $\mathbb{P}(M_n = k/n) = \mathbb{P}(Y_n = k)$ for $k \in \{0, 1, \dots, n\}$ as given [above](#), so let's move on to the mean and variance.

For $n \in \mathbb{N}_+$, the mean and variance of M_n are

1. $\mathbb{E}(M_n) = \frac{a}{a+b}$
2. $\text{var}(M_n) = \frac{1}{n} \frac{ab}{(a+b)^2} + \frac{n-1}{n} \frac{ab}{(a+b)^2(a+b+1)}$

Proof

These results follow from the mean and variance of Y_n [above](#) and properties of expected value and variance:

1. $\mathbb{E}(M_n) = \frac{1}{n} \mathbb{E}(Y_n)$
2. $\text{var}(M_n) = \frac{1}{n^2} \text{var}(Y_n)$

So $\mathbb{E}(M_n)$ is constant in $n \in \mathbb{N}_+$ while $\text{var}(M_n) \rightarrow ab/(a+b)^2(a+b+1)$ as $n \rightarrow \infty$. These results suggest that perhaps M_n has a limit, in some sense, as $n \rightarrow \infty$. For an ordinary sequence of Bernoulli trials with success parameter $p \in (0, 1)$, we know from the law of large numbers that $M_n \rightarrow p$ as $n \rightarrow \infty$ with probability 1 and in mean (and hence also in distribution). What happens here when the success probability P has been randomized with the beta distribution? The answer is what we might hope.

$M_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1 and in mean square, and hence also in in distribution.

Proof

Let g denote the PDF of P . For convergence with probability 1, we condition on P

$$\mathbb{P}(M_n \rightarrow P \text{ as } n \rightarrow \infty) = \mathbb{E}[\mathbb{P}(M_n \rightarrow P \text{ as } n \rightarrow \infty) \mid P] \quad (11.7.25)$$

$$= \int_0^1 \mathbb{P}(M_n \rightarrow p \text{ as } n \rightarrow \infty \mid P = p) g(p) dp = \int_0^1 g(p) dp = 1 \quad (11.7.26)$$

For convergence in mean square, once again we condition on P . Note that

$$\mathbb{E}[(M_n - P)^2 \mid P = p] = \mathbb{E}[(M_n - p)^2 \mid P = p] = \frac{p(1-p)}{n} \rightarrow 0 \text{ as } n \rightarrow \infty \quad (11.7.27)$$

Hence by the dominated convergence theorem,

$$\mathbb{E}[(M_n - P)^2] = \int_0^1 \frac{p(1-p)}{n} g(p) dp \rightarrow 0 \text{ as } n \rightarrow \infty \quad (11.7.28)$$

Proof of convergence in distribution

Convergence with probability 1 implies convergence in distribution, but it's interesting to give a direct proof. For $x \in (0, 1)$, note that

$$\mathbb{P}(M_n \leq x) = \mathbb{P}(Y_n \leq nx) = \sum_{k=0}^{\lfloor nx \rfloor} \binom{n}{k} \frac{a^{[k]} b^{[n-k]}}{(a+b)^{[n]}} \quad (11.7.29)$$

where $\lfloor \cdot \rfloor$ is the floor function. But recall that

$$\frac{a^{[k]}b^{[n-k]}}{(a+b)^{[n]}} = \frac{B(a+k, b+n-k)}{B(a, b)} = \frac{1}{B(a, b)} \int_0^1 p^{a+k-1} (1-p)^{b+n-k-1} dp \quad (11.7.30)$$

Substituting and doing some algebra we get

$$\mathbb{P}(M_n \leq x) = \frac{1}{B(a, b)} \int_0^1 \left[\sum_{k=0}^{\lfloor nx \rfloor} \binom{n}{k} p^k (1-p)^{n-k} \right] p^{a-1} (1-p)^{b-1} dp \quad (11.7.31)$$

The sum in the square brackets is $\mathbb{P}(W_n \leq nx) = \mathbb{P}(W_n/n \leq x)$ where W_n has the ordinary binomial distribution with parameters n and p . But W_n/n converges (in every sense) to p as $n \rightarrow \infty$ so $\mathbb{P}(W_n/n \leq x) \rightarrow \mathbf{1}(p \leq x)$ as $n \rightarrow \infty$. So by the dominated convergence theorem,

$$\mathbb{P}(M_n \leq x) \rightarrow \frac{1}{B(a, b)} \int_0^x p^{a-1} (1-p)^{b-1} dp = \mathbb{P}(P \leq x) \quad (11.7.32)$$

Recall again that the Bayesian estimator of p based on (X_1, X_2, \dots, X_n) is

$$\mathbb{E}(P | Y_n) = \frac{a + Y_n}{a + b + n} = \frac{a/n + M_n}{a/n + b/n + 1} \quad (11.7.33)$$

It follows from the last theorem that $\mathbb{E}(P | Y_n) \rightarrow P$ with probability 1, in mean square, and in distribution. The stochastic process $\mathbf{Z} = \{Z_n = (a + Y_n)/(a + b + n) : n \in \mathbb{N}\}$ that we have seen several times now is of fundamental importance, and turns out to be a martingale. The theory of martingales provides powerful tools for studying convergence in the beta-Bernoulli process.

The Trial Number of a Success

For $k \in \mathbb{N}_+$, let V_k denote the trial number of the k th success. As we have seen before in similar circumstances, the process $\mathbf{V} = (V_1, V_2, \dots)$ can be defined in terms of the process \mathbf{Y} :

$$V_k = \min\{n \in \mathbb{N}_+ : Y_n = k\}, \quad k \in \mathbb{N}_+ \quad (11.7.34)$$

Note that V_k takes values in $\{k, k+1, \dots\}$. The random processes $\mathbf{V} = (V_1, V_2, \dots)$ and $\mathbf{Y} = (Y_0, Y_1, \dots)$ are inverses of each other in a sense.

For $k \in \mathbb{N}$ and $n \in \mathbb{N}_+$ with $k \leq n$,

1. $V_k \leq n$ if and only if $Y_n \geq k$
2. $V_k = n$ if and only if $Y_{n-1} = k-1$ and $X_n = 1$

The probability density function of V_k is given by

$$\mathbb{P}(V_k = n) = \binom{n-1}{k-1} \frac{a^{[k]}b^{[n-k]}}{(a+b)^{[n]}}, \quad n \in \{k, k+1, \dots\} \quad (11.7.35)$$

Proof 1

As usual, we can condition on P and use known results for ordinary Bernoulli trials. Given $P = p$, random variable V_k has the negative binomial distribution with parameters k and p . Hence

$$\begin{aligned} \mathbb{P}(V_k = n) &= \int_0^1 \mathbb{P}(V_k = n | P = p) g(p) dp = \int_0^1 \binom{n-1}{k-1} p^k (1-p)^{n-k} \frac{1}{B(a, b)} p^{a-1} (1-p)^{b-1} dp \\ &= \binom{n-1}{k-1} \frac{1}{B(a, b)} \int_0^1 p^{a+k-1} (1-p)^{b+n-k-1} dp \\ &= \binom{n-1}{k-1} \frac{B(a+k, b+n-k)}{B(a, b)} = \binom{n-1}{k-1} \frac{a^{[k]}b^{[n-k]}}{(a+b)^{[n]}} \end{aligned}$$

Proof 2

In this proof, we condition on Y_{n-1} . Using the PDF of Y_{n-1} and the result [above](#),

$$\begin{aligned}\mathbb{P}(V_k = n) &= \mathbb{P}(Y_{n-1} = k-1, X_n = 1) = \mathbb{P}(Y_{n-1} = k-1)\mathbb{P}(X_n = 1 \mid Y_{n-1} = k-1) \\ &= \binom{n-1}{k-1} \frac{a^{[k-1]}b^{[(n-1)-(k-1)]}}{(a+b)^{[n-1]}} \frac{a+k-1}{a+b+(n-1)} = \binom{n-1}{k-1} \frac{a^{[k]}b^{[n-k]}}{(a+b)^{[n]}}\end{aligned}$$

The distribution of V_k is known as the *beta-negative binomial distribution* with parameters k , a , and b .

If $a = b = 1$ so that P is uniformly distributed on $(0, 1)$, then

$$\mathbb{P}(V_k = n) = \frac{k}{n(n+1)}, \quad n \in \{k, k+1, k+2, \dots\} \quad (11.7.36)$$

Proof

Recall again that $1^{[j]} = j!$ and $2^{[j]} = (j+1)!$ for $j \in \mathbb{N}$. Hence from the previous result,

$$\mathbb{P}(V_k = n) = \frac{(n-1)!}{(k-1)!(n-k)!} \frac{k!(n-k)!}{(n+1)!} = \frac{k}{n(n+1)}, \quad n \in \{k, k+1, \dots\} \quad (11.7.37)$$

In the simulation of the beta-negative binomial experiment, vary the parameters and note the shape of the probability density function. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The mean and variance of V_k are

1. $\mathbb{E}(V_k) = k \frac{a+b-1}{a-1}$ if $a > 1$.
2. $\text{var}(V_k) = k \frac{a+b-1}{(a-1)(a-2)} [b + k(a+b-2)] - k^2 \left(\frac{a+b-1}{a-1} \right)^2$

Proof

From our work with the negative binomial distribution we know that $\mathbb{E}(V_k | P = p) = k \frac{1}{p}$ and $\mathbb{E}(V_k^2 | P = p) = k \frac{1-p}{p^2} + \frac{k^2}{p^2}$.

Thus, conditioning on P we have

$$\mathbb{E}(V_k) = \mathbb{E}[\mathbb{E}(V_k | P)] = \int_0^1 \frac{k}{p} \frac{p^{a-1}(1-p)^{b-1}}{B(a, b)} = k \frac{B(a-1, b)}{B(a, b)} = k \frac{a+b-1}{a-1} \quad (11.7.38)$$

which gives part (a). Similarly

$$\mathbb{E}(V_k^2) = \mathbb{E}[\mathbb{E}(V_k^2 | P)] = \int_0^1 \left(k \frac{1-p}{p^2} + \frac{k^2}{p^2} \right) \frac{p^{a-1}(1-p)^{b-1}}{B(a, b)} \quad (11.7.39)$$

$$= k \frac{B(a-2, b+1)}{B(a, b)} + k^2 \frac{B(a-2, b)}{B(a, b)} = k \frac{b(a+b-2)}{(a-1)(a-2)} + k^2 \frac{(a+b-1)(a+b-2)}{(a-1)(a-2)} \quad (11.7.40)$$

Simplifying and using part (a) gives part (b).

In the simulation of the beta-negative binomial experiment, vary the parameters and note the location and size of the mean \pm standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical moments to the true moments.

This page titled [11.7: The Beta-Bernoulli Process](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

12: Finite Sampling Models

This chapter explores a number of models and problems based on sampling from a finite population. Sampling without replacement from a population of objects of various types leads to the hypergeometric and multivariate hypergeometric models. Sampling with replacement from a finite population leads naturally to the birthday and coupon-collector problems. Sampling without replacement from an ordered population leads naturally to the matching problem and to the study of order statistics.

- [12.1: Introduction to Finite Sampling Models](#)
- [12.2: The Hypergeometric Distribution](#)
- [12.3: The Multivariate Hypergeometric Distribution](#)
- [12.4: Order Statistics](#)
- [12.5: The Matching Problem](#)
- [12.6: The Birthday Problem](#)
- [12.7: The Coupon Collector Problem](#)
- [12.8: Pólya's Urn Process](#)
- [12.9: The Secretary Problem](#)

This page titled [12: Finite Sampling Models](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.1: Introduction to Finite Sampling Models

Basic Theory

Sampling Models

Suppose that we have a *population* D of m objects. The population could be a deck of cards, a set of people, an urn full of balls, or any number of other collections. In many cases, we simply label the objects from 1 to m , so that $D = \{1, 2, \dots, m\}$. In other cases (such as the card experiment), it may be more natural to label the objects with vectors. In any case, D is usually a finite subset of \mathbb{R}^k for some $k \in \mathbb{N}_+$.

Our basic experiment consists of selecting n objects from the population D at random and recording the sequence of objects chosen. Thus, the outcome is $\mathbf{X} = (X_1, X_2, \dots, X_n)$ where $X_i \in D$ is the i th object chosen. If the sampling is *with replacement*, the sample size n can be any positive integer. In this case, the sample space S is

$$S = D^n = \{(x_1, x_2, \dots, x_n) : x_i \in D \text{ for each } i\} \quad (12.1.1)$$

If the sampling is *without replacement*, the sample size n can be no larger than the population size m . In this case, the sample space S consists of all permutations of size n chosen from D :

$$S = D_n = \{(x_1, x_2, \dots, x_n) : x_i \in D \text{ for each } i \text{ and } x_i \neq x_j \text{ for all } i \neq j\} \quad (12.1.2)$$

From the multiplication principle of combinatorics,

1. $\#(D^n) = m^n$
2. $\#(D_n) = m^{(n)} = m(m-1) \cdots (m-n+1)$

With either type of sampling, we assume that the samples are equally likely and thus that the outcome variable \mathbf{X} is uniformly distributed on the appropriate sample space S ; this is the meaning of the phrase *random sample*:

$$\mathbb{P}(\mathbf{X} \in A) = \frac{\#(A)}{\#(S)}, \quad A \subseteq S \quad (12.1.3)$$

The Exchangeable Property

Suppose again that we select n objects at random from the population D , either with or without replacement and record the ordered sample $\mathbf{X} = (X_1, X_2, \dots, X_n)$

Any permutation of \mathbf{X} has the same distribution as \mathbf{X} itself, namely the uniform distribution on the appropriate sample space S :

1. D^n if the sampling is with replacement.
2. D_n if the sampling is without replacement.

A sequence of random variables with this property is said to be *exchangeable*. Although this property is very simple to understand, both intuitively and mathematically, it is nonetheless very important. We will use the exchangeable property often in this chapter.

More generally, any sequence of k of the n outcome variables is uniformly distributed on the appropriate sample space:

1. D^k if the sampling is with replacement.
2. D_k if the sampling is without replacement.

In particular, for either sampling method, X_i is uniformly distributed on D for each $i \in \{1, 2, \dots, n\}$.

If the sampling is with replacement then $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent random variables.

Thus, when the sampling is with replacement, the sample variables form a random sample from the uniform distribution, in statistical terminology.

If the sampling is without replacement, then the conditional distribution of a sequence of k of the outcome variables, given the values of a sequence of j other outcome variables, is the uniform distribution on the set of permutations of size k chosen from the population when the j known values are removed (of course, $j + k \leq n$).

In particular, X_i and X_j are dependent for any distinct i and j when the sampling is without replacement.

The Unordered Sample

In many cases when the sampling is without replacement, the *order* in which the objects are chosen is not important; all that matters is the (unordered) *set* of objects:

$$\mathbf{W} = \{X_1, X_2, \dots, X_n\} \quad (12.1.4)$$

The random set \mathbf{W} takes values in the set of combinations of size n chosen from D :

$$T = \{\{x_1, x_2, \dots, x_n\} : x_i \in D \text{ for each } i \text{ and } x_i \neq x_j \text{ for all } i \neq j\} \quad (12.1.5)$$

Recall that $\#(T) = \binom{m}{n}$.

\mathbf{W} is uniformly distributed over T :

$$\mathbb{P}(\mathbf{W} \in B) = \frac{\#(B)}{\#(T)} = \frac{\#(B)}{\binom{m}{n}}, \quad B \subseteq T \quad (12.1.6)$$

Proof

For any combination of size n from D , there are $n!$ permutations of size n .

Suppose now that the sampling is with replacement, and we again denote the unordered outcome by \mathbf{W} . In this case, \mathbf{W} takes values in the collection of *multisets* of size n from D . (A multiset is like an ordinary set, except that repeated elements are allowed).

$$T = \{\{x_1, x_2, \dots, x_n\} : x_i \in D \text{ for each } i\} \quad (12.1.7)$$

Recall that $\#(T) = \binom{m+n-1}{n}$.

\mathbf{W} is *not* uniformly distributed on T .

Summary of Sampling Formulas

The following table summarizes the formulas for the number of samples of size n chosen from a population of m elements, based on the criteria of order and replacement.

Sampling Formulas

Number of samples	With order	Without
With replacement	m^n	$\binom{m+n-1}{n}$
Without	$m^{(n)}$	$\binom{m}{n}$

Examples and Applications

Suppose that a sample of size 2 is chosen from the population $\{1, 2, 3, 4\}$. Explicitly list all samples in the following cases:

1. Ordered samples, with replacement.
2. Ordered samples, without replacement.
3. Unordered samples, with replacement.
4. Unordered samples, without replacement.

Answer

1. $\{(1, 1), (1, 2), (1, 3), (1, 4), (2, 1), (2, 2), (2, 3), (2, 4), (3, 1), (3, 2), (3, 3), (3, 4), (4, 1), (4, 2), (4, 3), (4, 4)\}$

2. $\{(1, 2), (1, 3), (1, 4), (2, 1), (2, 3), (2, 4), (3, 1), (3, 2), (3, 4), (4, 1), (4, 2), (4, 3)\}$
3. $\{\{1, 1\}, \{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 2\}, \{2, 3\}, \{2, 4\}, \{3, 3\}, \{3, 4\}, \{4, 4\}\}$
4. $\{\{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}$

Multi-type Populations

A *dichotomous population* consists of two types of objects.

Suppose that a batch of 100 components includes 10 that are defective. A random sample of 5 components is selected without replacement. Compute the probability that the sample contains at least one defective component.

Answer

0.4162

An urn contains 50 balls, 30 red and 20 green. A sample of 15 balls is chosen at random. Find the probability that the sample contains 10 red balls in each of the following cases:

1. The sampling is without replacement
2. The sampling is with replacement

Answer

1. 0.2070
2. 0.1859

In the ball and urn experiment select 50 balls with 30 red balls, and sample size 15. Run the experiment 100 times. Compute the relative frequency of the event that the sample has 10 red balls in each of the following cases, and compare with the respective probability in the previous exercise:

1. The sampling is without replacement
2. The sampling is with replacement

Suppose that a club has 100 members, 40 men and 60 women. A committee of 10 members is selected at random (and without replacement, of course).

1. Find the probability that both genders are represented on the committee.
2. If you observed the experiment and in fact the committee members are all of the same gender, would you believe that the sampling was random?

Answer

1. 0.9956
2. No

Suppose that a small pond contains 500 fish, 50 of them tagged. A fisherman catches 10 fish. Find the probability that the catch contains at least 2 tagged fish.

Answer

0.2635

The basic distribution that arises from sampling without replacement from a dichotomous population is studied in the section on the hypergeometric distribution. More generally, a *multi-type population* consists of objects of k different types.

Suppose that a legislative body consists of 60 republicans, 40 democrats, and 20 independents. A committee of 10 members is chosen at random. Find the probability that at least one party is not represented on the committee.

Answer

0.1633. Use the inclusion-exclusion law.

The basic distribution that arises from sampling without replacement from a multi-type population is studied in the section on the multivariate hypergeometric distribution.

Cards

Recall that a standard *card deck* can be modeled by the product set

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (12.1.8)$$

where the first coordinate encodes the *denomination* or *kind* (ace, 2-10, jack, queen, king) and where the second coordinate encodes the *suit* (clubs, diamonds, hearts, spades). The general *card experiment* consists of drawing n cards at random and without replacement from the deck D . Thus, the i th card is $X_i = (Y_i, Z_i)$ where Y_i is the denomination and Z_i is the suit. The special case $n = 5$ is the *poker experiment* and the special case $n = 13$ is the *bridge experiment*. Note that with respect to the denominations or with respect to the suits, a deck of cards is a multi-type population as discussed above.

In the card experiment with $n = 5$ cards (poker), there are

1. 311,875,200 ordered hands
2. 2,598,960 unordered hands

In the card experiment with $n = 13$ cards (bridge), there are

1. 3,954,242,643,911,239,680,000 ordered hands
2. 635,013,559,600 unordered hands

In the card experiment, set $n = 5$. Run the simulation 5 times and on each run, list all of the (ordered) sequences of cards that would give the same unordered hand as the one you observed.

In the card experiment,

1. Y_i is uniformly distributed on $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\}$ for each i .
2. Z_j is uniformly distributed on $\{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\}$ for each i .

In the card experiment, Y_i and Z_j are independent for any i and j .

In the card experiment, (Y_1, Y_2) and (Z_1, Z_2) are dependent.

Suppose that a sequence of 5 cards is dealt. Find each of the following:

1. The probability that the third card is a spade.
2. The probability that the second and fourth cards are queens.
3. The conditional probability that the second card is a heart given that the fifth card is a heart.
4. The probability that the third card is a queen and the fourth card is a heart.

Answer

1. $\frac{1}{4}$
2. $\frac{1}{221}$
3. $\frac{4}{17}$
4. $\frac{1}{52}$

Run the card experiment 500 time. Compute the relative frequency corresponding to each probability in the previous exercise.

Find the probability that a bridge hand will contain no *honor cards* that is, no cards of denomination 10, jack, queen, king, or ace. Such a hand is called a *Yarborough*, in honor of the second Earl of Yarborough.

Answer

0.000547

Dice

Rolling n fair, six-sided *dice* is equivalent to choosing a random sample of size n with replacement from the population $\{1, 2, 3, 4, 5, 6\}$. Generally, selecting a random sample of size n with replacement from $D = \{1, 2, \dots, m\}$ is equivalent to rolling n fair, m -sided dice.

In the game of *poker dice*, 5 standard, fair dice are thrown. Find each of the following:

1. The probability that all dice show the same score.
2. The probability that the scores are distinct.
3. The probability that 1 occurs twice and 6 occurs 3 times.

Answer

1. $\frac{1}{1296}$
2. $\frac{5}{24}$
3. $\frac{5}{3888}$

Run the poker dice experiment 500 times. Compute the relative frequency of each event in the previous exercise and compare with the corresponding probability.

The game of poker dice is treated in more detail in the chapter on Games of Chance.

Birthdays

Supposes that we select n persons at random and record their *birthdays*. If we assume that birthdays are uniformly distributed throughout the year, and if we ignore leap years, then this experiment is equivalent to selecting a sample of size n with replacement from $D = \{1, 2, \dots, 365\}$. Similarly, we could record birth months or birth weeks.

Suppose that a probability class has 30 students. Find each of the following:

1. The probability that the birthdays are distinct.
2. The probability that there is at least one duplicate birthday.

Answer

1. 0.2937
2. 0.7063

In the birthday experiment, set $m = 365$ and $n = 30$. Run the experiment 1000 times and compare the relative frequency of each event in the previous exercise to the corresponding probability.

The birthday problem is treated in more detail later in this chapter.

Balls into Cells

Suppose that we distribute n distinct *balls* into m distinct *cells* at random. This experiment also fits the basic model, where D is the population of cells and X_i is the cell containing the i th ball. Sampling with replacement means that a cell may contain more than one ball; sampling without replacement means that a cell may contain at most one ball.

Suppose that 5 balls are distributed into 10 cells (with no restrictions). Find each of the following:

1. The probability that the balls are all in different cells.
2. The probability that the balls are all in the same cell.

Answer

1. $\frac{189}{625}$
2. $\frac{1}{10000}$

Coupons

Suppose that when we purchase a certain product (bubble gum, or cereal for example), we receive a *coupon* (a baseball card or small toy, for example), which is equally likely to be any one of m types. We can think of this experiment as sampling with replacement from the population of coupon types; X_i is the coupon that we receive on the i th purchase.

Suppose that a kid's meal at a fast food restaurant comes with a toy. The toy is equally likely to be any of 5 types. Suppose that a mom buys a kid's meal for each of her 3 kids. Find each of the following:

1. The probability that the toys are all the same.
2. The probability that the toys are all different.

Answer

1. $\frac{1}{25}$
2. $\frac{12}{25}$

The coupon collector problem is studied in more detail later in this chapter.

The Key Problem

Suppose that a person has n keys, only one of which opens a certain door. The person tries the keys at random. We will let N denote the trial number when the person finds the correct key.

Suppose that unsuccessful keys are discarded (the rational thing to do, of course). Then N has the uniform distribution on $\{1, 2, \dots, n\}$.

1. $\mathbb{P}(N = i) = \frac{1}{n}, \quad i \in \{1, 2, \dots, n\}.$
2. $\mathbb{E}(N) = \frac{n+1}{2}.$
3. $\text{var}(N) = \frac{n^2-1}{12}.$

Suppose that unsuccessful keys are not discarded (perhaps the person has had a bit too much to drink). Then N has a geometric distribution on \mathbb{N}_+ .

1. $\mathbb{P}(N = i) = \frac{1}{n} \left(\frac{n-1}{n} \right)^{i-1}, \quad i \in \mathbb{N}_+.$
2. $\mathbb{E}(N) = n.$
3. $\text{var}(N) = n(n-1).$

Simulating a Random Samples

It's very easy to simulate a random sample of size n , with replacement from $D = \{1, 2, \dots, m\}$. Recall that the *ceiling function* $\lceil x \rceil$ gives the smallest integer that is at least as large as x .

Let $\mathbf{U} = (U_1, U_2, \dots, U_n)$ be a sequence of be a random numbers. Recall that these are independent random variables, each uniformly distributed on the interval $[0, 1]$ (the standard uniform distribution). Then $X_i = \lceil m U_i \rceil$ for $i \in \{1, 2, \dots, n\}$ simulates a random sample, with replacement, from D .

It's a bit harder to simulate a random sample of size n , without replacement, since we need to remove each sample value before the next draw.

The following algorithm generates a random sample of size n , without replacement, from D .

1. For $i = 1$ to m , let $b_i = i$.
2. For $i = 1$ to n ,
 - a. let $j = m - i + 1$
 - b. let U_i be a random number
 - c. let $J = \lfloor j U_i \rfloor$
 - d. let $X_i = b_J$
 - e. let $k = b_j$

- f. let $b_j = b_J$
- g. let $b_J = k$
- 3. Return $\mathbf{X} = (X_1, X_2, \dots, X_n)$

This page titled [12.1: Introduction to Finite Sampling Models](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.2: The Hypergeometric Distribution

Basic Theory

Dichotomous Populations

Suppose that we have a *dichotomous* population D . That is, a population that consists of two types of objects, which we will refer to as type 1 and type 0. For example, we could have

- balls in an urn that are either *red* or *green*
- a batch of components that are either *good* or *defective*
- a population of people who are either *male* or *female*
- a population of animals that are either *tagged* or *untagged*
- voters who are either *democrats* or *republicans*

Let R denote the subset of D consisting of the type 1 objects, and suppose that $\#(D) = m$ and $\#(R) = r$. As in the basic sampling model, we sample n objects at random from D . In this section, our only concern is in the types of the objects, so let X_i denote the type of the i th object chosen (1 or 0). The random vector of types is

$$\mathbf{X} = (X_1, X_2, \dots, X_n) \quad (12.2.1)$$

Our main interest is the random variable Y that gives the number of type 1 objects in the sample. Note that Y is a counting variable, and thus like all counting variables, can be written as a sum of indicator variables, in this case the type variables:

$$Y = \sum_{i=1}^n X_i \quad (12.2.2)$$

We will assume initially that the sampling is without replacement, which is usually the realistic setting with dichotomous populations.

The Probability Density Function

Recall that since the sampling is without replacement, the unordered sample is uniformly distributed over the set of all combinations of size n chosen from D . This observation leads to a simple combinatorial derivation of the probability density function of Y .

The probability density function of Y is given by

$$\mathbb{P}(Y = y) = \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}}, \quad y \in \{\max\{0, n - (m - r)\}, \dots, \min\{n, r\}\} \quad (12.2.3)$$

Proof

Consider the unordered outcome, which is uniformly distributed on the set of combinations of size n chosen from the population of size m . The number of ways to select y type 1 objects from the r type 1 objects in the population is $\binom{r}{y}$. Similarly the number of ways to select the remaining $n - y$ type 0 objects from the $m - r$ type 0 objects in the population is $\binom{m-r}{n-y}$. Finally the number of ways to select the sample of size n from the population of size m is $\binom{m}{n}$.

This distribution defined by this probability density function is known as the *hypergeometric distribution* with parameters m , r , and n .

Another form of the probability density function of Y is

$$\mathbb{P}(Y = y) = \binom{n}{y} \frac{r^{(y)} (m-r)^{(n-y)}}{m^{(n)}}, \quad y \in \{\max\{0, n - (m - r)\}, \dots, \min\{n, r\}\} \quad (12.2.4)$$

Combinatorial Proof

The combinatorial proof is much like the [previous proof](#), except that we consider the ordered sample, which is uniformly distributed on the set of permutations of size n chosen from the population of m objects. The binomial coefficient $\binom{n}{y}$ is the number of ways to select the coordinates where the type 1 objects will go; $r^{(y)}$ is the number of ways to select an ordered sequence of y type 1 objects; and $(m-r)^{(n-y)}$ is the number of ways to select an ordered sequence of $n - y$ type 0 objects. Finally $m^{(n)}$ is the number of ways to select an ordered sequence of n objects from the population.

Algebraic Proof

The new form of the PDF can also be derived algebraically by starting with the [previous form of the PDF](#). Use the formula $\binom{k}{j} = k^{(j)} / j!$ for each binomial coefficient, and then rearrange things a bit.

Recall our convention that $j^{(i)} = \binom{j}{i} = 0$ for $i > j$. With this convention, the two formulas for the probability density function are correct for $y \in \{0, 1, \dots, n\}$. We usually use this simpler set as the set of values for the hypergeometric distribution.

The hypergeometric distribution is unimodal. Let $v = \frac{(r+1)(n+1)}{m+2}$. Then

1. $\mathbb{P}(Y = y) > \mathbb{P}(Y = y - 1)$ if and only if $y < v$.
2. The mode occurs at $\lfloor v \rfloor$ if v is not an integer, and at v and $v - 1$ if v is an integer greater than 0.

In the ball and urn experiment, select sampling without replacement. Vary the parameters and note the shape of the probability density function. For selected values of the parameters, run the experiment 1000 times and compare the relative frequency function to the probability density function.

You may wonder about the rather exotic name *hypergeometric distribution*, which seems to have nothing to do with sampling from a dichotomous population. The name comes from a power series, which was studied by Leonhard Euler, Carl Friedrich Gauss, Bernhard Riemann, and others.

A (generalized) *hypergeometric series* is a power series

$$\sum_{k=0}^{\infty} a_k x^k \quad (12.2.5)$$

where $k \mapsto a_{k+1}/a_k$ is a rational function (that is, a ratio of polynomials).

Many of the basic power series studied in calculus are hypergeometric series, including the ordinary geometric series and the exponential series.

The probability generating function of the hypergeometric distribution is a hypergeometric series.

Proof

The PGF is $P(t) = \sum_{k=0}^n f(k)t^k$ where f is the hypergeometric PDF, given [above](#). Simple algebra shows that

$$\frac{f(k+1)}{f(k)} = \frac{(r-k)(n-k)}{(k+1)(N-r-n+k+1)} \quad (12.2.6)$$

In addition, the hypergeometric distribution function can be expressed in terms of a hypergeometric *series*. These representations are not particularly helpful, so basically were stuck with the non-descriptive term for historical reasons.

Moments

Next we will derive the mean and variance of Y . The exchangeable property of the indicator variables, and properties of covariance and correlation will play a key role.

$\mathbb{E}(X_i) = \frac{r}{m}$ for each i .

Proof

Recall that X_i is an indicator variable with $\mathbb{P}(X_i = 1) = r/m$ for each i .

From the representation of Y as the sum of indicator variables, the expected value of Y is trivial to compute. But just for fun, we give the derivation from the probability density function as well.

$\mathbb{E}(Y) = n \frac{r}{m}$.

Proof

This follows from the [previous result](#) and the additive property of expected value.

Proof from the definition

Using the hypergeometric PDF,

$$\mathbb{E}(Y) = \sum_{y=0}^n y \frac{\binom{r}{y} \binom{m-r}{n-y}}{\binom{m}{n}} \quad (12.2.7)$$

Note that the $y = 0$ term is 0. For the other terms, we can use the identity $y \binom{r}{y} = r \binom{r-1}{y-1}$ to get

$$\mathbb{E}(Y) = \frac{r}{\binom{m}{n}} \sum_{y=1}^n \binom{r-1}{y-1} \binom{m-r}{n-y} \quad (12.2.8)$$

But substituting $k = y - 1$ and using another fundamental identity,

$$\sum_{y=1}^n \binom{r-1}{y-1} \binom{m-r}{n-y} = \sum_{k=0}^{n-1} \binom{r-1}{k} \binom{m-r}{n-1-k} = \binom{m-1}{n-1} \quad (12.2.9)$$

So substituting and doing a bit of algebra gives $\mathbb{E}(Y) = n \frac{r}{m}$.

Next we turn to the variance of the hypergeometric distribution. For that, we will need not only the variances of the indicator variables, but their covariances as well.

$\text{var}(X_i) = \frac{r}{m} \left(1 - \frac{r}{m}\right)$ for each i .

Proof

Again this follows because X_i is an indicator variable with $\mathbb{P}(X_i = 1) = r/m$ for each i .

For distinct i, j ,

1. $\text{cov}(X_i, X_j) = -\frac{r}{m} \left(1 - \frac{r}{m}\right) \frac{1}{m-1}$
2. $\text{cor}(X_i, X_j) = -\frac{1}{m-1}$

Proof

Note that $X_i X_j$ is an indicator variable that indicates the event that the i th and j th objects are both type 1. By the exchangeable property, $\mathbb{P}(X_i X_j = 1) = \mathbb{P}(X_i = 1) \mathbb{P}(X_j = 1 | X_i = 1) = \frac{r}{m} \frac{r-1}{m-1}$. Part (a) then follows from $\text{cov}(X_i, X_j) = \mathbb{E}(X_i X_j) - \mathbb{E}(X_i) \mathbb{E}(X_j)$. Part (b) follows from part (a) and the definition of correlation.

Note that the event of a type 1 object on draw i and the event of a type 1 object on draw j are negatively correlated, but the correlation depends only on the population size and not on the number of type 1 objects. Note also that the correlation is perfect if $m = 2$, which must be the case.

$\text{var}(Y) = n \frac{r}{m} \left(1 - \frac{r}{m}\right) \frac{m-n}{m-1}$.

Proof

This result follows from the previous results on the [variance](#) and [covariance](#) of the indicator variables. Recall that the variance of Y is the sum of $\text{cov}(X_i, X_j)$ over all i and j .

Note that $\text{var}(Y) = 0$ if $r = 0$ or $r = m$ or $n = m$, which must be true since Y is deterministic in each of these cases.

In the ball and urn experiment, select sampling without replacement. Vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

Sampling with Replacement

Suppose now that the sampling is *with* replacement, even though this is usually not realistic in applications.

(X_1, X_2, \dots, X_n) is a sequence of n Bernoulli trials with success parameter $\frac{r}{m}$.

The following results now follow immediately from the general theory of Bernoulli trials, although modifications of the arguments above could also be used.

Y has the binomial distribution with parameters n and $\frac{r}{m}$:

$$\mathbb{P}(Y = y) = \binom{n}{y} \left(\frac{r}{m}\right)^y \left(1 - \frac{r}{m}\right)^{n-y}, \quad y \in \{0, 1, \dots, n\} \quad (12.2.10)$$

The mean and variance of Y are

1. $\mathbb{E}(Y) = n \frac{r}{m}$
2. $\text{var}(Y) = n \frac{r}{m} \left(1 - \frac{r}{m}\right)$

Note that for any values of the parameters, the mean of Y is the same, whether the sampling is with or without replacement. On the other hand, the variance of Y is smaller, by a factor of $\frac{m-n}{m-1}$, when the sampling is without replacement than with replacement. It certainly makes sense that the variance of Y should be smaller when sampling without replacement, since each selection reduces the variability in the population that remains. The factor $\frac{m-n}{m-1}$ is sometimes called the *finite population correction factor*.

In the ball and urn experiment, vary the parameters and switch between sampling without replacement and sampling with replacement. Note the difference between the graphs of the hypergeometric probability density function and the binomial probability density function. Note also the difference between the mean \pm standard deviation bars. For selected values of the parameters and for the two different sampling modes, run the simulation 1000 times.

Convergence of the Hypergeometric Distribution to the Binomial

Suppose that the population size m is very large compared to the sample size n . In this case, it seems reasonable that sampling *without* replacement is not too much different than sampling *with* replacement, and hence the hypergeometric distribution should be well approximated by the binomial. The following exercise makes this observation precise. Practically, it is a valuable result, since the binomial distribution has fewer parameters. More specifically, we do not need to know the population size m and the number of type 1 objects r individually, but only in the ratio r/m .

Suppose that $r_m \in \{0, 1, \dots, m\}$ for each $m \in \mathbb{N}_+$ and that $r_m/m \rightarrow p \in [0, 1]$ as $m \rightarrow \infty$. Then for fixed n , the hypergeometric probability density function with parameters m , r_m , and n converges to the binomial probability density function with parameters n and p as $m \rightarrow \infty$.

Proof

Consider the [second version](#) of the hypergeometric PDF above. In the fraction, note that there are n factors in the numerator and n in the denominator. Suppose we pair the factors to write the original fraction as the product of n fractions. The first y fractions have the form $\frac{r_m - i}{m - i}$ where i does not depend on m . Hence each of these fractions converge to p as $m \rightarrow \infty$. The remaining $n - y$ fractions have the form $\frac{m - r_m - j}{m - y - j}$, where again, j does not depend on m . Hence each of these fractions converges to $1 - p$ as $m \rightarrow \infty$.

The type of convergence in the previous exercise is known as convergence in distribution.

In the ball and urn experiment, vary the parameters and switch between sampling without replacement and sampling with replacement. Note the difference between the graphs of the hypergeometric probability density function and the binomial probability density function. In particular, note the similarity when m is large and n small. For selected values of the parameters, and for both sampling modes, run the experiment 1000 times.

In the setting of the [convergence result](#) above, note that the mean and variance of the hypergeometric distribution converge to the mean and variance of the binomial distribution as $m \rightarrow \infty$.

Inferences in the Hypergeometric Model

In many real problems, the parameters r or m (or both) may be unknown. In this case we are interested in drawing inferences about the unknown parameters based on our observation of Y , the number of type 1 objects in the sample. We will assume initially that the sampling is without replacement, the realistic setting in most applications.

Estimation of r with m Known

Suppose that the size of the population m is known but that the number of type 1 objects r is unknown. This type of problem could arise, for example, if we had a batch of m manufactured items containing an unknown number r of defective items. It would be too costly to test all m items (perhaps even destructive), so we might instead select n items at random and test those.

A simple estimator of r can be derived by hoping that the *sample* proportion of type 1 objects is close to the *population* proportion of type 1 objects. That is,

$$\frac{Y}{n} \approx \frac{r}{m} \implies r \approx \frac{m}{n} Y \quad (12.2.11)$$

Thus, our estimator of r is $\frac{m}{n} Y$. This method of deriving an estimator is known as the method of moments.

$$\mathbb{E}\left(\frac{m}{n} Y\right) = r$$

Proof

This follows from the [expected value](#) of Y above, and the scale property of expected value.

The result in the previous exercise means that $\frac{m}{n} Y$ is an *unbiased* estimator of r . Hence the variance is a measure of the quality of the estimator, in the mean square sense.

$$\text{var}\left(\frac{m}{n} Y\right) = (m - r) \frac{r}{n} \frac{m - n}{m - 1}.$$

Proof

This follows from [variance](#) of Y above, and standard properties of variance.

For fixed m and r , $\text{var}\left(\frac{m}{n} Y\right) \downarrow 0$ as $n \uparrow m$.

Thus, the estimator improves as the sample size increases; this property is known as *consistency*.

In the ball and urn experiment, select sampling without replacement. For selected values of the parameters, run the experiment 100 times and note the estimate of r on each run.

1. Compute the average error and the average squared error over the 100 runs.
2. Compare the average squared error with the variance in [mean square error](#) given above.

Often we just want to estimate the *ratio* r/m (particularly if we don't know m either). In this case, the natural estimator is the sample proportion Y/n .

The estimator of $\frac{r}{m}$ has the following properties:

1. $\mathbb{E}\left(\frac{Y}{n}\right) = \frac{r}{m}$, so the estimator is unbiased.
2. $\text{var}\left(\frac{Y}{n}\right) = \frac{1}{n} \frac{r}{m} \left(1 - \frac{r}{m} \frac{m - n}{m - 1}\right)$
3. $\text{var}\left(\frac{Y}{n}\right) \downarrow 0$ as $n \uparrow m$ so the estimator is consistent.

Estimation of m with r Known

Suppose now that the number of type 1 objects r is known, but the population size m is unknown. As an example of this type of problem, suppose that we have a lake containing m fish where m is unknown. We capture r of the fish, tag them, and return them to the lake. Next we capture n of the fish and observe Y , the number of tagged fish in the sample. We wish to estimate m from this data. In this context, the estimation problem is sometimes called the *capture-recapture problem*.

Do you think that the main assumption of the sampling model, namely equally likely samples, would be satisfied for a real capture-recapture problem? Explain.

Once again, we can use the method of moments to derive a simple estimate of m , by hoping that the *sample* proportion of type 1 objects is close the *population* proportion of type 1 objects. That is,

$$\frac{Y}{n} \approx \frac{r}{m} \implies m \approx \frac{nr}{Y} \quad (12.2.12)$$

Thus, our estimator of m is $\frac{nr}{Y}$ if $Y > 0$ and is ∞ if $Y = 0$.

In the ball and urn experiment, select sampling without replacement. For selected values of the parameters, run the experiment 100 times.

1. On each run, compare the true value of m with the estimated value.
2. Compute the average error and the average squared error over the 100 runs.

If $y > 0$ then $\frac{nr}{y}$ maximizes $\mathbb{P}(Y = y)$ as a function of m for fixed r and n . This means that $\frac{nr}{Y}$ is a maximum likelihood estimator of m .

$$\mathbb{E}\left(\frac{nr}{Y}\right) \geq m.$$

Proof

This result follows from Jensen's inequality since $y \mapsto \frac{nr}{y}$ is a convex function on $(0, \infty)$.

Thus, the estimator is *positively biased* and tends to over-estimate m . Indeed, if $n \leq m - r$, so that $\mathbb{P}(Y = 0) > 0$ then $\mathbb{E}\left(\frac{nr}{Y}\right) = \infty$. For another approach to estimating the population size m , see the section on Order Statistics.

Sampling with Replacement

Suppose now that the sampling is with replacement, even though this is unrealistic in most applications. In this case, Y has the binomial distribution with parameters n and $\frac{r}{m}$. The estimators of r with m known, $\frac{r}{m}$, and m with r known make sense, just as before, but have slightly different properties.

The estimator $\frac{m}{n}Y$ of r with m known satisfies

1. $\mathbb{E}\left(\frac{m}{n}Y\right) = r$
2. $\text{var}\left(\frac{m}{n}Y\right) = \frac{r(m-r)}{n}$

The estimator $\frac{1}{n}Y$ of $\frac{r}{m}$ satisfies

1. $\mathbb{E}\left(\frac{1}{n}Y\right) = \frac{r}{m}$
2. $\text{var}\left(\frac{1}{n}Y\right) = \frac{1}{n} \frac{r}{m} \left(1 - \frac{r}{m}\right)$

Thus, the estimators are still unbiased and consistent, but have larger mean square error than before. Thus, sampling without replacement works better, for any values of the parameters, than sampling with replacement.

In the ball and urn experiment, select sampling with replacement. For selected values of the parameters, run the experiment 100 times.

1. On each run, compare the true value of r with the estimated value.
2. Compute the average error and the average squared error over the 100 runs.

Examples and Applications

A batch of 100 computer chips contains 10 defective chips. Five chips are chosen at random, without replacement. Find each of the following:

1. The probability density function of the number of defective chips in the sample.
2. The mean and variance of the number of defective chips in the sample
3. The probability that the sample contains at least one defective chip.

Answer

Let Y denote the number of defective chips in the sample

1. $\mathbb{P}(Y = y) = \frac{\binom{10}{y} \binom{90}{5-y}}{\binom{100}{5}}, \quad y \in \{0, 1, 2, 3, 4, 5\}$
2. $\mathbb{E}(Y) = 0.5, \text{var}(Y) = 0.432$
3. $\mathbb{P}(Y > 0) = 0.416$

A club contains 50 members; 20 are men and 30 are women. A committee of 10 members is chosen at random. Find each of the following:

1. The probability density function of the number of women on the committee.
2. The mean and variance of the number of women on the committee.
3. The mean and variance of the number of men on the committee.
4. The probability that the committee members are all the same gender.

Answer

Let Y denote the number of women, so that $Z = 10 - Y$ is the number of men.

1. $\mathbb{P}(Y = y) = \frac{\binom{30}{y} \binom{20}{10-y}}{\binom{50}{10}}, \quad y \in \{0, 1, \dots, 10\}$
2. $\mathbb{E}(Y) = 6, \text{var}(Y) = 1.959$
3. $\mathbb{E}(Z) = 4, \text{var}(Z) = 1.959$

4. $\mathbb{P}(Y = 0) + \mathbb{P}(Y = 10) = 0.00294$

A small pond contains 1000 fish; 100 are tagged. Suppose that 20 fish are caught. Find each of the following:

1. The probability density function of the number of tagged fish in the sample.
2. The mean and variance of the number of tagged fish in the sample.
3. The probability that the sample contains at least 2 tagged fish.
4. The binomial approximation to the probability in (c).

Answer

Let Y denote the number of tagged fish in the sample

1. $\mathbb{P}(Y = y) = \frac{\binom{100}{y} \binom{900}{20-y}}{\binom{1000}{20}}, \quad y \in \{0, 1, \dots, 20\}$
2. $\mathbb{E}(Y) = 2, \text{ var}(Y) = \frac{196}{111}$
3. $\mathbb{P}(Y \geq 2) = 0.6108$
4. $\mathbb{P}(Y \geq 2) = 0.6083$

Forty percent of the registered voters in a certain district prefer candidate A . Suppose that 10 voters are chosen at random. Find each of the following:

1. The probability density function of the number of voters in the sample who prefer A .
2. The mean and variance of the number of voters in the sample who prefer A .
3. The probability that at least 5 voters in the sample prefer A .

Answer

1. $\mathbb{P}(Y = y) = \binom{10}{y} (0.4)^y (0.6)^{10-y}, \quad y \in \{0, 1, \dots, 10\}$
2. $\mathbb{E}(Y) = 4, \text{ var}(Y) = 2.4$
3. $\mathbb{P}(Y \geq 5) = 0.3669$

Suppose that 10 memory chips are sampled at random and without replacement from a batch of 100 chips. The chips are tested and 2 are defective. Estimate the number of defective chips in the entire batch.

Answer

20

A voting district has 5000 registered voters. Suppose that 100 voters are selected at random and polled, and that 40 prefer candidate A . Estimate the number of voters in the district who prefer candidate A .

Answer

2000

From a certain lake, 200 fish are caught, tagged and returned to the lake. Then 100 fish are caught and it turns out that 10 are tagged. Estimate the population of fish in the lake.

Answer

2000

Cards

Recall that the general *card experiment* is to select n cards at random and without replacement from a standard deck of 52 cards. The special case $n = 5$ is the *poker experiment* and the special case $n = 13$ is the *bridge experiment*.

In a poker hand, find the probability density function, mean, and variance of the following random variables:

1. The number of spades
2. The number of aces

Answer

Let U denote the number of spades and V the number of aces.

1. $\mathbb{P}(U = u) = \frac{\binom{13}{u} \binom{39}{5-u}}{\binom{52}{5}}, \quad u \in \{0, 1, \dots, 5\}, \mathbb{E}(U) = \frac{5}{4}, \text{var}(U) = \frac{235}{272}$
2. $\mathbb{P}(V = v) = \frac{\binom{4}{v} \binom{48}{5-v}}{\binom{52}{5}}, \quad v \in \{0, 1, 2, 3, 4\}, \mathbb{E}(V) = \frac{5}{13}, \text{var}(V) = \frac{940}{2873}$

In a bridge hand, find each of the following:

1. The probability density function, mean, and variance of the number of hearts
2. The probability density function, mean, and variance of the number of honor cards (ace, king, queen, jack, or 10).
3. The probability that the hand has no honor cards. A hand of this kind is known as a Yarborough, in honor of Second Earl of Yarborough.

Answer

Let U denote the number of hearts and V the number of honor cards.

1. $\mathbb{P}(U = u) = \frac{\binom{13}{u} \binom{39}{13-u}}{\binom{52}{13}}, \quad u \in \{0, 1, \dots, 13\}, \mathbb{E}(U) = \frac{13}{4}, \text{var}(U) = \frac{507}{272}$
2. $\mathbb{P}(V = v) = \frac{\binom{20}{v} \binom{32}{13-v}}{\binom{52}{13}}, \quad v \in \{0, 1, \dots, 13\}, \mathbb{E}(V) = 5, \text{var}(V) = 2.353$
3. $\frac{5394}{9\,860\,459} \approx 0.000547$

The Randomized Urn

An interesting thing to do in almost any parametric probability model is to randomize one or more of the parameters. Done in the right way, this often leads to an interesting new parametric model, since the distribution of the randomized parameter will often itself belong to a parametric family. This is also the natural setting to apply Bayes' theorem.

In this section, we will randomize the number of type 1 objects in the basic hypergeometric model. Specifically, we assume that we have m objects in the population, as before. However, instead of a fixed number r of type 1 objects, we assume that each of the m objects in the population, independently of the others, is type 1 with probability p and type 0 with probability $1 - p$. We have eliminated one parameter, r , in favor of a new parameter p with values in the interval $[0, 1]$. Let U_i denote the type of the i th object in the population, so that $\mathbf{U} = (U_1, U_2, \dots, U_n)$ is a sequence of Bernoulli trials with success parameter p . Let $V = \sum_{i=1}^m U_i$ denote the number of type 1 objects in the population, so that V has the binomial distribution with parameters m and p .

As before, we sample n object from the population. Again we let X_i denote the type of the i th object sampled, and we let $Y = \sum_{i=1}^n X_i$ denote the number of type 1 objects in the sample. We will consider sampling with and without replacement. In the first case, the sample size can be any positive integer, but in the second case, the sample size cannot exceed the population size. The key technique in the analysis of the randomized urn is to *condition on V* . If we know that $V = r$, then the model reduces to the model studied above: a population of size m with r type 1 objects, and a sample of size n .

With either type of sampling, $\mathbb{P}(X_i = 1) = p$

Proof

$$\mathbb{P}(X_i = 1) = \mathbb{E}[\mathbb{P}(X_i = 1 \mid V)] = \mathbb{E}(V/m) = p$$

Thus, in either model, \mathbf{X} is a sequence of identically distributed indicator variables. Ah, but what about dependence?

Suppose that the sampling is without replacement. Let $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ and let $y = \sum_{i=1}^n x_i$. Then

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = p^y (1-p)^{n-y} \quad (12.2.13)$$

Proof

Conditioning on V gives

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \mathbb{E}[\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid V)] = \mathbb{E}\left[\frac{V^{(y)}(m-V)^{(n-y)}}{m^{(n)}}\right] \quad (12.2.14)$$

Now let $G(s, t) = \mathbb{E}(s^V t^{m-V})$. Note that G is a probability generating function of sorts. From the binomial theorem, $G(s, t) = [ps + (1-p)t]^m$. Let $G_{j,k}$ denote the partial derivative of G of order $j+k$, with j derivatives with respect to the first argument and k derivatives with respect to the second argument. From the definition of G , $G_{j,k}(1, 1) = \mathbb{E}[V^{(j)}(m-V)^{(k)}]$. But from the binomial representation, $G_{j,k}(1, 1) = m^{j+k} p^j (1-p)^k$

From the joint distribution in the previous exercise, we see that \mathbf{X} is a sequence of Bernoulli trials with success parameter p , and hence Y has the binomial distribution with parameters n and p . We could also argue that \mathbf{X} is a Bernoulli trials sequence directly, by noting that $\{X_1, X_2, \dots, X_n\}$ is a randomly chosen subset of $\{U_1, U_2, \dots, U_m\}$.

Suppose now that the sampling is with replacement. Again, let $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ and let $y = \sum_{i=1}^n x_i$. Then

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \mathbb{E} \left[\frac{V^y (m - V)^{n-y}}{m^n} \right] \quad (12.2.15)$$

Proof

The result follows as before by conditioning on V :

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \mathbb{E} [\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid V)] = \mathbb{E} \left[\frac{V^y (m - V)^{n-y}}{m^n} \right] \quad (12.2.16)$$

A closed form expression for the joint distribution of \mathbf{X} , in terms of the parameters m , n , and p is not easy, but it is at least clear that the joint distribution will not be the same as the one when the sampling is without replacement. In particular, \mathbf{X} is a dependent sequence. Note however that \mathbf{X} is an exchangeable sequence, since the joint distribution is invariant under a permutation of the coordinates (this is a simple consequence of the fact that the joint distribution depends only on the sum y).

The probability density function of Y is given by

$$\mathbb{P}(Y = y) = \binom{n}{y} \mathbb{E} \left[\frac{V^y (m - V)^{n-y}}{m^n} \right], \quad y \in \{0, 1, \dots, n\} \quad (12.2.17)$$

Suppose that i and j are distinct indices. The covariance and correlation of (X_i, X_j) are

1. $\text{cov}(X_i, X_j) = \frac{p(1-p)}{m}$
2. $\text{cor}(X_i, X_j) = \frac{1}{m}$

Proof

Conditioning on V once again we have $\mathbb{P}(X_i = 1, X_j = 1) = \mathbb{E} \left[\left(\frac{V}{m} \right)^2 \right] = \frac{p(1-p)}{m} + p^2$. The results now follow from standard formulas for covariance and correlation.

The mean and variance of Y are

1. $\mathbb{E}(Y) = np$
2. $\text{var}(Y) = np(1-p) \frac{m+n-1}{m}$

Proof

Part (a) follows from the [distribution of the indicator variables](#) above, and the additive property of expected value. Part (b) follows from the previous result on [covariance](#). Recall again that the variance of Y is the sum of $\text{cov}(X_i, X_j)$ over all i and j .

Let's conclude with an interesting observation: For the randomized urn, \mathbf{X} is a sequence of independent variables when the sampling is without replacement but a sequence of dependent variables when the sampling is with replacement—just the opposite of the situation for the deterministic urn with a fixed number of type 1 objects.

This page titled [12.2: The Hypergeometric Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.3: The Multivariate Hypergeometric Distribution

Basic Theory

The Multitype Model

As in the basic sampling model, we start with a finite population D consisting of m objects. In this section, we suppose in addition that each object is one of k types; that is, we have a *multitype population*. For example, we could have an urn with balls of several different colors, or a population of voters who are either *democrat*, *republican*, or *independent*. Let D_i denote the subset of all type i objects and let $m_i = \#(D_i)$ for $i \in \{1, 2, \dots, k\}$. Thus $D = \bigcup_{i=1}^k D_i$ and $m = \sum_{i=1}^k m_i$. The dichotomous model considered earlier is clearly a special case, with $k = 2$.

As in the basic sampling model, we sample n objects at random from D . Thus the outcome of the experiment is $\mathbf{X} = (X_1, X_2, \dots, X_n)$ where $X_i \in D$ is the i th object chosen. Now let Y_i denote the number of type i objects in the sample, for $i \in \{1, 2, \dots, k\}$. Note that $\sum_{i=1}^k Y_i = n$ so if we know the values of $k-1$ of the counting variables, we can find the value of the remaining counting variable. As with any counting variable, we can express Y_i as a sum of indicator variables:

For $i \in \{1, 2, \dots, k\}$

$$Y_i = \sum_{j=1}^n \mathbf{1}(X_j \in D_i) \quad (12.3.1)$$

We assume initially that the sampling is without replacement, since this is the realistic case in most applications.

The Joint Distribution

Basic combinatorial arguments can be used to derive the probability density function of the random vector of counting variables. Recall that since the sampling is without replacement, the unordered sample is uniformly distributed over the combinations of size n chosen from D .

The probability density function of (Y_1, Y_2, \dots, Y_k) is given by

$$\mathbb{P}(Y_1 = y_1, Y_2 = y_2, \dots, Y_k = y_k) = \frac{\binom{m_1}{y_1} \binom{m_2}{y_2} \dots \binom{m_k}{y_k}}{\binom{m}{n}}, \quad (y_1, y_2, \dots, y_k) \in \mathbb{N}^k \text{ with } \sum_{i=1}^k y_i = n \quad (12.3.2)$$

Proof

The binomial coefficient $\binom{m_i}{y_i}$ is the number of unordered subsets of D_i (the type i objects) of size y_i . The binomial coefficient $\binom{m}{n}$ is the number of unordered samples of size n chosen from D . Thus the result follows from the multiplication principle of combinatorics and the uniform distribution of the unordered sample

The distribution of (Y_1, Y_2, \dots, Y_k) is called the *multivariate hypergeometric distribution* with parameters $m, (m_1, m_2, \dots, m_k)$, and n . We also say that $(Y_1, Y_2, \dots, Y_{k-1})$ has this distribution (recall again that the values of any $k-1$ of the variables determines the value of the remaining variable). Usually it is clear from context which meaning is intended. The ordinary hypergeometric distribution corresponds to $k = 2$.

An alternate form of the probability density function of Y_1, Y_2, \dots, Y_k is

$$\mathbb{P}(Y_1 = y_1, Y_2 = y_2, \dots, Y_k = y_k) = \binom{n}{y_1, y_2, \dots, y_k} \frac{m_1^{(y_1)} m_2^{(y_2)} \dots m_k^{(y_k)}}{m^{(n)}}, \quad (y_1, y_2, \dots, y_k) \in \mathbb{N}^k \text{ with } \sum_{i=1}^k y_i = n \quad (12.3.3)$$

Combinatorial Proof

The combinatorial proof is to consider the ordered sample, which is uniformly distributed on the set of permutations of size n from D . The multinomial coefficient on the right is the number of ways to partition the index set $\{1, 2, \dots, n\}$ into k groups where group i has y_i elements (these are the coordinates of the type i objects). The number of (ordered) ways to select the type i objects is $m_i^{(y_i)}$. The denominator $m^{(n)}$ is the number of ordered samples of size n chosen from D .

Algebraic Proof

There is also a simple algebraic proof, starting from the [first version](#) of probability density function above. Write each binomial coefficient $\binom{a}{j} = a^{(j)} / j!$ and rearrange a bit.

The Marginal Distributions

For $i \in \{1, 2, \dots, k\}$, Y_i has the hypergeometric distribution with parameters m, m_i , and n

$$\mathbb{P}(Y_i = y) = \frac{\binom{m_i}{y} \binom{m-m_i}{n-y}}{\binom{m}{n}}, \quad y \in \{0, 1, \dots, n\} \quad (12.3.4)$$

Proof

An analytic proof is possible, by starting with the [first version](#) or the [second version](#) of the joint PDF and summing over the unwanted variables. However, a probabilistic proof is much better: Y_i is the number of type i objects in a sample of size n chosen at random (and without replacement) from a population of m objects, with m_i of type i and the remaining $m - m_i$ not of this type.

Grouping

The multivariate hypergeometric distribution is preserved when the counting variables are combined. Specifically, suppose that (A_1, A_2, \dots, A_l) is a partition of the index set $\{1, 2, \dots, k\}$ into nonempty, disjoint subsets. Let $W_j = \sum_{i \in A_j} Y_i$ and $r_j = \sum_{i \in A_j} m_i$ for $j \in \{1, 2, \dots, l\}$

(W_1, W_2, \dots, W_l) has the multivariate hypergeometric distribution with parameters $m, (r_1, r_2, \dots, r_l)$, and n .

Proof

Again, an analytic proof is possible, but a probabilistic proof is much better. Effectively, we now have a population of m objects with l types, and r_i is the number of objects of the new type i . As before we sample n objects without replacement, and W_i is the number of objects in the sample of the new type i .

Note that the [marginal distribution](#) of Y_i given above is a special case of grouping. We have two types: type i and not type i . More generally, the marginal distribution of any subsequence of (Y_1, Y_2, \dots, Y_n) is hypergeometric, with the appropriate parameters.

Conditioning

The multivariate hypergeometric distribution is also preserved when some of the counting variables are observed. Specifically, suppose that (A, B) is a partition of the index set $\{1, 2, \dots, k\}$ into nonempty, disjoint subsets. Suppose that we observe $Y_j = y_j$ for $j \in B$. Let $z = n - \sum_{j \in B} y_j$ and $r = \sum_{i \in A} m_i$.

The conditional distribution of $(Y_i : i \in A)$ given $(Y_j = y_j : j \in B)$ is multivariate hypergeometric with parameters $r, (m_i : i \in A)$, and z .

Proof

Once again, an analytic argument is possible using the definition of conditional probability and the appropriate joint distributions. A probabilistic argument is much better. Effectively, we are selecting a sample of size z from a population of size r , with m_i objects of type i for each $i \in A$.

Combinations of the [grouping result](#) and the [conditioning result](#) can be used to compute any marginal or conditional distributions of the counting variables.

Moments

We will compute the mean, variance, covariance, and correlation of the counting variables. Results from the hypergeometric distribution and the representation in terms of [indicator variables](#) are the main tools.

For $i \in \{1, 2, \dots, k\}$,

1. $\mathbb{E}(Y_i) = n \frac{m_i}{m}$
2. $\text{var}(Y_i) = n \frac{m_i}{m} \frac{m-m_i}{m} \frac{m-n}{m-1}$

Proof

This follows immediately, since Y_i has the hypergeometric distribution with parameters m, m_i , and n .

Now let $I_{ti} = \mathbf{1}(X_t \in D_i)$, the indicator variable of the event that the t th object selected is type i , for $t \in \{1, 2, \dots, n\}$ and $i \in \{1, 2, \dots, k\}$.

Suppose that r and s are distinct elements of $\{1, 2, \dots, n\}$ and i and j are distinct elements of $\{1, 2, \dots, k\}$. Then

$$\text{cov}(I_{ri}, I_{rj}) = -\frac{m_i}{m} \frac{m_j}{m} \quad (12.3.5)$$

$$\text{cov}(I_{ri}, I_{sj}) = \frac{1}{m-1} \frac{m_i}{m} \frac{m_j}{m} \quad (12.3.6)$$

Proof

Recall that if A and B are events, then $\text{cov}(A, B) = \mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)$. In the first case the events are that sample item r is type i and that sample item r is type j . These events are disjoint, and the individual probabilities are $\frac{m_i}{m}$ and $\frac{m_j}{m}$. In the second case, the events are that sample item r is type i and that sample item s is type j . The probability that both events occur is $\frac{m_i}{m} \frac{m_j}{m-1}$ while the individual probabilities are the same as in the first case.

Suppose again that r and s are distinct elements of $\{1, 2, \dots, n\}$ and i and j are distinct elements of $\{1, 2, \dots, k\}$. Then

$$\text{cor}(I_{ri}, I_{rj}) = -\sqrt{\frac{m_i}{m-m_i} \frac{m_j}{m-m_j}} \quad (12.3.7)$$

$$\text{cor}(I_{ri}, I_{sj}) = \frac{1}{m-1} \sqrt{\frac{m_i}{m-m_i} \frac{m_j}{m-m_j}} \quad (12.3.8)$$

Proof

This follows from the previous result and the definition of correlation. Recall that if I is an indicator variable with parameter p then $\text{var}(I) = p(1-p)$.

In particular, I_{ri} and I_{rj} are negatively correlated while I_{ri} and I_{sj} are positively correlated.

For distinct $i, j \in \{1, 2, \dots, k\}$,

$$\text{cov}(Y_i, Y_j) = -n \frac{m_i}{m} \frac{m_j}{m} \frac{m-n}{m-1} \quad (12.3.9)$$

$$\text{cor}(Y_i, Y_j) = -\sqrt{\frac{m_i}{m-m_i} \frac{m_j}{m-m_j}} \quad (12.3.10)$$

Sampling with Replacement

Suppose now that the sampling is with replacement, even though this is usually not realistic in applications.

The types of the objects in the sample form a sequence of n multinomial trials with parameters $(m_1/m, m_2/m, \dots, m_k/m)$

The following results now follow immediately from the general theory of multinomial trials, although modifications of the arguments above could also be used.

(Y_1, Y_2, \dots, Y_k) has the multinomial distribution with parameters n and $(m_1/m, m_2/m, \dots, m_k/m)$

$$\mathbb{P}(Y_1 = y_1, Y_2 = y_2, \dots, Y_k = y_k) = \binom{n}{y_1, y_2, \dots, y_k} \frac{m_1^{y_1} m_2^{y_2} \dots m_k^{y_k}}{m^n}, \quad (y_1, y_2, \dots, y_k) \in \mathbb{N}^k \text{ with } \sum_{i=1}^k y_i = n \quad (12.3.11)$$

For distinct $i, j \in \{1, 2, \dots, k\}$,

1. $\mathbb{E}(Y_i) = n \frac{m_i}{m}$
2. $\text{var}(Y_i) = n \frac{m_i}{m} \frac{m-m_i}{m}$
3. $\text{cov}(Y_i, Y_j) = -n \frac{m_i}{m} \frac{m_j}{m}$
4. $\text{cor}(Y_i, Y_j) = -\sqrt{\frac{m_i}{m-m_i} \frac{m_j}{m-m_j}}$

Comparing with our previous results, note that the means and correlations are the same, whether sampling with or without replacement. The variances and covariances are smaller when sampling without replacement, by a factor of the *finite population correction factor* $(m-n)/(m-1)$

Convergence to the Multinomial Distribution

Suppose that the population size m is very large compared to the sample size n . In this case, it seems reasonable that sampling *without* replacement is not too much different than sampling *with* replacement, and hence the multivariate hypergeometric distribution should be well approximated by the multinomial. The following exercise makes this observation precise. Practically, it is a valuable result, since in many cases we do not know the population size exactly. For the approximate multinomial distribution, we do not need to know m_i and m individually, but only in the ratio m_i/m .

Suppose that m_i depends on m and that $m_i/m \rightarrow p_i$ as $m \rightarrow \infty$ for $i \in \{1, 2, \dots, k\}$. For fixed n , the multivariate hypergeometric probability density function with parameters $m, (m_1, m_2, \dots, m_k)$, and n converges to the multinomial probability density function with

parameters n and (p_1, p_2, \dots, p_k) .

Proof

Consider the [second version](#) of the hypergeometric probability density function. In the fraction, there are n factors in the denominator and n in the numerator. If we group the factors to form a product of n fractions, then each fraction in group i converges to p_i .

Examples and Applications

A population of 100 voters consists of 40 republicans, 35 democrats and 25 independents. A random sample of 10 voters is chosen. Find each of the following:

1. The joint density function of the number of republicans, number of democrats, and number of independents in the sample
2. The mean of each variable in (a).
3. The variance of each variable in (a).
4. The covariance of each pair of variables in (a).
5. The probability that the sample contains at least 4 republicans, at least 3 democrats, and at least 2 independents.

Answer

1. $\mathbb{P}(X = x, Y = y, Z = z) = \frac{\binom{40}{x} \binom{35}{y} \binom{25}{z}}{\binom{100}{10}}$ for $x, y, z \in \mathbb{N}$ with $x + y + z = 10$
2. $\mathbb{E}(X) = 4, \mathbb{E}(Y) = 3.5, \mathbb{E}(Z) = 2.5$
3. $\text{var}(X) = 2.1818, \text{var}(Y) = 2.0682, \text{var}(Z) = 1.7045$
4. $\text{cov}(X, Y) = -1.6346, \text{cov}(X, Z) = -0.9091, \text{cov}(Y, Z) = -0.7955$
5. 0.2474

Cards

Recall that the *general card experiment* is to select n cards at random and without replacement from a standard deck of 52 cards. The special case $n = 5$ is the *poker experiment* and the special case $n = 13$ is the *bridge experiment*.

In a bridge hand, find the probability density function of

1. The number of spades, number of hearts, and number of diamonds.
2. The number of spades and number of hearts.
3. The number of spades.
4. The number of red cards and the number of black cards.

Answer

Let X, Y, Z, U , and V denote the number of spades, hearts, diamonds, red cards, and black cards, respectively, in the hand.

1. $\mathbb{P}(X = x, Y = y, Z = z) = \frac{\binom{13}{x} \binom{13}{y} \binom{13}{z} \binom{13}{13-x-y-z}}{\binom{52}{13}}$ for $x, y, z \in \mathbb{N}$ with $x + y + z \leq 13$
2. $\mathbb{P}(X = x, Y = y) = \frac{\binom{13}{x} \binom{13}{y} \binom{26}{13-x-y}}{\binom{52}{13}}$ for $x, y \in \mathbb{N}$ with $x + y \leq 13$
3. $\mathbb{P}(X = x) = \frac{\binom{13}{x} \binom{39}{13-x}}{\binom{52}{13}}$ for $x \in \{0, 1, \dots, 13\}$
4. $\mathbb{P}(U = u, V = v) = \frac{\binom{26}{u} \binom{26}{v}}{\binom{52}{13}}$ for $u, v \in \mathbb{N}$ with $u + v = 13$

In a bridge hand, find each of the following:

1. The mean and variance of the number of spades.
2. The covariance and correlation between the number of spades and the number of hearts.
3. The mean and variance of the number of red cards.

Answer

Let X, Y , and U denote the number of spades, hearts, and red cards, respectively, in the hand.

1. $\mathbb{E}(X) = \frac{13}{4}, \text{var}(X) = \frac{507}{272}$
2. $\text{cov}(X, Y) = -\frac{169}{272}$
3. $\mathbb{E}(U) = \frac{13}{2}, \text{var}(U) = \frac{169}{272}$

In a bridge hand, find each of the following:

1. The conditional probability density function of the number of spades and the number of hearts, given that the hand has 4 diamonds.
2. The conditional probability density function of the number of spades given that the hand has 3 hearts and 2 diamonds.

Answer

Let X , Y and Z denote the number of spades, hearts, and diamonds respectively, in the hand.

1. $\mathbb{P}(X = x, Y = y, | Z = 4) = \frac{\binom{13}{x} \binom{13}{y} \binom{22}{9-x-y}}{\binom{48}{9}}$ for $x, y \in \mathbb{N}$ with $x + y \leq 9$
2. $\mathbb{P}(X = x | Y = 3, Z = 2) = \frac{\binom{13}{x} \binom{34}{8-x}}{\binom{47}{8}}$ for $x \in \{0, 1, \dots, 8\}$

In the card experiment, a hand that does not contain any cards of a particular suit is said to be *void* in that suit.

Use the inclusion-exclusion rule to show that the probability that a poker hand is void in at least one suit is

$$\frac{1913496}{2598960} \approx 0.736 \quad (12.3.12)$$

In the card experiment, set $n = 5$. Run the simulation 1000 times and compute the relative frequency of the event that the hand is void in at least one suit. Compare the relative frequency with the true probability given in the previous exercise.

Use the inclusion-exclusion rule to show that the probability that a bridge hand is void in at least one suit is

$$\frac{32427298180}{635013559600} \approx 0.051 \quad (12.3.13)$$

This page titled [12.3: The Multivariate Hypergeometric Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.4: Order Statistics

Basic Theory

Definitions

Suppose that the objects in our population are numbered from 1 to m , so that $D = \{1, 2, \dots, m\}$. For example, the population might consist of manufactured items, and the labels might correspond to serial numbers. As in the basic sampling model we select n objects at random, without replacement from D . Thus the outcome is $\mathbf{X} = (X_1, X_2, \dots, X_n)$ where $X_i \in S$ is the i th object chosen. Recall that \mathbf{X} is uniformly distributed over the set of permutations of size n chosen from D . Recall also that $\mathbf{W} = \{X_1, X_2, \dots, X_n\}$ is the unordered sample, which is uniformly distributed on the set of combinations of size n chosen from D .

For $i \in \{1, 2, \dots, n\}$ let $X_{(i)} = i$ th smallest element of $\{X_1, X_2, \dots, X_n\}$. The random variable $X_{(i)}$ is known as the *order statistic* of order i for the sample \mathbf{X} . In particular, the extreme order statistics are

$$X_{(1)} = \min\{X_1, X_2, \dots, X_n\} \quad (12.4.1)$$

$$X_{(n)} = \max\{X_1, X_2, \dots, X_n\} \quad (12.4.2)$$

Random variable $X_{(i)}$ takes values in $\{i, i+1, \dots, m-n+i\}$ for $i \in \{1, 2, \dots, n\}$.

We will denote the vector of order statistics by $\mathbf{Y} = (X_{(1)}, X_{(2)}, \dots, X_{(n)})$. Note that \mathbf{Y} takes values in

$$L = \{(y_1, y_2, \dots, y_n) \in D^n : y_1 < y_2 < \dots < y_n\} \quad (12.4.3)$$

Run the order statistic experiment. Note that you can vary the population size m and the sample size n . The order statistics are recorded on each update.

Distributions

L has $\binom{m}{n}$ elements and \mathbf{Y} is uniformly distributed on L .

Proof

For $\mathbf{y} = (y_1, y_2, \dots, y_n) \in L$, $\mathbf{Y} = \mathbf{y}$ if and only if \mathbf{X} is one of the $n!$ permutations of \mathbf{y} . Hence $\mathbb{P}(\mathbf{Y} = \mathbf{y}) = n! / m^{(n)} = 1 / \binom{m}{n}$.

The probability density function of $X_{(i)}$ is

$$\mathbb{P}[X_{(i)} = x] = \frac{\binom{x-1}{i-1} \binom{m-x}{n-i}}{\binom{m}{n}}, \quad x \in \{i, i+1, \dots, m-n+i\} \quad (12.4.4)$$

Proof

The event that the i th order statistic is x means that $i-1$ sample values are less than x and $n-i$ are greater than x , and of course, one of the sample values is x . By the multiplication principle of combinatorics, the number of unordered samples corresponding to this event is $\binom{x-1}{i-1} \binom{m-x}{n-i}$. The total number of unordered samples is $\binom{m}{n}$.

In the order statistic experiment, vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the experiment 1000 times and compare the relative frequency function to the probability density function.

Moments

The [probability density function](#) of $X_{(i)}$ above can be used to obtain an interesting identity involving the binomial coefficients. This identity, in turn, can be used to find the mean and variance of $X_{(i)}$.

For $i, n, m \in \mathbb{N}_+$ with $i \leq n \leq m$,

$$\sum_{k=i}^{m-n+i} \binom{k-1}{i-1} \binom{m-k}{n-i} = \binom{m}{n} \quad (12.4.5)$$

Proof

This result follows immediately from the [probability density function](#) of $X_{(i)}$ above

The expected value of $X_{(i)}$ is

$$\mathbb{E}[X_{(i)}] = i \frac{m+1}{n+1} \quad (12.4.6)$$

Proof

We start with the definition of expected value. Recall that $x \binom{x-1}{i-1} = i \binom{x}{i}$. Next we use the [identity](#) above with m replaced with $m+1$, n replaced with $n+1$, and i replaced with $i+1$. Simplifying gives the result.

The variance of $X_{(i)}$ is

$$\text{var}[X_{(i)}] = i(n-i+1) \frac{(m+1)(m-n)}{(n+1)^2(n+2)} \quad (12.4.7)$$

Proof

The result follows from another application of the [identity](#) above.

In the order statistic experiment, vary the parameters and note the size and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

Estimators of m Based on Order Statistics

Suppose that the population size m is unknown. In this subsection we consider estimators of m constructed from the various order statistics.

For $i \in \{1, 2, \dots, n\}$, the following statistic is an unbiased estimator of m :

$$U_i = \frac{n+1}{i} X_{(i)} - 1 \quad (12.4.8)$$

Proof

From the [expected value](#) of $X_{(i)}$ above and the linear property of expected value, note that $\mathbb{E}(U_i) = m$.

Since U_i is unbiased, its variance is the *mean square error*, a measure of the quality of the estimator.

The variance of U_i is

$$\text{var}(U_i) = \frac{(m+1)(m-n)(n-i+1)}{i(n+2)} \quad (12.4.9)$$

Proof

This result follows from [variance](#) of $X_{(i)}$ given above and standard properties of variance.

For fixed m and n , $\text{var}(U_i)$ decreases as i increases. Thus, the estimators improve as i increases; in particular, U_n is the best and U_1 the worst.

The *relative efficiency* of U_j with respect to U_i is

$$\frac{\text{var}(U_i)}{\text{var}(U_j)} = \frac{j(n-i+1)}{i(n-j+1)} \quad (12.4.10)$$

Note that the relative efficiency depends only on the orders i and j and the sample size n , but not on the population size m (the unknown parameter). In particular, the relative efficiency of U_n with respect to U_1 is n^2 . For fixed i and j , the asymptotic relative efficiency of U_j to U_i is j/i . Usually, we hope that an estimator improves (in the sense of mean square error) as the sample size n increases (the more information we have, the better our estimate should be). This general idea is known as *consistency*.

$\text{var}(U_n)$ decreases to 0 as n increases from 1 to m , and so U_n is consistent:

$$\text{var}(U_n) = \frac{(m+1)(m-n)}{n(n+2)} \quad (12.4.11)$$

For fixed i , $\text{var}(U_i)$ at first increases and then decreases to 0 as n increases from i to m . Thus, U_i is inconsistent.




Figure 12.4.1: $\text{var}(U_1)$ as a function of n for $m = 100$

An Estimator of m Based on the Sample Mean

In this subsection, we will derive another estimator of the parameter m based on the average of the sample variables $M = \frac{1}{n} \sum_{i=1}^n x_i$, (the sample mean) and compare this estimator with the estimator based on the maximum of the variables (the largest order statistic).

$$\mathbb{E}(M) = \frac{m+1}{2}.$$

Proof

Recall that X_i is uniformly distributed on D for each i and hence $\mathbb{E}(X_i) = \frac{m+1}{2}$.

It follows that $V = 2M - 1$ is an unbiased estimator of m . Moreover, it seems that superficially at least, V uses more information from the sample (since it involves all of the sample variables) than U_n . Could it be better? To find out, we need to compute the variance of the estimator (which, since it is unbiased, is the mean square error). This computation is a bit complicated since the sample variables are dependent. We will compute the variance of the sum as the sum of all of the pairwise covariances.

$$\text{For distinct } i, j \in \{1, 2, \dots, n\}, \text{cov}(X_i, X_j) = -\frac{m+1}{12}.$$

Proof

First recall that given $X_i = x$, X_j is uniformly distributed on $D \setminus \{x\}$. Hence $\mathbb{E}(X_j | X_i = x) = \frac{m(m+1)}{2(m-1)} - \frac{x}{m-1}$. Thus conditioning on X_i gives $\mathbb{E}(X_i X_j) = \frac{(m+1)(3m+2)}{12}$. The result now follows from the standard formula $\text{cov}(X_i, X_j) = \mathbb{E}(X_i X_j) - \mathbb{E}(X_i)\mathbb{E}(X_j)$.

$$\text{For } i \in \{1, 2, \dots, n\}, \text{var}(X_i) = \frac{m^2-1}{12}.$$

Proof

This follows since X_i is uniformly distributed on D .

$$\text{var}(M) = \frac{(m+1)(m-n)}{12n}.$$

Proof

The variance of M is $\frac{1}{n^2}$ times the sum of $\text{cov}(X_i, X_j)$ over all $i, j \in \{1, 2, \dots, n\}$. There are n covariance terms with the value given in the [variance result](#) above (corresponding to $i = j$) and $n^2 - n$ terms with the value given in the [pure covariance](#) result above (corresponding to $i \neq j$). Simplifying gives the result.

$$\text{var}(V) = \frac{(m+1)(m-n)}{3n}.$$

Proof

This follows from the [variance](#) of M above and standard properties of variance.

The variance of V is decreasing with n , so V is also consistent. Let's compute the relative efficiency of the estimator based on the maximum to the estimator based on the mean.

$$\text{var}(V)/\text{var}(U_n) = (n+2)/3.$$

Thus, once again, the estimator based on the maximum is better. In addition to the mathematical analysis, all of the estimators except U_n can sometimes be manifestly worthless by giving estimates that are smaller than some of the sample values.

Sampling with Replacement

If the sampling is *with* replacement, then the sample $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent and identically distributed random variables. The order statistics from such samples are studied in the chapter on Random Samples.

Examples and Applications

Suppose that in a lottery, tickets numbered from 1 to 25 are placed in a bowl. Five tickets are chosen at random and without replacement.

1. Find the probability density function of $X_{(3)}$.
2. Find $\mathbb{E}[X_{(3)}]$.
3. Find $\text{var}[X_{(3)}]$.

Answer

1. $\mathbb{P}[X_{(3)} = x] = \frac{\binom{x-1}{2} \binom{25-x}{2}}{\binom{25}{5}} \text{ for } x \in \{3, 4, \dots, 23\}$
2. $\mathbb{E}[X_{(3)}] = 13$
3. $\text{var}[X_{(3)}] = \frac{130}{7}$

The German Tank Problem

The estimator U_n was used by the Allies during World War II to estimate the number of German tanks m that had been produced. German tanks had serial numbers, and captured German tanks and records formed the sample data. The statistical estimates turned out to be much more accurate than intelligence estimates. Some of the data are given in the table below.

German Tank Data. Source: [Wikipedia](#)

Date	Statistical Estimate	Intelligence Estimate	German Records
June 1940	169	1000	122
June 1941	244	1550	271
August 1942	327	1550	342

One of the morals, evidently, is not to put serial numbers on your weapons!

Suppose that in a certain war, 5 enemy tanks have been captured. The serial numbers are 51, 3, 27, 82, 65. Compute the estimate of m , the total number of tanks, using all of the estimators discussed above.

Answer

1. $u_1 = 17$
2. $u_2 = 80$
3. $u_3 = 101$
4. $u_4 = 96.5$
5. $u_5 = 97.4$
6. $v = 90.2$

In the order statistic experiment, and set $m = 100$ and $n = 10$. Run the experiment 50 times. For each run, compute the estimate of m based on each order statistic. For each estimator, compute the square root of the average of the squares of the errors over the 50 runs. Based on these empirical error estimates, rank the estimators of m in terms of quality.

Suppose that in a certain war, 10 enemy tanks have been captured. The serial numbers are 304, 125, 417, 226, 192, 340, 468, 499, 87, 352. Compute the estimate of m , the total number of tanks, using the estimator based on the maximum and the estimator based on the mean.

Answer

1. $u = 548$
2. $v = 601$

This page titled [12.4: Order Statistics](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.5: The Matching Problem

Definitions and Notation

The Matching Experiment

The *matching experiment* is a random experiment that can be formulated in a number of colorful ways:

- Suppose that n male-female couples are at a party and that the males and females are randomly paired for a dance. A match occurs if a couple happens to be paired together.
- An absent-minded secretary prepares n letters and envelopes to send to n different people, but then randomly stuffs the letters into the envelopes. A match occurs if a letter is inserted in the proper envelope.
- n people with hats have had a bit too much to drink at a party. As they leave the party, each person randomly grabs a hat. A match occurs if a person gets his or her own hat.

These experiments are clearly equivalent from a mathematical point of view, and correspond to selecting a random permutation $\mathbf{X} = (X_1, X_2, \dots, X_n)$ of the population $D_n = \{1, 2, \dots, n\}$. Here are the interpretations for the examples above:

- Number the couples from 1 to n . Then X_i is the number of the woman paired with the i th man.
- Number the letters and corresponding envelopes from 1 to n . Then X_i is the number of the envelope containing the i th letter.
- Number the people and their corresponding hats from 1 to n . Then X_i is the number of the hat chosen by the i th person.

Our modeling assumption, of course, is that \mathbf{X} is uniformly distributed on the sample space of permutations of D_n . The number of objects n is the basic parameter of the experiment. We will also consider the case of sampling *with* replacement from the population D_n , because the analysis is much easier but still provides insight. In this case, \mathbf{X} is a sequence of independent random variables, each uniformly distributed over D_n .

Matches

We will say that a *match* occurs at position j if $X_j = j$. Thus, number of matches is the random variable N defined mathematically by

$$N_n = \sum_{j=1}^n I_j \quad (12.5.1)$$

where $I_j = \mathbf{1}(X_j = j)$ is the indicator variable for the event of match at position j . Our problem is to compute the probability distribution of the number of matches. This is an old and famous problem in probability that was first considered by Pierre-Remond Montmort; it is sometimes referred to as *Montmort's matching problem* in his honor.

Sampling With Replacement

First let's solve the matching problem in the easy case, when the sampling is with replacement. Of course, this is not the way that the matching game is usually played, but the analysis will give us some insight.

(I_1, I_2, \dots, I_n) is a sequence of n Bernoulli trials, with success probability $\frac{1}{n}$.

Proof

The variables are independent since the sampling is with replacement. Since X_j is uniformly distributed, $\mathbb{P}(I_j = 1) = \mathbb{P}(X_j = j) = \frac{1}{n}$.

The number of matches N_n has the binomial distribution with trial parameter n and success parameter $\frac{1}{n}$.

$$\mathbb{P}(N_n = k) = \binom{n}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{n-k}, \quad k \in \{0, 1, \dots, n\} \quad (12.5.2)$$

Proof

This follows immediately from the previous result on [Bernoulli trials](#).

The mean and variance of the number of matches are

1. $\mathbb{E}(N_n) = 1$
2. $\text{var}(N_n) = \frac{n-1}{n}$

Proof

These results follow from the previous result on the [binomial distribution](#) of N_n . Recall that the binomial distribution with parameters n and p has mean np and variance $np(1-p)$.

The distribution of the number of matches converges to the Poisson distribution with parameter 1 as $n \rightarrow \infty$:

$$\mathbb{P}(N_n = k) \rightarrow \frac{e^{-1}}{k!} \text{ as } n \rightarrow \infty \text{ for } k \in \mathbb{N} \quad (12.5.3)$$

Proof

This is a special case of the convergence of the binomial distribution to the Poisson. For a direct proof, note that

$$\mathbb{P}(N_n = k) = \frac{1}{k!} \frac{n^{(k)}}{n^k} \left(1 - \frac{1}{n}\right)^{n-k} \quad (12.5.4)$$

But $\frac{n^{(k)}}{n^k} \rightarrow 1$ as $n \rightarrow \infty$ and $\left(1 - \frac{1}{n}\right)^{n-k} \rightarrow e^{-1}$ as $n \rightarrow \infty$ by a famous limit from calculus.

Sampling Without Replacement

Now let's consider the case of real interest, when the sampling is without replacement, so that \mathbf{X} is a random permutation of the elements of $D_n = \{1, 2, \dots, n\}$.

Counting Permutations with Matches

To find the probability density function of N_n , we need to count the number of permutations of D_n with a specified number of matches. This will turn out to be easy once we have counted the number of permutations with no matches; these are called *derangements* of D_n . We will denote the number of permutations of D_n with exactly k matches by $b_n(k) = \#\{N_n = k\}$ for $k \in \{0, 1, \dots, n\}$. In particular, $b_n(0)$ is the number of derangements of D_n .

The number of derangements is

$$b_n(0) = n! \sum_{j=0}^n \frac{(-1)^j}{j!} \quad (12.5.5)$$

Proof

By the complement rule for counting measure $b_n(0) = n! - \#(\bigcup_{i=1}^n \{X_i = i\})$. From the inclusion-exclusion formula,

$$b_n(0) = n! - \sum_{j=1}^n (-1)^{j-1} \sum_{J \subseteq D_n, \#(J)=j} \#\{X_i = i \text{ for all } i \in J\} \quad (12.5.6)$$

But if $J \subseteq D_n$ with $\#(J) = j$ then $\#\{X_i = i \text{ for all } i \in J\} = (n-j)!$. Finally, the number of subsets J of D_n with $\#(J) = j$ is $\binom{n}{j}$. Substituting into the displayed equation and simplifying gives the result.

The number of permutations with exactly k matches is

$$b_n(k) = \frac{n!}{k!} \sum_{j=0}^{n-k} \frac{(-1)^j}{j!}, \quad k \in \{0, 1, \dots, n\} \quad (12.5.7)$$

Proof

The following is two-step procedure that generates all permutations with exactly k matches: First select the k integers that will match. The number of ways of performing this step is $\binom{n}{k}$. Second, select a permutation of the remaining $n-k$ integers with

no matches. The number of ways of performing this step is $b_{n-k}(0)$. By the multiplication principle of combinatorics it follows that $b_n(k) = \binom{n}{k} b_{n-k}(0)$. Using the result above for [derangements](#) and simplifying gives the results.

The Probability Density Function

The probability density function of the number of matches is

$$\mathbb{P}(N_n = k) = \frac{1}{k!} \sum_{j=0}^{n-k} \frac{(-1)^j}{j!}, \quad k \in \{0, 1, \dots, n\} \quad (12.5.8)$$

Proof

This follows directly from the result above on [permutations with matches](#), since $\mathbb{P}(N_n = k) = \#\{N_n = k\} / n!$.

In the matching experiment, vary the parameter n and note the shape and location of the probability density function. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.

$$\mathbb{P}(N_n = n - 1) = 0.$$

Proof

A simple *probabilistic* proof is to note that the event is impossible—if there are $n - 1$ matches, then there must be n matches. An algebraic proof can also be constructed from the [probability density function](#) of N_n above.

The distribution of the number of matches converges to the Poisson distribution with parameter 1 as $n \rightarrow \infty$:

$$\mathbb{P}(N_n = k) \rightarrow \frac{e^{-1}}{k!} \text{ as } n \rightarrow \infty, \quad k \in \mathbb{N} \quad (12.5.9)$$

Proof

From the power series for the exponential function,

$$\sum_{j=0}^{n-k} \frac{(-1)^j}{j!} \rightarrow \sum_{j=0}^{\infty} \frac{(-1)^j}{j!} = e^{-1} \text{ as } n \rightarrow \infty \quad (12.5.10)$$

So the result follows from the [probability density function](#) of N_n above.

The convergence is remarkably rapid.

In the matching experiment, increase n and note how the probability density function stabilizes rapidly. For selected values of n , run the simulation 1000 times and compare the relative frequency function to the probability density function.

Moments

The mean and variance of the number of matches could be computed directly from the distribution. However, it is much better to use the representation in terms of indicator variables. The exchangeable property is an important tool in this section.

$$\mathbb{E}(I_j) = \frac{1}{n} \text{ for } j \in \{1, 2, \dots, n\}.$$

Proof

X_j is uniformly distributed on D_n for each j so $\mathbb{P}(I_j = 1) = \mathbb{P}(X_j = x) = \frac{1}{n}$.

$$\mathbb{E}(N_n) = 1 \text{ for each } n$$

Proof

This follows from the [previous result](#) and basic properties of expected value.

Thus, the expected number of matches is 1, regardless of n , just as when the sampling is [with replacement](#).

$$\text{var}(I_j) = \frac{n-1}{n^2} \text{ for } j \in \{1, 2, \dots, n\}.$$

Proof

This follows from $\mathbb{P}(I_j = 1) = \frac{1}{n}$.

A match in one position would seem to make it more likely that there would be a match in another position. Thus, we might guess that the indicator variables are positively correlated.

For distinct $j, k \in \{1, 2, \dots, n\}$,

$$1. \text{cov}(I_j, I_k) = \frac{1}{n^2(n-1)}$$

$$2. \text{cor}(I_j, I_k) = \frac{1}{(n-1)^2}$$

Proof

Note that $I_j I_k$ is the indicator variable of the event of a match in position j and a match in position k . Hence by the exchangeable property $\mathbb{P}(I_j I_k = 1) = \mathbb{P}(I_j = 1) \mathbb{P}(I_k = 1 | I_j = 1) = \frac{1}{n} \frac{1}{n-1}$. As before, $\mathbb{P}(I_j = 1) = \mathbb{P}(I_k = 1) = \frac{1}{n}$. The results now follow from standard computational formulas for covariance and correlation.

Note that when $n = 2$, the event that there is a match in position 1 is perfectly correlated with the event that there is a match in position 2. This makes sense, since there will either be 0 matches or 2 matches.

$$\text{var}(N_n) = 1 \text{ for every } n \in \{2, 3, \dots\}.$$

Proof

This follows from the previous two results on the [variance](#) and the [covariance](#) of the indicator variables, and basic properties of covariance. Recall that $\text{var}(N_n) = \sum_{j=1}^n \sum_{k=1}^n \text{cov}(I_j, I_k)$.

In the matching experiment, vary the parameter n and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameter, run the simulation 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

For distinct $j, k \in \{1, 2, \dots, n\}$, $\text{cov}(I_j, I_k) \rightarrow 0$ as $n \rightarrow \infty$.

Thus, the event that a match occurs in position j is nearly independent of the event that a match occurs in position k if n is large. For large n , the indicator variables behave nearly like n Bernoulli trials with success probability $\frac{1}{n}$, which of course, is what happens when the sampling is [with replacement](#).

A Recursion Relation

In this subsection, we will give an alternate derivation of the distribution of the number of matches, in a sense by *embedding* the experiment with parameter n into the experiment with parameter $n + 1$.

The probability density function of the number of matches satisfies the following recursion relation and initial condition:

$$1. \mathbb{P}(N_n = k) = (k+1) \mathbb{P}(N_{n+1} = k+1) \text{ for } k \in \{0, 1, \dots, n\}.$$

$$2. \mathbb{P}(N_1 = 1) = 1.$$

Proof

First, consider the random permutation $(X_1, X_2, \dots, X_n, X_{n+1})$ of D_{n+1} . Note that (X_1, X_2, \dots, X_n) is a random permutation of D_n if and only if $X_{n+1} = n+1$ if and only if $I_{n+1} = 1$. It follows that

$$\mathbb{P}(N_n = k) = \mathbb{P}(N_{n+1} = k+1 | I_{n+1} = 1), \quad k \in \{0, 1, \dots, n\} \quad (12.5.11)$$

From the definition of conditional probability argument we have

$$\mathbb{P}(N_n = k) = \mathbb{P}(N_{n+1} = k+1) \frac{\mathbb{P}(I_{n+1} = 1 \mid N_{n+1} = k+1)}{\mathbb{P}(I_{n+1} = 1)}, \quad k \in \{0, 1, \dots, n\} \quad (12.5.12)$$

But $\mathbb{P}(I_{n+1} = 1) = \frac{1}{n+1}$ and $\mathbb{P}(I_{n+1} = 1 \mid N_{n+1} = k+1) = \frac{k+1}{n+1}$. Substituting into the last displayed equation gives the recurrence relation. The initial condition is obvious, since if $n = 1$ we must have one match.

This result can be used to obtain the probability density function of N_n recursively for any n .

The Probability Generating Function

Next recall that the probability generating function of N_n is given by

$$G_n(t) = \mathbb{E}(t^{N_n}) = \sum_{j=0}^n \mathbb{P}(N_n = j)t^j, \quad t \in \mathbb{R} \quad (12.5.13)$$

The family of probability generating functions satisfies the following differential equations and ancillary conditions:

1. $G'_{n+1}(t) = G_n(t)$ for $t \in \mathbb{R}$ and $n \in \mathbb{N}_+$
2. $G_n(1) = 1$ for $n \in \mathbb{N}_+$

Note also that $G_1(t) = t$ for $t \in \mathbb{R}$. Thus, the system of differential equations can be used to compute G_n for any $n \in \mathbb{N}_+$.

In particular, for $t \in \mathbb{R}$,

1. $G_2(t) = \frac{1}{2} + \frac{1}{2}t^2$
2. $G_3(t) = \frac{1}{3} + \frac{1}{2}t + \frac{1}{6}t^3$
3. $G_4(t) = \frac{3}{8} + \frac{1}{3}t + \frac{1}{4}t^2 + \frac{1}{24}t^4$

For $k, n \in \mathbb{N}_+$ with $k < n$,

$$G_n^{(k)}(t) = G_{n-k}(t), \quad t \in \mathbb{R} \quad (12.5.14)$$

Proof

This follows from [differential equation](#) for the PGF given above.

For $n \in \mathbb{N}_+$,

$$\mathbb{P}(N_n = k) = \frac{1}{k!} \mathbb{P}(N_{n-k} = 0), \quad k \in \{0, 1, \dots, n-1\} \quad (12.5.15)$$

Proof

This follows from the [previous result](#) and basic properties of generating functions.

Examples and Applications

A secretary randomly stuffs 5 letters into 5 envelopes. Find each of the following:

1. The number of outcomes with exactly k matches, for each $k \in \{0, 1, 2, 3, 4, 5\}$
2. The probability density function of the number of matches.
3. The covariance and correlation of a match in one envelope and a match in another envelope.

Answer

1.	k	0	1	2	3	4	5
	$b_5(k)$	44	45	20	10	0	1

2.	k	0	1	2	3	4	5

$\mathbb{P}(N_5 = k)$	0.3667	0.3750	0.1667	0.0833	0	0.0083
-----------------------	--------	--------	--------	--------	---	--------

3. Covariance: $\frac{1}{100}$, correlation $\frac{1}{16}$

Ten married couples are randomly paired for a dance. Find each of the following:

1. The probability density function of the number of matches.
2. The mean and variance of the number of matches.
3. The probability of at least 3 matches.

Answer

1. k	$\mathbb{P}(N_{10} = k)$
0	$\frac{16\,481}{44\,800} \approx 0.3678795$
1	$\frac{16\,687}{45\,360} \approx 0.3678792$
2	$\frac{2119}{11\,520} \approx 0.1839410$
3	$\frac{103}{1680} \approx 0.06130952$
4	$\frac{53}{3456} \approx 0.01533565$
5	$\frac{11}{3600} \approx 0.003055556$
6	$\frac{1}{1920} \approx 0.0005208333$
7	$\frac{1}{15\,120} \approx 0.00006613757$
8	$\frac{1}{80\,640} \approx 0.00001240079$
9	0
10	$\frac{1}{3\,628\,800} \approx 2.755732 \times 10^{-7}$

2. $\mathbb{E}(N_{10}) = 1$, $\text{var}(N_{10}) = 1$

3. $\mathbb{P}(N_{10} \geq 3) = \frac{145\,697}{1\,814\,400} \approx 0.08030037$

In the matching experiment, set $n = 10$. Run the experiment 1000 times and compare the following for the number of matches:

1. The true probabilities
2. The relative frequencies from the simulation
3. The limiting Poisson probabilities

Answer

1. See part (a) of the [previous problem](#).

3. k	$e^{-1} \frac{1}{k!}$
0	0.3678794
1	0.3678794
2	0.1839397
3	0.06131324
4	0.01532831
5	0.003065662
6	0.0005109437
7	0.00007299195

8	9.123994×10^{-6}
9	1.013777×10^{-6}
10	1.013777×10^{-7}

This page titled [12.5: The Matching Problem](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.6: The Birthday Problem

Introduction

The Sampling Model

As in the basic sampling model, suppose that we select n numbers at random, *with* replacement, from the population $D = \{1, 2, \dots, m\}$. Thus, our outcome vector is $\mathbf{X} = (X_1, X_2, \dots, X_n)$ where X_i is the i th number chosen. Recall that our basic modeling assumption is that \mathbf{X} is uniformly distributed on the sample space $S = D^n = \{1, 2, \dots, m\}^n$

In this section, we are interested in the number of population values missing from the sample, and the number of (distinct) population values in the sample. The computation of probabilities related to these random variables are generally referred to as *birthday problems*. Often, we will interpret the sampling experiment as a distribution of n balls into m cells; X_i is the cell number of ball i . In this interpretation, our interest is in the number of empty cells and the number of occupied cells.

For $i \in D$, let Y_i denote the number of times that i occurs in the sample:

$$Y_i = \# \{j \in \{1, 2, \dots, n\} : X_j = i\} = \sum_{j=1}^n \mathbf{1}(X_j = i) \quad (12.6.1)$$

$\mathbf{Y} = (Y_1, Y_2, \dots, Y_m)$ has the multinomial distribution with parameters n and $(1/m, 1/m, \dots, 1/m)$

$$\mathbb{P}(Y_1 = y_1, Y_2 = y_2, \dots, Y_m = y_m) = \binom{n}{y_1, y_2, \dots, y_m} \frac{1}{m^n}, \quad (y_1, y_2, \dots, y_m) \in \mathbb{N}^n \text{ with } \sum_{i=1}^m y_i = n \quad (12.6.2)$$

Proof

This follows immediately from the definition of the multinomial distribution, since (X_1, X_2, \dots, X_n) is an independent sequence, and X_i is uniformly distributed on $\{1, 2, \dots, m\}$ for each i .

We will now define the main random variables of interest.

The number of population values missing in the sample is

$$U = \# \{j \in \{1, 2, \dots, m\} : Y_j = 0\} = \sum_{j=1}^m \mathbf{1}(Y_j = 0) \quad (12.6.3)$$

and the number of (distinct) population values that occur in the sample is

$$V = \# \{j \in \{1, 2, \dots, m\} : Y_j > 0\} = \sum_{j=1}^m \mathbf{1}(Y_j > 0) \quad (12.6.4)$$

Also, U takes values in $\{\max\{m - n, 0\}, \dots, m - 1\}$ and V takes values in $\{1, 2, \dots, \min\{m, n\}\}$

Clearly we must have $U + V = m$ so once we have the probability distribution and moments of one variable, we can easily find them for the other variable. However, we will first solve the simplest version of the birthday problem.

The Simple Birthday Problem

The event that there is at least one duplication when a sample of size n is chosen from a population of size m is

$$B_{m,n} = \{V < n\} = \{U > m - n\} \quad (12.6.5)$$

The (simple) *birthday problem* is to compute the probability of this event. For example, suppose that we choose n people at random and note their birthdays. If we ignore leap years and assume that birthdays are uniformly distributed throughout the year, then our sampling model applies with $m = 365$. In this setting, the birthday problem is to compute the probability that at least two people have the same birthday (this special case is the origin of the name).

The solution of the birthday problem is an easy exercise in combinatorial probability.

The probability of the birthday event is

$$\mathbb{P}(B_{m,n}) = 1 - \frac{m^{(n)}}{m^n}, \quad n \leq m \quad (12.6.6)$$

and $\mathbb{P}(B_{m,n}) = 1$ for $n > m$

Proof

The complementary event B^c occurs if and only if the outcome vector \mathbf{X} forms a permutation of size n from $\{1, 2, \dots, m\}$. The number of permutations is $m^{(n)}$ and of course the number of samples is m^n .

The fact that the probability is 1 for $n > m$ is sometimes referred to as the *pigeonhole principle*: if more than m pigeons are placed into m holes then at least one hole has 2 or more pigeons. The following result gives a recurrence relation for the probability of distinct sample values and thus gives another way to compute the birthday probability.

Let $p_{m,n}$ denote the probability of the complementary birthday event B^c , that the sample variables are distinct, with population size m and sample size n . Then $p_{m,n}$ satisfies the following recursion relation and initial condition:

1. $p_{m,n+1} = \frac{m-n}{m} p_{m,n}$
2. $p_{m,1} = 1$

Examples

Let $m = 365$ (the standard birthday problem).

1. $\mathbb{P}(B_{365,10}) = 0.117$
2. $\mathbb{P}(B_{365,20}) = 0.411$
3. $\mathbb{P}(B_{365,30}) = 0.706$
4. $\mathbb{P}(B_{365,40}) = 0.891$
5. $\mathbb{P}(B_{365,50}) = 0.970$
6. $\mathbb{P}(B_{365,60}) = 0.994$




Figure 12.6.1: $\mathbb{P}(B_{365,n})$ as a function of n , smoothed for the sake of appearance

In the birthday experiment, set $n = 365$ and select the indicator variable I . For $n \in \{10, 20, 30, 40, 50, 60\}$ run the experiment 1000 times each and compare the relative frequencies with the true probabilities.

In spite of its easy solution, the birthday problem is famous because, numerically, the probabilities can be a bit surprising. Note that with a just 60 people, the event is almost certain! With just 23 people, the birthday event is about $\frac{1}{2}$; specifically $\mathbb{P}(B_{365,23}) = 0.507$. Mathematically, the rapid increase in the birthday probability, as n increases, is due to the fact that m^n grows much faster than $m^{(n)}$.

Four fair, standard dice are rolled. Find the probability that the scores are distinct.

Answer

$$\frac{5}{18}$$

In the birthday experiment, set $m = 6$ and select the indicator variable I . Vary n with the scrollbar and note graphically how the probabilities change. Now with $n = 4$, run the experiment 1000 times and compare the relative frequency of the event to the corresponding probability.

Five persons are chosen at random.

1. Find the probability that at least 2 have the same birth *month*.
2. Criticize the sampling model in this setting

Answer

1. $\frac{89}{144}$
2. The number of days in a month varies, so the assumption that a person's birth month is uniformly distributed over the 12 months not quite accurate.

In the birthday experiment, set $m = 12$ and select the indicator variable I . Vary n with the scrollbar and note graphically how the probabilities change. Now with $n = 5$, run the experiment 1000 times and compare the relative frequency of the event to the corresponding probability.

A fast-food restaurant gives away one of 10 different toys with the purchase of a kid's meal. A family with 5 children buys 5 kid's meals. Find the probability that the 5 toys are different.

Answer

$$\frac{189}{625}$$

In the birthday experiment, set $m = 10$ and select the indicator variable I . Vary n with the scrollbar and note graphically how the probabilities change. Now with $n = 5$, run the experiment 1000 times and compare the relative frequency of the event to the corresponding probability.

Let $m = 52$. Find the smallest value of n such that the probability of a duplication is at least $\frac{1}{2}$.

Answer

$$n = 9$$

The General Birthday Problem

We now return to the more general problem of finding the distribution of the number of distinct sample values and the distribution of the number of excluded sample values.

The Probability Density Function

The number of samples with exactly j values excluded is

$$\#\{U = j\} = \binom{m}{j} \sum_{k=0}^{m-j} (-1)^k \binom{m-j}{k} (m-j-k)^n, \quad j \in \{\max\{m-n, 0\}, \dots, m-1\} \quad (12.6.7)$$

Proof

For $i \in D$, consider the event that i does not occur in the sample: $A_i = \{Y_i = 0\}$. Now let $K \subseteq D$ with $\#(K) = k$. Using the multiplication rule of combinatorics, it is easy to count the number of samples that do not contain any elements of K :

$$\#\left(\bigcap_{i \in K} A_i\right) = (m-k)^n \quad (12.6.8)$$

Now the inclusion-exclusion rule of combinatorics can be used to count the number samples that are missing at least one population value:

$$\# \left(\bigcup_{i=1}^m A_i \right) = \sum_{k=1}^m (-1)^{k-1} \binom{m}{k} (m-k)^n \quad (12.6.9)$$

Once we have this, we can use DeMorgan's law to count the number samples that contain all population values:

$$\# \left(\bigcap_{i=1}^m A_i^c \right) = \sum_{k=0}^m (-1)^k \binom{m}{k} (m-k)^n \quad (12.6.10)$$

Now we can use a two-step procedure to generate all samples that exclude exactly j population values: First, choose the j values that are to be excluded. The number of ways to perform this step is $\binom{m}{j}$. Next select a sample of size n from the remaining population values so that none are excluded. The number of ways to perform this step is the result in the last displayed equation, but with $m-j$ replacing m . The multiplication principle of combinatorics gives the result.

The distributions of the number of excluded values and the number of distinct values are now easy.

The probability density function of U is given by

$$\mathbb{P}(U=j) = \binom{m}{j} \sum_{k=0}^{m-j} (-1)^k \binom{m-j}{k} \left(1 - \frac{j+k}{m}\right)^n, \quad j \in \{\max\{m-n, 0\}, \dots, m-1\} \quad (12.6.11)$$

Proof

Since the samples are uniformly distributed, $\mathbb{P}(U=j) = \#\{U=j\}/m^n$ and so the result follows from the previous exercise.

The probability density function of the number of distinct values V is given by

$$\mathbb{P}(V=j) = \binom{m}{j} \sum_{k=0}^j (-1)^k \binom{j}{k} \left(\frac{j-k}{m}\right)^n, \quad j \in \{1, 2, \dots, \min\{m, n\}\} \quad (12.6.12)$$

Proof

This follows from the previous theorem since $\mathbb{P}(V=j) = \mathbb{P}(U=m-j)$.

In the birthday experiment, select the number of distinct sample values. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 and compare the relative frequency function to the probability density function.

The distribution of the number of excluded values can also be obtained by a recursion argument.

Let $f_{m,n}$ denote the probability density function of the number of excluded values U , when the population size is m and the sample size is n . Then

1. $f_{m,1}(m-1) = 1$
2. $f_{m,n+1}(j) = \frac{m-j}{m} f_{m,n}(j) + \frac{j+1}{m} f_{m,n}(j+1)$

Moments

Now we will find the means and variances. The number of excluded values and the number of distinct values are counting variables and hence can be written as sums of indicator variables. As we have seen in many other models, this representation is frequently the best for computing moments.

For $j \in \{0, 1, \dots, m\}$, let $I_j = \mathbf{1}(Y_j = 0)$, the indicator variable of the event that j is not in the sample. Note that the number of population values missing in the sample can be written as the sum of the indicator variables:

$$U = \sum_{j=1}^m I_j \quad (12.6.13)$$

For distinct $i, j \in \{1, 2, \dots, m\}$,

1. $E(I_j) = \left(1 - \frac{1}{m}\right)^n$
2. $\text{var}(I_j) = \left(1 - \frac{1}{m}\right)^n - \left(1 - \frac{1}{m}\right)^{2n}$
3. $\text{cov}(I_i, I_j) = \left(1 - \frac{2}{m}\right)^n - \left(1 - \frac{1}{m}\right)^{2n}$

Proof

Since each population value is equally likely to be chosen, $\mathbb{P}(I_j = 1) = (1 - 1/m)^n$. Thus, parts (a) and (b) follow from standard results for the mean and variance of an indicator variable. Next, $I_i I_j$ is the indicator variable of the event that i and j are both excluded, so $\mathbb{P}(I_i I_j = 1) = (1 - 2/m)^n$. Part (c) then follows from the standard formula for covariance.

The expected number of excluded values and the expected number of distinct values are

1. $\mathbb{E}(U) = m \left(1 - \frac{1}{m}\right)^n$
2. $\mathbb{E}(V) = m \left[1 - \left(1 - \frac{1}{m}\right)^n\right]$

Proof

Part (a) follows from the previous exercise and the representation $U = \sum_{j=1}^n I_j$. Part (b) follows from part (a) since $U + V = m$.

The variance of the number of excluded values and the variance of the number of distinct values are

$$\text{var}(U) = \text{var}(V) = m(m-1) \left(1 - \frac{2}{m}\right)^n + m \left(1 - \frac{1}{m}\right)^n - m^2 \left(1 - \frac{1}{m}\right)^{2n} \quad (12.6.14)$$

Proof

Recall that $\text{var}(U) = \sum_{i=1}^m \sum_{j=1}^m \text{cov}(I_i, I_j)$. Using the results above on the [covariance](#) of the indicator variables and simplifying gives the variance of U . Also, $\text{var}(V) = \text{var}(U)$ since $U + V = m$.

In the birthday experiment, select the number of distinct sample values. Vary the parameters and note the size and location of the mean \pm standard-deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the sample mean and variance to the distribution mean and variance.

Examples and Applications

Suppose that 30 persons are chosen at random. Find each of the following:

1. The probability density function of the number of distinct birthdays.
2. The mean of the number of distinct birthdays.
3. The variance of the number of distinct birthdays.
4. The probability that there are at least 28 different birthdays represented.

Answer

1. $\mathbb{P}(V = j) = \binom{30}{j} \sum_{k=0}^j (-1)^k \left(\frac{j-k}{365}\right)^{30}, \quad j \in \{1, 2, \dots, 30\}$
2. $\mathbb{E}(V) = 28.8381$
3. $\text{var}(V) = 1.0458$
4. $\mathbb{P}(V \geq 28) = 0.89767$

In the birthday experiment, set $m = 365$ and $n = 30$. Run the experiment 1000 times with an update frequency of 10 and compute the relative frequency of the event in part (d) of the last exercise.

Suppose that 10 fair dice are rolled. Find each of the following:

1. The probability density function of the number of distinct scores.
2. The mean of the number of distinct scores.
3. The variance of the number of distinct scores.

4. The probability that there will 4 or fewer distinct scores.

Answer

1. $\mathbb{P}(V = j) = \binom{10}{j} \sum_{k=0}^j (-1)^k \binom{j}{k} \left(\frac{j-k}{6}\right)^{10}, \quad j \in \{1, 2, \dots, 6\}$
2. $\mathbb{E}(V) = 5.0310$
3. $\text{var}(V) = 0.5503$
4. $\mathbb{P}(V \leq 4) = 0.22182$

In the birthday experiment, set $m = 6$ and $n = 10$. Run the experiment 1000 times and compute the relative frequency of the event in part (d) of the last exercise.

A fast food restaurant gives away one of 10 different toys with the purchase of each kid's meal. A family buys 15 kid's meals. Find each of the following:

1. The probability density function of the number of toys that are missing.
2. The mean of the number of toys that are missing.
3. The variance of the number of toys that are missing.
4. The probability that at least 3 toys are missing.

Answer

1. $\mathbb{P}(U = j) = \binom{15}{j} \sum_{k=0}^{10-j} (-1)^k \binom{10-j}{k} \left(1 - \frac{j+k}{10}\right)^{15}, \quad j \in \{0, 1, \dots, 9\}$
2. $\mathbb{E}(U) = 2.0589$
3. $\text{var}(U) = 0.9864$
4. $\mathbb{P}(U \geq 3) = 0.3174$

In the birthday experiment, set $m = 10$ and $n = 15$. Run the experiment 1000 times and compute the relative frequency of the event in part (d).

The lying students problem. Suppose that 3 students, who ride together, miss a mathematics exam. They decide to lie to the instructor by saying that the car had a flat tire. The instructor separates the students and asks each of them which tire was flat. The students, who did not anticipate this, select their answers independently and at random. Find each of the following:

1. The probability density function of the number of distinct answers.
2. The probability that the students get away with their deception.
3. The mean of the number of distinct answers.
4. The standard deviation of the number of distinct answers.

Answer

1. j	1	2	3
$\mathbb{P}(V = j)$	$\frac{1}{16}$	$\frac{9}{16}$	$\frac{6}{16}$

2. $\mathbb{P}(V = 1) = \frac{1}{16}$
3. $\mathbb{E}(V) = \frac{37}{16}$
4. $\text{sd}(V) = \sqrt{\frac{87}{256}} \approx 0.58296$

The duck hunter problem. Suppose that there are 5 duck hunters, each a perfect shot. A flock of 10 ducks fly over, and each hunter selects one duck at random and shoots. Find each of the following:

1. The probability density function of the number of ducks that are killed.
2. The mean of the number of ducks that are killed.
3. The standard deviation of the number of ducks that are killed.

Answer

1.	j	1	2	3	4	5
	$\mathbb{P}(V = j)$	$\frac{1}{10\,000}$	$\frac{27}{2000}$	$\frac{9}{50}$	$\frac{63}{125}$	$\frac{189}{625}$

2. $\mathbb{E}(V) = \frac{40\,951}{10\,000} = 4.0951$

3. $\text{sd}(V) = 0.72768$

This page titled [12.6: The Birthday Problem](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.7: The Coupon Collector Problem

Basic Theory

Definitions

In this section, our random experiment is to sample repeatedly, *with replacement*, from the population $D = \{1, 2, \dots, m\}$. This generates a sequence of independent random variables $\mathbf{X} = (X_1, X_2, \dots)$, each uniformly distributed on D .

We will often interpret the sampling in terms of a *coupon collector*: each time the collector buys a certain product (bubble gum or Cracker Jack, for example) she receives a coupon (a baseball card or a toy, for example) which is equally likely to be any one of m types. Thus, in this setting, $X_i \in D$ is the coupon type received on the i th purchase.

Let V_n denote the number of distinct values in the first n selections, for $n \in \mathbb{N}_+$. This is the random variable studied in the last section on the Birthday Problem. Our interest in this section is the sample size needed to get a specified number of distinct sample values

For $k \in \{1, 2, \dots, m\}$, let

$$W_k = \min\{n \in \mathbb{N}_+ : V_n = k\} \quad (12.7.1)$$

the sample size needed to get k distinct sample values.

In terms of the coupon collector, this random variable gives the number of products required to get k distinct coupon types. Note that the set of possible values of W_k is $\{k, k+1, \dots\}$. We will be particularly interested in W_m , the sample size needed to get the entire population. In terms of the coupon collector, this is the number of products required to get the entire set of coupons.

In the coupon collector experiment, run the experiment in single-step mode a few times for selected values of the parameters.

The Probability Density Function

Now let's find the distribution of W_k . The results of the previous section will be very helpful

For $k \in \{1, 2, \dots, m\}$, the probability density function of W_k is given by

$$\mathbb{P}(W_k = n) = \binom{m-1}{k-1} \sum_{j=0}^{k-1} (-1)^j \binom{k-1}{j} \left(\frac{k-j-1}{m}\right)^{n-1}, \quad n \in \{k, k+1, \dots\} \quad (12.7.2)$$

Proof

Note first that $W_k = n$ if and only if $V_{n-1} = k-1$ and $V_n = k$. Hence

$$\mathbb{P}(W_k = n) = \mathbb{P}(V_{n-1} = k-1) \mathbb{P}(V_n = k \mid V_{n-1} = k-1) = \frac{m-k+1}{m} \mathbb{P}(V_{n-1} = k-1) \quad (12.7.3)$$

Using the PDF of V_{n-1} from the previous section gives the result.

In the coupon collector experiment, vary the parameters and note the shape of and position of the probability density function. For selected values of the parameters, run the experiment 1000 times and compare the relative frequency function to the probability density function.

An alternate approach to the probability density function of W_k is via a recursion formula.

For fixed m , let g_k denote the probability density function of W_k . Then

1. $g_k(n+1) = \frac{k-1}{m} g_k(n) + \frac{m-k+1}{m} g_{k-1}(n)$
2. $g_1(1) = 1$

Decomposition as a Sum

We will now show that W_k can be decomposed as a sum of k independent, geometrically distributed random variables. This will provide some additional insight into the nature of the distribution and will make the computation of the mean and variance easy.

For $i \in \{1, 2, \dots, m\}$, let Z_i denote the number of additional samples needed to go from $i - 1$ distinct values to i distinct values. Then $\mathbf{Z} = (Z_1, Z_2, \dots, Z_m)$ is a sequence of independent random variables, and Z_i has the geometric distribution on \mathbb{N}_+ with parameter $p_i = \frac{m-i+1}{m}$. Moreover,

$$W_k = \sum_{i=1}^k Z_i, \quad k \in \{1, 2, \dots, m\} \quad (12.7.4)$$

This result shows clearly that each time a new coupon is obtained, it becomes harder to get the next new coupon.

In the coupon collector experiment, run the experiment in single-step mode a few times for selected values of the parameters. In particular, try this with m large and k near m .

Moments

The decomposition as a sum of independent variables provides an easy way to compute the mean and other moments of W_k .

The mean and variance of the sample size needed to get k distinct values are

1. $\mathbb{E}(W_k) = \sum_{i=1}^k \frac{m}{m-i+1}$
2. $\text{var}(W_k) = \sum_{i=1}^k \frac{(i-1)m}{(m-i+1)^2}$

Proof

These results follow from the decomposition of W_k as a [sum of independent variables](#) and standard results for the geometric distribution, since $\mathbb{E}(W_k) = \sum_{i=1}^k \mathbb{E}(Z_i)$ and $\text{var}(W_k) = \sum_{i=1}^k \text{var}(Z_i)$.

In the coupon collector experiment, vary the parameters and note the shape and location of the mean \pm standard deviation bar. For selected values of the parameters, run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation.

The probability generating function of W_k is given by

$$\mathbb{E}(t^{W_k}) = \prod_{i=1}^k \frac{m-i+1}{m-(i-1)t}, \quad |t| < \frac{m}{k-1} \quad (12.7.5)$$

Proof

This follows from the decomposition of W_k as a [sum of independent variables](#) and standard results for the geometric distribution on \mathbb{N}_+ , since $\mathbb{E}(t^{W_k}) = \prod_{i=1}^k \mathbb{E}(t^{Z_i})$.

Examples and Applications

Suppose that people are sampled at random until 40 distinct birthdays are obtained. Find each of the following:

1. The probability density function of the sample size.
2. The mean of the sample size.
3. The variance of the sample size.
4. The probability generating function of the sample size.

Answer

Let W denote the sample size.

1. $\mathbb{P}(W = n) = \binom{364}{n} \sum_{j=0}^{30} (-1)^j \binom{39}{j} \left(\frac{39-j}{365}\right)^{n-1} \quad \text{for } n \in \{40, 41, \dots\}$

2. $\mathbb{E}(W) = 42.3049$
3. $\text{var}(W) = 2.4878$
4. $\mathbb{E}(t^W) = \prod_{i=1}^{40} \frac{366-i}{365-(i-1)t}$ for $|t| < \frac{365}{39}$

Suppose that a standard, fair die is thrown until all 6 scores have occurred. Find each of the following:

1. The probability density function of the number of throws.
2. The mean of the number of throws.
3. The variance of the number of throws.
4. The probability that at least 10 throws are required.

Answer

Let W denote the number of throws.

1. $\mathbb{P}(W = n) = \sum_{j=0}^5 (-1)^j \binom{5}{j} \left(\frac{5-j}{6}\right)^{n-1}$ for $n \in \{6, 7, \dots\}$
2. $\mathbb{E}(W) = 14.7$
3. $\text{var}(W) = 38.99$
4. $\mathbb{P}(W \geq 10) = \frac{1051}{1296} \approx 0.81096$

A box of a certain brand of cereal comes with a special toy. There are 10 different toys in all. A collector buys boxes of cereal until she has all 10 toys. Find each of the following:

1. The probability density function of the number boxes purchased.
2. The mean of the number of boxes purchased.
3. The variance of the number of boxes purchased.
4. The probability that no more than 15 boxes were purchased.

Answer

Let W denote the number of boxes purchased.

1. $\mathbb{P}(W = n) = \sum_{j=0}^9 (-1)^j \binom{9}{j} \left(\frac{9-j}{10}\right)^{n-1}$, for $n \in \{10, 11, \dots\}$
2. $\mathbb{E}(W) = 29.2897$
3. $\text{var}(W) = 125.6871$
4. $\mathbb{P}(W \leq 15) = 0.04595$

This page titled [12.7: The Coupon Collector Problem](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.8: Pólya's Urn Process

Basic Theory

The Model

Pólya's urn scheme is a dichotomous sampling model that generalizes the hypergeometric model (sampling without replacement) and the Bernoulli model (sampling with replacement). Pólya's urn process leads to a famous example of a sequence of random variables that is exchangeable, but not independent, and has deep connections with the beta-Bernoulli process.

Suppose that we have an urn (what else!) that initially contains a red and b green balls, where a and b are positive integers. At each discrete time (trial), we select a ball from the urn and then return the ball to the urn along with c new balls of the same color. Ordinarily, the parameter c is a nonnegative integer. However, the model actually makes sense if c is a negative integer, if we interpret this to mean that we *remove* the balls rather than add them, and assuming that there are enough balls of the proper color in the urn to perform this action. In any case, the random process is known as *Pólya's urn process*, named for George Pólya.

In terms of the colors of the selected balls, Pólya's urn scheme generalizes the standard models of sampling with and without replacement.

1. $c = 0$ corresponds to sampling with replacement.
2. $c = -1$ corresponds to sampling without replacement.

For the most part, we will assume that c is nonnegative so that the process can be continued indefinitely. Occasionally we consider the case $c = -1$ so that we can interpret the results in terms of sampling without replacement.

The Outcome Variables

Let X_i denote the color of the ball selected at time i , where 0 denotes green and 1 denotes red. Mathematically, our basic random process is the sequence of indicator variables $\mathbf{X} = (X_1, X_2, \dots)$, known as the *Pólya process*. As with any random process, our first goal is to compute the *finite dimensional distributions* of \mathbf{X} . That is, we want to compute the joint distribution of (X_1, X_2, \dots, X_n) for each $n \in \mathbb{N}_+$. Some additional notation will really help. Recall the generalized permutation formula in our study of combinatorial structures: for $r, s \in \mathbb{R}$ and $j \in \mathbb{N}$, we defined

$$r^{(s,j)} = r(r+s)(r+2s) \cdots [r+(j-1)s] \quad (12.8.1)$$

Note that the expression has j factors, starting with r , and with each factor obtained by adding s to the previous factor. As usual, we adopt the convention that a product over an empty index set is 1. Hence $r^{(s,0)} = 1$ for every r and s .

Recall that

1. $r^{(0,j)} = r^j$, an ordinary power
2. $r^{(-1,j)} = r^{(j)} = r(r-1) \cdots (r-j+1)$, a descending power
3. $r^{(1,j)} = r^{[j]} = r(r+1) \cdots (r+j-1)$, an ascending power
4. $r^{(r,j)} = j!r^j$
5. $1^{(1,j)} = j!$

The following simple result will turn out to be quite useful.

Suppose that $r, s \in (0, \infty)$ and $j \in \mathbb{N}$. Then

$$\frac{r^{(s,j)}}{s^j} = \left(\frac{r}{s}\right)^{[j]} \quad (12.8.2)$$

Proof

It's just a matter of grouping the factors:

$$\begin{aligned} \frac{r^{(s,j)}}{s^j} &= \frac{r(r+s)(r+2s) \cdots [r+(j-1)s]}{s^j} \\ &= \left(\frac{r}{s}\right) \left(\frac{r}{s} + 1\right) \left(\frac{r}{s} + 2\right) \cdots \left[\frac{r}{s} + (j-1)\right] = \left(\frac{r}{s}\right)^{[j]} \end{aligned}$$

The finite dimensional distributions are easy to compute using the multiplication rule of conditional probability. If we know the contents of the urn at any given time, then the probability of an outcome at the next time is all but trivial.

Let $n \in \mathbb{N}_+$, $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ and let $k = x_1 + x_2 + \cdots + x_n$. Then

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a^{(c,k)} b^{(c,n-k)}}{(a+b)^{(c,n)}} \quad (12.8.3)$$

Proof

By the multiplication rule for conditional probability,

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \mathbb{P}(X_1 = x_1) \mathbb{P}(X_2 = x_2 | X_1 = x_1) \cdots \mathbb{P}(X_n = x_n | X_1 = x_1, \dots, X_{n-1} = x_{n-1}) \quad (12.8.4)$$

Of course, if we know that the urn has, say, r red and g green balls at a particular time, then the probability of a red ball on the next draw is $r/(r+g)$ while the probability of a green ball is $g/(r+g)$. The right side of the displayed equation above has n factors. The denominators are the total number of balls at the n times, and form the product $(a+b)(a+b+c) \cdots [a+b+(n-1)c] = (a+b)^{(c,n)}$. In the numerators, k of the factors correspond to probabilities of selecting red balls; these factors form the product $a(a+c) \cdots [a+(k-1)c] = a^{(c,k)}$. The remaining $n-k$ factors in the numerators correspond to selecting green balls; these factors form the product $b(b+c) \cdots [b+(n-k-1)c] = b^{(c,n-k)}$.

The joint probability in the previous exercise depends on (x_1, x_2, \dots, x_n) only through the number of red balls $k = \sum_{i=1}^n x_i$ in the sample. Thus, the joint distribution is invariant under a permutation of (x_1, x_2, \dots, x_n) , and hence \mathbf{X} is an exchangeable sequence of random variables. This means that for each n , all permutations of (X_1, X_2, \dots, X_n) have the same distribution. Of course the joint distribution reduces to the formulas we have obtained earlier in the special cases of sampling with replacement ($c = 0$) or sampling without replacement ($c = -1$), although in the latter case we must have $n \leq a+b$. When $c > 0$, the Pólya process is a special case of the beta-Bernoulli process, studied in the chapter on Bernoulli trials.

The Pólya process $\mathbf{X} = (X_1, X_2, \dots)$ with parameters $a, b, c \in \mathbb{N}_+$ is the beta-Bernoulli process with parameters a/c and b/c . That is, for $n \in \mathbb{N}_+$, $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, and with $k = x_1 + x_2 + \cdots + x_n$,

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{(a/c)^{[k]} (b/c)^{[n-k]}}{(a/c + b/c)^{[n]}} \quad (12.8.5)$$

Proof

From the previous two results,

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a^{(c,k)} b^{(c,n-k)}}{(a+b)^{(c,n)}} = \frac{[a^{(c,k)/c}] [b^{(c,n-k)/c} c^{n-k}]}{(a+b)^{(c,n)}/c^n} = \frac{(a/c)^{[k]} (b/c)^{[n-k]}}{(a/c + b/c)^{[n]}} \quad (12.8.6)$$

and this is the corresponding finite dimensional distribution of the beta-Bernoulli distribution with parameters a/c and b/c .

Recall that the beta-Bernoulli process is obtained, in the usual formulation, by randomizing the success parameter in a Bernoulli trials sequence, giving the success parameter a beta distribution. So specifically, suppose $a, b, c \in \mathbb{N}_+$ and that random variable P has the beta distribution with parameters a/c and b/c . Suppose also that given $P = p \in (0, 1)$, the random process $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter p . Then \mathbf{X} is the Pólya process with parameters a, b, c . This is a fascinating connection between two processes that at first, seem to have little in common. In fact however, every exchangeable sequence of indicator random variables can be obtained by randomizing the success parameter in a sequence of Bernoulli trials. This is *de Finetti's theorem*, named for Bruno de Finetti, which is studied in the section on backwards martingales. When $c \in \mathbb{N}_+$, all of the results in this section are special cases of the corresponding results for the beta-Bernoulli process, but it's still interesting to interpret the results in terms of the urn model.

For each $i \in \mathbb{N}_+$

1. $\mathbb{E}(X_i) = \frac{a}{a+b}$
2. $\text{var}(X_i) = \frac{a}{a+b} \frac{b}{a+b}$

Proof

Since the sequence is exchangeable, X_i has the same distribution as X_1 , so $\mathbb{P}(X_i = 1) = \frac{a}{a+b}$. The mean and variance now follow from standard results for indicator variables.

Thus \mathbf{X} is a sequence of identically distributed variables, quite surprising at first but of course inevitable for any exchangeable sequence. Compare the joint and marginal distributions. Note that \mathbf{X} is an independent sequence if and only if $c = 0$, when we have simple sampling with replacement. Pólya's urn is one of the most famous examples of a random process in which the outcome variables are exchangeable, but dependent (in general).

Next, let's compute the covariance and correlation of a pair of outcome variables.

Suppose that $i, j \in \mathbb{N}_+$ are distinct. Then

1. $\text{cov}(X_i, X_j) = \frac{abc}{(a+b)^2(a+b+c)}$
2. $\text{cor}(X_i, X_j) = \frac{c}{a+b+c}$

Proof

Since the variables are exchangeable, $\mathbb{P}(X_i = 1, X_j = 1) = \mathbb{P}(X_1 = 1, X_2 = 1) = \frac{a}{a+b} \frac{a+c}{a+b+c}$. The results now follow from standard formulas for covariance and correlation.

Thus, the variables are positively correlated if $c > 0$, negatively correlated if $c < 0$, and uncorrelated (in fact, independent), if $c = 0$. These results certainly make sense when we recall the dynamics of Pólya's urn. It turns out that in any *infinite* sequence of exchangeable variables, the variables must be nonnegatively correlated. Here is another result that explores how the variables are related.

Suppose that $n \in \mathbb{N}_+$ and $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$. Let $k = \sum_{i=1}^n x_i$. Then

$$\mathbb{P}(X_{n+1} = 1 \mid X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a + kc}{a + b + nc} \quad (12.8.7)$$

Proof

Using the joint distribution,

$$\begin{aligned} \mathbb{P}(X_{n+1} = 1 \mid X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) &= \frac{\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n, X_{n+1} = 1)}{\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)} \\ &= \frac{a^{(c,k+1)} b^{(c,n-k)}}{(a+b)^{(c,n+1)}} \frac{(a+b)^{(c,n)}}{a^{(c,k)} b^{(c,n-k)}} = \frac{a + ck}{a + b + cn} \end{aligned}$$

Pólya's urn is described by a sequence of indicator variables. We can study the same derived random processes that we studied with Bernoulli trials: the number of red balls in the first n trials, the trial number of the k th red ball, and so forth.

The Number of Red Balls

For $n \in \mathbb{N}$, the number of red balls selected in the first n trials is

$$Y_n = \sum_{i=1}^n X_i \quad (12.8.8)$$

so that $\mathbf{Y} = (Y_0, Y_1, \dots)$ is the partial sum process associated with $\mathbf{X} = (X_1, X_2, \dots)$.

Note that

1. The number of green balls selected in the first n trials is $n - Y_n$.
2. The number of red balls in the urn after the first n trials is $a + c Y_n$.
3. The number of green balls in the urn after the first n trials is $b + c(n - Y_n)$.
4. The number of balls in the urn after the first n trials is $a + b + c n$.

The basic analysis of \mathbf{Y} follows easily from our work with \mathbf{X} .

The probability density function of Y_n is given by

$$\mathbb{P}(Y_n = k) = \binom{n}{k} \frac{a^{(c,k)} b^{(c,n-k)}}{(a+b)^{(c,n)}}, \quad k \in \{0, 1, \dots, n\} \quad (12.8.9)$$

Proof

$\mathbb{P}(Y_n = y)$ is the sum of $\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)$ over all $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$ with $\sum_{i=1}^n x_i = k$. There are $\binom{n}{k}$ such sequences, and each has the probability given above.

The distribution defined by this probability density function is known, appropriately enough, as the *Pólya distribution* with parameters n, a, b , and c . Of course, the distribution reduces to the binomial distribution with parameters n and $a/(a+b)$ in the case of sampling with replacement ($c = 0$) and to the hypergeometric distribution with parameters n, a , and b in the case of sampling without replacement ($c = -1$), although again in this case we need $n \leq a + b$. When $c > 0$, the Pólya distribution is a special case of the beta-binomial distribution.

If $a, b, c \in \mathbb{N}_+$ then the Pólya distribution with parameters a, b, c is the beta-binomial distribution with parameters a/c and b/c . That is,

$$P(Y_n = k) = \binom{n}{k} \frac{(a/c)^{[k]} (b/c)^{[n-k]}}{(a/c + b/c)^{[n]}}, \quad k \in \{0, 1, \dots, n\} \quad (12.8.10)$$

Proof

This follows immediately from the result above that $\mathbf{X} = (X_1, X_2, \dots)$ is the beta-Bernoulli process with parameters a/c and b/c . So by definition, $Y_n = \sum_{i=1}^n X_i$ has the beta-binomial distribution with parameters $n, a/c$, and b/c . A direct proof is also simple using the permutation formula above:

$$\mathbb{P}(Y_n = k) = \binom{n}{k} \frac{a^{(c,k)} b^{(c,n-k)}}{(a+b)^{(c,n)}} = \binom{n}{k} \frac{[a^{(c,k)} / c^k] [b^{(c,n-k)} / c^{n-k}]}{(a+b)^{(c,n)} / c^n} = \binom{n}{k} \frac{(a/c)^{[k]} (b/c)^{[n-k]}}{(a/c + b/c)^{[n]}}, \quad k \in \{0, 1, \dots, n\} \quad (12.8.11)$$

The case where all three parameters are equal is particularly interesting.

If $a = b = c$ then Y_n is uniformly distributed on $\{0, 1, \dots, n\}$.

Proof

This follows from the previous result, since the beta-binomial distribution with parameters n , 1, and 1 reduces to the uniform distribution. Specifically, note that $1^{[k]} = k!$, $1^{[n-k]} = (n-k)!$ and $2^{[n]} = (n+1)!$. So substituting gives

$$\mathbb{P}(Y_n = k) = \frac{n!}{k!(n-k)!} \frac{k!(n-k)!}{(n+1)!} = \frac{1}{n+1}, \quad k \in \{0, 1, \dots, n\} \quad (12.8.12)$$

In general, the Pólya family of distributions has a diverse collection of shapes.

Start the simulation of the Pólya Urn Experiment. Vary the parameters and note the shape of the probability density function. In particular, note when the function is skewed, when the function is symmetric, when the function is unimodal, when the function is monotone, and when the function is U-shaped. For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

The Pólya probability density function is

1. unimodal if $a > b > c$ and $n > \frac{a-c}{b-c}$
2. unimodal if $b > a > c$ and $n > \frac{b-c}{a-c}$
3. U-shaped if $c > a > b$ and $n > \frac{c-b}{c-a}$
4. U-shaped if $c > b > a$ and $n > \frac{c-a}{c-b}$
5. increasing if $b < c < a$
6. decreasing if $a < c < b$

Proof

These results follow from solving the inequality $\mathbb{P}(Y_n = k) > \mathbb{P}(Y_n = k-1)$.

Next, let's find the mean and variance. Curiously, the mean does not depend on the parameter c .

The mean and variance of the number of red balls selected are

1. $\mathbb{E}(Y_n) = n \frac{a}{a+b}$
2. $\text{var}(Y_n) = n \frac{ab}{(a+b)^2} \left[1 + (n-1) \frac{c}{a+b+c} \right]$

Proof

These results follow from the mean and covariance of the indicator variables given above, and basic properties of expected value and variance.

1. $\mathbb{E}(Y_n) = \sum_{i=1}^n \mathbb{E}(X_i)$
2. $\text{var}(Y_n) = \sum_{i=1}^n \sum_{j=1}^n \text{cov}(X_i, X_j)$

Start the simulation of the Pólya Urn Experiment. Vary the parameters and note the shape and location of the mean \pm standard deviation bar. For various values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

Explicitly compute the probability density function, mean, and variance of Y_5 when $a = 6$, $b = 4$, and for the values of $c \in \{-1, 0, 1, 2, 3, 10\}$. Sketch the graph of the density function in each case.

Fix a , b , and n , and let $c \rightarrow \infty$. Then

1. $\mathbb{P}(Y_n = 0) \rightarrow \frac{b}{a+b}$
2. $\mathbb{P}(Y_n = n) \rightarrow \frac{a}{a+b}$
3. $\mathbb{P}(Y_n \in \{1, 2, \dots, n-1\}) \rightarrow 0$

Proof

Note that $\mathbb{P}(Y_n = 0) = \frac{b^{(c,n)}}{(a+b)^{(c,n)}}$. The numerator and denominator each have n factors. If these factors are grouped into a product of n fractions, then the first is $\frac{b}{a+b}$. The rest have the form $\frac{a+jc}{a+b+jc}$ where $j \in \{1, 2, \dots, n-1\}$. Each of these converges to 1 as $c \rightarrow \infty$. Part (b) follows by a similar argument. Part (c) follows from (a) and (b) and the complement rule.

Thus, the limiting distribution of Y_n as $c \rightarrow \infty$ is concentrated on 0 and n . The limiting probabilities are just the initial proportion of green and red balls, respectively. Interpret this result in terms of the dynamics of Pólya's urn scheme.

Our next result gives the conditional distribution of X_{n+1} given Y_n .

Suppose that $n \in \mathbb{N}$ and $k \in \{0, 1, \dots, n\}$. Then

$$\mathbb{P}(X_{n+1} = 1 \mid Y_n = k) = \frac{a + ck}{a + b + cn} \quad (12.8.13)$$

Proof

Let $S = \{(x_1, x_2, \dots, x_n) \in \{0, 1\}^n : \sum_{i=1}^n x_i = k\}$ and let $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$. Note that the events $\{\mathbf{X}_n = \mathbf{x}\}$ over $\mathbf{x} \in S$ partition the event $\{Y_n = k\}$. Conditioning on \mathbf{X}_n ,

$$\begin{aligned} \mathbb{P}(X_{n+1} = 1 \mid Y_n = k) &= \sum_{\mathbf{x} \in S} \mathbb{P}(X_{n+1} = 1 \mid Y_n = k, \mathbf{X}_n = \mathbf{x}) \mathbb{P}(\mathbf{X}_n = \mathbf{x} \mid Y_n = k) \\ &= \sum_{\mathbf{x} \in S} \mathbb{P}(X_{n+1} = 1 \mid \mathbf{X}_n = \mathbf{x}) \mathbb{P}(\mathbf{X}_n = \mathbf{x} \mid Y_n = k) \end{aligned}$$

But from our result above, $\mathbb{P}(X_{n+1} = 1 \mid \mathbf{X}_n = \mathbf{x}) = (a + ck)/(a + b + cn)$ for every $\mathbf{x} \in S$. Hence

$$\mathbb{P}(X_{n+1} = 1 \mid Y_n = k) = \frac{a + ck}{a + b + cn} \sum_{\mathbf{x} \in S} \mathbb{P}(\mathbf{X}_n = \mathbf{x} \mid Y_n = k) \quad (12.8.14)$$

The last sum is 1.

In particular, if $a = b = c$ then $\mathbb{P}(X_{n+1} = 1 \mid Y_n = n) = \frac{n+1}{n+2}$. This is *Laplace's rule of succession*, another interesting connection. The rule is named for Pierre Simon Laplace, and is studied from a different point of view in the section on Independence.

The Proportion of Red Balls

Suppose that $c \in \mathbb{N}$, so that the process continues indefinitely. For $n \in \mathbb{N}_+$, the proportion of red balls selected in the first n trials is

$$M_n = \frac{Y_n}{n} = \frac{1}{n} \sum_{i=1}^n X_i \quad (12.8.15)$$

This is an interesting variable, since a little reflection suggests that it may have a limit as $n \rightarrow \infty$. Indeed, if $c = 0$, then M_n is just the sample mean corresponding to n Bernoulli trials. Thus, by the law of large numbers, M_n converges to the success parameter $\frac{a}{a+b}$ as $n \rightarrow \infty$ with probability 1. On the other hand, the proportion of red balls in the urn after n trials is

$$Z_n = \frac{a + cY_n}{a + b + cn} \quad (12.8.16)$$

When $c = 0$, of course, $Z_n = \frac{a}{a+b}$ so that in this case, Z_n and M_n have the same limiting behavior. Note that

$$Z_n = \frac{a}{a + b + cn} + \frac{cn}{a + b + cn} M_n \quad (12.8.17)$$

Since the constant term converges to 0 as $n \rightarrow \infty$ and the coefficient of M_n converges to 1 as $n \rightarrow \infty$, it follows that the limits of M_n and Z_n as $n \rightarrow \infty$ will be the same, if the limit exists, for any mode of convergence: with probability 1, in mean, or in distribution. Here is the general result when $c > 0$.

Suppose that $a, b, c \in \mathbb{N}_+$. There exists a random variable P having the beta distribution with parameters a/c and b/c such that $M_n \rightarrow P$ and $Z_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1 and in mean square, and hence also in distribution.

Proof

As noted earlier, the urn process is equivalent to the beta-Bernoulli process with parameters a/c and b/c . We showed in that section that $M_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1 and in mean square, where P is the beta random variable used in the construction.

It turns out that the random process $\mathbf{Z} = \{Z_n = (a + cY_n)/(a + b + cn) : n \in \mathbb{N}\}$ is a martingale. The theory of martingales provides powerful tools for studying convergence in Pólya's urn process. As an interesting special case, note that if $a = b = c$ then the limiting distribution is the uniform distribution on $(0, 1)$.

The Trial Number of the k th Red Ball

Suppose again that $c \in \mathbb{N}$, so that the process continues indefinitely. For $k \in \mathbb{N}_+$ let V_k denote the trial number of the k th red ball selected. Thus

$$V_k = \min\{n \in \mathbb{N}_+ : Y_n = k\} \quad (12.8.18)$$

Note that V_k takes values in $\{k, k+1, \dots\}$. The random processes $\mathbf{V} = (V_1, V_2, \dots)$ and $\mathbf{Y} = (Y_1, Y_2, \dots)$ are inverses of each other in a sense.

For $k, n \in \mathbb{N}_+$ with $k \leq n$,

1. $V_k \leq n$ if and only if $Y_n \geq k$
2. $V_k = n$ if and only if $Y_{n-1} = k-1$ and $X_n = 1$

The probability density function of V_k is given by

$$\mathbb{P}(V_k = n) = \binom{n-1}{k-1} \frac{a^{(c,k)} b^{(c,n-k)}}{(a+b)^{(c,n)}}, \quad n \in \{k, k+1, \dots\} \quad (12.8.19)$$

Proof

We condition on Y_{n-1} . Using the PDF of Y_{n-1} and the result above,

$$\begin{aligned} \mathbb{P}(V_k = n) &= \mathbb{P}(Y_{n-1} = k-1, X_n = 1) = \mathbb{P}(Y_{n-1} = k-1) \mathbb{P}(X_n = 1 \mid Y_{n-1} = k-1) \\ &= \binom{n-1}{k-1} \frac{a^{(c,k-1)} b^{(c,(n-1)-(k-1))}}{(a+b)^{(c,n-1)}} \frac{a+c(k-1)}{a+b+c(n-1)} = \binom{n-1}{k-1} \frac{a^{(c,k)} b^{(c,n-k)}}{(a+b)^{(c,n)}} \end{aligned}$$

Of course this probability density function reduces to the negative binomial density function with trial parameter k and success parameter $p = \frac{a}{a+b}$ when $c = 0$ (sampling with replacement). When $c > 0$, the distribution is a special case of the beta-negative binomial distribution.

If $a, b, c \in \mathbb{N}_+$ then V_k has the beta-negative binomial distribution with parameters $k, a/c$, and b/c . That is,

$$\mathbb{P}(V_k = n) = \binom{n-1}{k-1} \frac{(a/c)^{[k]} (b/c)^{[n-k]}}{(a/c + b/c)^{[n]}}, \quad n \in \{k, k+1, \dots\} \quad (12.8.20)$$

Proof

As with previous proofs, this result follows since the underlying process $\mathbf{X} = (X_1, X_2, \dots)$ is the beta-Bernoulli process with parameters a/c and b/c . The form of the PDF also follows easily from the previous result by dividing the numerator and denominator c^n .

If $a = b = c$ then

$$\mathbb{P}(V_k = n) = \frac{k}{n(n+1)}, \quad n \in \{k, k+1, k+2, \dots\} \quad (12.8.21)$$

Proof

As in the corresponding proof for the number of red balls, the fraction in the PDF of V_k in the previous result reduces to $\frac{k!(n-k)!}{(n+1)!}$, while the binomial coefficient is $\frac{(n-1)!}{(k-1)!(n-k)!}$.

Fix a, b , and k , and let $c \rightarrow \infty$. Then

1. $\mathbb{P}(V_k = k) \rightarrow \frac{a}{a+b}$
2. $\mathbb{P}(V_k \in \{k+1, k+2, \dots\}) \rightarrow 0$

Thus, the limiting distribution of V_k is concentrated on k and ∞ . The limiting probabilities at these two points are just the initial proportion of red and green balls, respectively. Interpret this result in terms of the dynamics of Pólya's urn scheme.

This page titled [12.8: Pólya's Urn Process](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

12.9: The Secretary Problem

In this section we will study a nice problem known variously as the *secretary problem* or the *marriage problem*. It is simple to state and not difficult to solve, but the solution is interesting and a bit surprising. Also, the problem serves as a nice introduction to the general area of statistical decision making.

Statement of the Problem

As always, we must start with a clear statement of the problem.

We have n candidates (perhaps applicants for a job or possible marriage partners). The assumptions are

1. The candidates are totally ordered from best to worst with no ties.
2. The candidates arrive sequentially in random order.
3. We can only determine the *relative ranks* of the candidates as they arrive. We cannot observe the *absolute ranks*.
4. Our goal is choose the very best candidate; no one less will do.
5. Once a candidate is rejected, she is gone forever and cannot be recalled.
6. The number of candidates n is known.

The assumptions, of course, are not entirely reasonable in real applications. The last assumption, for example, that n is known, is more appropriate for the secretary interpretation than for the marriage interpretation.

What is an optimal strategy? What is the probability of success with this strategy? What happens to the strategy and the probability of success as n increases? In particular, when n is large, is there any reasonable hope of finding the best candidate?

Strategies

Play the secretary game several times with $n = 10$ candidates. See if you can find a good strategy just by trial and error.

After playing the secretary game a few times, it should be clear that the only reasonable type of strategy is to let a certain number $k - 1$ of the candidates go by, and then select the first candidate we see who is better than all of the previous candidates (if she exists). If she does not exist (that is, if no candidate better than all previous candidates appears), we will agree to accept the last candidate, even though this means failure. The parameter k must be between 1 and n ; if $k = 1$, we select the first candidate; if $k = n$, we select the last candidate; for any other value of k , the selected candidate is random, distributed on $\{k, k + 1, \dots, n\}$. We will refer to this “let $k - 1$ go by” strategy as *strategy k* .

Thus, we need to compute the probability of success $p_n(k)$ using strategy k with n candidates. Then we can maximize the probability over k to find the optimal strategy, and then take the limit over n to study the asymptotic behavior.

Analysis

First, let's do some basic computations.

For the case $n = 3$, list the 6 permutations of $\{1, 2, 3\}$ and verify the probabilities in the table below. Note that $k = 2$ is optimal.

k	1	2	3
$p_3(k)$	$\frac{2}{6}$	$\frac{3}{6}$	$\frac{2}{6}$

Answer

The following table gives the $3! = 6$ permutations of the candidates $(1, 2, 3)$, and the candidate selected by each strategy. The last row gives the total number of successes for each strategy.

Permutation	$k = 1$	$k = 2$	$k = 3$
$(1, 2, 3)$	1	3	3

(1, 3, 2)	1	2	2
(2, 1, 3)	2	1	3
(2, 3, 1)	2	1	1
(3, 1, 2)	3	1	2
(3, 2, 1)	3	2	1
Total	2	3	2

In the secretary experiment, set the number of candidates to $n = 3$. Run the experiment 1000 times with each strategy $k \in \{1, 2, 3\}$

For the case $n = 4$, list the 24 permutations of $\{1, 2, 3, 4\}$ and verify the probabilities in the table below. Note that $k = 2$ is optimal. The last row gives the total number of successes for each strategy.

k	1	2	3	4
$p_4(k)$	$\frac{6}{24}$	$\frac{11}{24}$	$\frac{10}{24}$	$\frac{6}{24}$

Answer

The following table gives the $4! = 24$ permutations of the candidates (1, 2, 3, 4), and the candidate selected by each strategy.

Permutation	$k = 1$	$k = 2$	$k = 3$	$k = 4$
(1, 2, 3, 4)	1	4	4	4
(1, 2, 4, 3)	1	3	3	3
(1, 3, 2, 4)	1	4	4	4
(1, 3, 4, 2)	1	3	2	2
(1, 4, 2, 3)	1	3	3	3
(1, 4, 3, 2)	1	2	2	2
(2, 1, 3, 4)	2	1	4	4
(2, 1, 4, 3)	2	1	3	3
(2, 3, 1, 4)	2	1	1	4
(2, 3, 4, 1)	2	1	1	1
(2, 4, 1, 3)	2	1	1	3
(2, 4, 3, 1)	2	1	1	1
(3, 1, 2, 4)	3	1	4	4
(3, 1, 4, 2)	3	1	2	2
(3, 2, 1, 4)	3	2	1	4
(3, 2, 4, 1)	3	2	1	1
(3, 4, 1, 2)	3	1	1	2
(3, 4, 2, 1)	3	2	2	1
(4, 1, 2, 3)	4	1	3	3

(4, 1, 3, 2)	4	1	2	2
(4, 2, 1, 3)	4	2	1	3
(4, 2, 3, 1)	4	2	1	1
(4, 3, 1, 2)	4	3	1	2
(4, 3, 2, 1)	4	3	2	1
Total	6	11	10	6

In the secretary experiment, set the number of candidates to $n = 4$. Run the experiment 1000 times with each strategy $k \in \{1, 2, 3, 4\}$

For the case $n = 5$, list the 120 permutations of $\{1, 2, 3, 4, 5\}$ and verify the probabilities in the table below. Note that $k = 3$ is optimal.

k	1	2	3	4	5
$p_5(k)$	$\frac{24}{120}$	$\frac{50}{120}$	$\frac{52}{120}$	$\frac{42}{120}$	$\frac{24}{120}$

In the secretary experiment, set the number of candidates to $n = 5$. Run the experiment 1000 times with each strategy $k \in \{1, 2, 3, 4, 5\}$

Well, clearly we don't want to keep doing *this*. Let's see if we can find a general analysis. With n candidates, let X_n denote the number (arrival order) of the best candidate, and let $S_{n,k}$ denote the event of success for strategy k (we select the best candidate).

X_n is uniformly distributed on $\{1, 2, \dots, n\}$.

Proof

This follows since the candidates arrive in random order.

Next we will compute the conditional probability of success given the arrival order of the best candidate.

For $n \in \mathbb{N}_+$ and $k \in \{2, 3, \dots, n\}$,

$$\mathbb{P}(S_{n,k} | X_n = j) = \begin{cases} 0, & j \in \{1, 2, \dots, k-1\} \\ \frac{k-1}{j-1}, & j \in \{k, k+1, \dots, n\} \end{cases} \quad (12.9.1)$$

Proof

For the first case, note that if the arrival number of the best candidate is $j < k$, then strategy k will certainly fail. For the second cases, note that if the arrival order of the best candidate is $j \geq k$, then strategy k will succeed if and only if one of the first $k-1$ candidates (the ones that are automatically rejected) is the best among the first $j-1$

The two cases are illustrated below. The large dot indicates the best candidate. Red dots indicate candidates that are rejected out of hand, while blue dots indicate candidates that are considered.




Figure 12.9.1: The case when $X_n = j < k$




Figure 12.9.2: The case when $X_n = j \geq k$

Now we can compute the probability of success with strategy k .

For $n \in \mathbb{N}_+$

$$p_n(k) = \mathbb{P}(S_{n,k}) = \begin{cases} \frac{1}{n}, & k = 1 \\ \frac{k-1}{n} \sum_{j=k}^n \frac{1}{j-1}, & k \in \{2, 3, \dots, n\} \end{cases} \quad (12.9.2)$$

Proof

When $k = 1$ we simply select the first candidate. This candidate will be the best one with probability $1/n$. The result for $k \in \{2, 3, \dots, n\}$ follows from the previous two results, by conditioning on X_n :

$$\mathbb{P}(S_{n,k}) = \sum_{j=1}^n \mathbb{P}(X_n = j) \mathbb{P}(S_{n,k} | X_n = j) = \sum_{j=k}^n \frac{1}{n} \frac{k-1}{j-1} \quad (12.9.3)$$

Values of the function p_n can be computed by hand for small n and by a computer algebra system for moderate n . The graph of p_{100} is shown below. Note the concave downward shape of the graph and the optimal value of k , which turns out to be 38. The optimal probability is about 0.37104.




Figure 12.9.3: The graph of p_{100}

The optimal strategy k_n that maximizes $k \mapsto p_n(k)$, the ratio k_n/n , and the optimal probability $p_n(k_n)$ of finding the best candidate, as functions of $n \in \{3, 4, \dots, 20\}$ are given in the following table:

Candidates n	Optimal strategy k_n	Ratio k_n/n	Optimal probability $p_n(k_n)$
3	2	0.6667	0.5000
4	2	0.5000	0.4583
5	3	0.6000	0.4333
6	3	0.5000	0.4278
7	3	0.4286	0.4143
8	4	0.5000	0.4098
9	4	0.4444	0.4060
10	4	0.4000	0.3987
11	5	0.4545	0.3984
12	5	0.4167	0.3955
13	6	0.4615	0.3923
14	6	0.4286	0.3917
15	6	0.4000	0.3894

Candidates n	Optimal strategy k_n	Ratio k_n/n	Optimal probability $p_n(k_n)$
16	7	0.4375	0.3881
17	7	0.4118	0.3873
18	7	0.3889	0.3854
19	8	0.4211	0.3850
20	8	0.4000	0.3842

Apparently, as we might expect, the optimal strategy k_n increases and the optimal probability $p_n(k_n)$ decreases as $n \rightarrow \infty$. On the other hand, it's encouraging, and a bit surprising, that the optimal probability does not appear to be decreasing to 0. It's perhaps least clear what's going on with the ratio. Graphical displays of some of the information in the table may help:




Figure 12.9.4 The optimal probability $p_n(k_n)$




Figure 12.9.5 The optimal ratio k_n/n

Could it be that the ratio k_n/n and the probability $p_n(k_n)$ are both converging, and moreover, are converging to the same number? First let's try to establish rigorously some of the trends observed in the table.

The success probability p_n satisfies

$$p_n(k-1) < p_n(k) \text{ if and only if } \sum_{j=k}^n \frac{1}{j-1} > 1 \quad (12.9.4)$$

It follows that for each $n \in \mathbb{N}_+$, the function p_n at first increases and then decreases. The maximum value of p_n occurs at the largest k with $\sum_{j=k}^n \frac{1}{j-1} > 1$. This is the optimal strategy with n candidates, which we have denoted by k_n .

As n increases, k_n increases and the optimal probability $p_n(k_n)$ decreases.

Asymptotic Analysis

We are naturally interested in the asymptotic behavior of the function p_n , and the optimal strategy as $n \rightarrow \infty$. The key is recognizing p_n as a *Riemann sum* for a simple integral. (Riemann sums, of course, are named for Georg Riemann.)

If $k(n)$ depends on n and $k(n)/n \rightarrow x \in (0, 1)$ as $n \rightarrow \infty$ then $p_n[k(n)] \rightarrow -x \ln x$ as $n \rightarrow \infty$.

Proof

First note that

$$p_n(k) = \frac{k-1}{n} \sum_{j=k}^n \frac{1}{n} \frac{n}{j-1} \quad (12.9.5)$$

We recognize the sum above as the left Riemann sum for the the function $f(t) = \frac{1}{t}$ corresponding to the partition of the interval $\left[\frac{k-1}{n}, 1\right]$ into $(n-k)+1$ subintervals of length $\frac{1}{n}$ each: $\left(\frac{k-1}{n}, \frac{k}{n}, \dots, \frac{n-1}{n}, 1\right)$. It follows that

$$p_n[k(n)] \rightarrow x \int_x^1 \frac{1}{t} dt = -x \ln x \text{ as } n \rightarrow \infty \quad (12.9.6)$$

The optimal strategy k_n that maximizes $k \mapsto p_n(k)$, the ratio k_n/n , and the optimal probability $p_n(k_n)$ of finding the best candidate, as functions of $n \in \{10, 20, \dots, 100\}$ are given in the following table:

Candidates n	Optimal strategy k_n	Ratio k_n/n	Optimal probability $p_n(k_n)$
10	4	0.4000	0.3987
20	8	0.4000	0.3842
30	12	0.4000	0.3786
40	16	0.4000	0.3757
50	19	0.3800	0.3743
60	23	0.3833	0.3732
70	27	0.3857	0.3724
80	30	0.3750	0.3719
90	34	0.3778	0.3714
100	38	0.3800	0.3710

The graph below shows the true probabilities $p_n(k)$ and the limiting values $-\frac{k}{n} \ln\left(\frac{k}{n}\right)$ as a function of k with $n = 100$.




Figure 12.9.6: True and approximate probabilities of success as a function of k with $n = 100$

For the optimal strategy k_n , there exists $x_0 \in (0, 1)$ such that $k_n/n \rightarrow x_0$ as $n \rightarrow \infty$. Thus, $x_0 \in (0, 1)$ is the limiting proportion of the candidates that we reject out of hand. Moreover, x_0 maximizes $x \mapsto -x \ln x$ on $(0, 1)$.

The maximum value of $-x \ln x$ occurs at $x_0 = 1/e$ and the maximum value is also $1/e$.

Proof




Figure 12.9.7: The graph of $x \ln x$ on the interval $(0, 1)$

Thus, the magic number $1/e \approx 0.3679$ occurs twice in the problem. For large n :

- Our approximate optimal strategy is to reject out of hand the first 37% of the candidates and then select the first candidate (if she appears) that is better than all of the previous candidates.
- Our probability of finding the best candidate is about 0.37.

The article “[Who Solved the Secretary Problem?](#)” by Tom Ferguson (1989) has an interesting historical discussion of the problem, including speculation that Johannes Kepler may have used the optimal strategy to choose his second wife. The article also discusses many interesting generalizations of the problem. A different version of the secretary problem, in which the candidates are assigned a score in $[0, 1]$, rather than a relative rank, is discussed in the section on Stopping Times in the chapter on Martingales

This page titled [12.9: The Secretary Problem](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

13: Games of Chance

Games of chance hold an honored place in probability theory, because of their conceptual clarity and because of their fundamental influence on the early development of the subject. In this chapter, we explore some of the most common and basic games of chance. Roulette, craps, and Keno are casino games. The Monty Hall problem is based on a TV game show, and has become famous because of the controversy that it generated. Lotteries are now basic ways that governments and other institutions raise money. In the last four sections on the game of red and black, we study various types of gambling strategies, a study which leads to some deep and fascinating mathematics.

[13.1: Introduction to Games of Chance](#)

[13.2: Poker](#)

[13.3: Simple Dice Games](#)

[13.4: Craps](#)

[13.5: Roulette](#)

[13.6: The Monty Hall Problem](#)

[13.7: Lotteries](#)

[13.8: The Red and Black Game](#)

[13.9: Timid Play](#)

[13.10: Bold Play](#)

[13.11: Optimal Strategies](#)

This page titled [13: Games of Chance](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.1: Introduction to Games of Chance

Gambling and Probability

Games of chance are among the oldest of human inventions. The use of a certain type of animal heel bone (called the *astragalus* or colloquially the *knucklebone*) as a crude die dates to about 3600 BCE. The modern six-sided die dates to 2000 BCE, and the term *bones* is used as a slang expression for dice to this day (as in *roll the bones*). It is because of these ancient origins, by the way, that we use the die as the fundamental symbol in this project.



Figure 13.1.1: An artificial knucklebone made of steatite, from the [Arjan Verweij Dice Website](#)

Gambling is intimately interwoven with the development of probability as a mathematical theory. Most of the early development of probability, in particular, was stimulated by special gambling problems, such as

- DeMere's problem
- Pepy's problem
- the problem of points
- the Petersburg problem

Some of the very first books on probability theory were written to analyze games of chance, for example *Liber de Ludo Aleae* (The Book on Games of Chance), by Girolamo Cardano, and *Essay d'Analyse sur les Jeux de Hazard* (Analytical Essay on Games of Chance), by Pierre-Remond Montmort. Gambling problems continue to be a source of interesting and deep problems in probability to this day (see the discussion of Red and Black for an example).



Figure 13.1.2: Allegory of Fortune by Dosso Dossi (c. 1591), Getty Museum. For more depictions of gambling in paintings, see the ancillary material on art.

Of course, it is important to keep in mind that breakthroughs in probability, even when they are originally motivated by gambling problems, are often profoundly important in the natural sciences, the social sciences, law, and medicine. Also, games of chance provide some of the conceptually clearest and cleanest examples of random experiments, and thus their analysis can be very helpful to students of probability.

However, nothing in this chapter should be construed as encouraging you, gentle reader, to gamble. On the contrary, our analysis will show that, in the long run, only the gambling houses prosper. The gambler, inevitably, is a sad victim of the law of large numbers.

In this chapter we will study some interesting games of chance. Poker, poker dice, craps, and roulette are popular parlor and casino games. The Monty Hall problem, on the other hand, is interesting because of the controversy that it generated. The lottery is a basic way that many states and nations use to raise money (a voluntary tax, of sorts).

Terminology

Let us discuss some of the basic terminology that will be used in several sections of this chapter. Suppose that A is an event in a random experiment. The *mathematical odds* concerning A refer to the probability of A .

If a and b are positive numbers, then by definition, the following are equivalent:

1. the *odds in favor* of A are $a : b$.
2. $\mathbb{P}(A) = \frac{a}{a+b}$.
3. the *odds against* A are $b : a$.
4. $\mathbb{P}(A^c) = \frac{b}{a+b}$.

In many cases, a and b can be given as positive integers with no common factors.

Similarly, suppose that $p \in [0, 1]$. The following are equivalent:

1. $\mathbb{P}(A) = p$.
2. The odds in favor of A are $p : 1 - p$.
3. $\mathbb{P}(A^c) = 1 - p$.
4. The odds against A are $1 - p : p$.

On the other hand, the *house odds* of an event refer to the *payout* when a bet is made on the event.

A bet on event A pays $n : m$ means that if a gambler bets m units on A then

1. If A occurs, the gambler receives the m units back and an additional n units (for a net profit of n)
2. If A does not occur, the gambler loses the bet of m units (for a net profit of $-m$).

Equivalently, the gambler puts up m units (betting on A), the house puts up n units, (betting on A^c) and the winner takes the pot. Of course, it is usually not necessary for the gambler to bet exactly m ; a smaller or larger is bet is scaled appropriately. Thus, if the gambler bets k units and wins, his payout is $k \frac{n}{m}$.

Naturally, our main interest is in the net *winnings* if we make a bet on an event. The following result gives the probability density function, mean, and variance for a unit bet. The expected value is particularly interesting, because by the law of large numbers, it gives the long term gain or loss, per unit bet.

Suppose that the odds in favor of event A are $a : b$ and that a bet on event A pays $n : m$. Let W denote the winnings from a unit bet on A . Then

1. $\mathbb{P}(W = -1) = \frac{b}{a+b}$, $\mathbb{P}(W = \frac{n}{m}) = \frac{a}{a+b}$
2. $\mathbb{E}(W) = \frac{a n - b m}{m(a+b)}$
3. $\text{var}(W) = \frac{ab(n+m)^2}{m^2(a+b)^2}$

In particular, the expected value of the bet is zero if and only if $an = bm$, positive if and only if $an > bm$, and negative if and only if $an < bm$. The first case means that the bet is *fair*, and occurs when the payoff is the same as the odds *against* the event. The second means that the bet is *favorable* to the gambler, and occurs when the payoff is greater than the odds against the event. The third case means that the bet is *unfair* to the gambler, and occurs when the payoff is less than the odds against the event. Unfortunately, all casino games fall into the third category.

More About Dice

Shapes of Dice

The *standard die*, of course, is a cube with six sides. A bit more generally, most real dice are in the shape of *Platonic solids*, named for Plato naturally. The faces of a Platonic solid are congruent regular polygons. Moreover, the same number of faces meet at each vertex so all of the edges and angles are congruent as well.

The five Platonic solids are

1. The *tetrahedron*, with 4 sides.
2. The *hexahedron* (cube), with 6 sides
3. The *octahedron*, with 8 sides
4. The *dodecahedron*, with 12 sides
5. The *icosahedron*, with 20 sides



Figure 13.1.3: Blue Platonic Dice from [Wikipedia](#)

Note that the 4-sided die is the only Platonic die in which the outcome is the face that is *down* rather than *up* (or perhaps it's better to think of the vertex that is up as the outcome).

Fair and Crooked Dice

Recall that a *fair* die is one in which the faces are equally likely. In addition to fair dice, there are various types of crooked dice. For the standard six-sided die, there are three crooked types that we use frequently in this project. To understand the geometry, recall that with the standard six-sided die, opposite faces sum to 7.

Flat Dice

1. An *ace-six flat die* is a six-sided die in which faces 1 and 6 have probability $\frac{1}{4}$ each while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each.
2. A *two-five flat die* is a six-sided die in which faces 2 and 5 have probability $\frac{1}{4}$ each while faces 1, 3, 4, and 6 have probability $\frac{1}{8}$ each.
3. A *three-four flat die* is a six-sided die in which faces 3 and 4 have probability $\frac{1}{4}$ each while faces 1, 2, 5, and 6 have probability $\frac{1}{8}$ each.

A flat die, as the name suggests, is a die that is not a cube, but rather is shorter in one of the three directions. The particular probabilities that we use ($\frac{1}{4}$ and $\frac{1}{8}$) are fictitious, but the essential property of a flat die is that the opposite faces on the shorter axis have slightly larger probabilities (because they have slightly larger areas) than the other four faces. Flat dice are sometimes used by gamblers to cheat.

In the Dice Experiment, select one die. Run the experiment 1000 times in each of the following cases and observe the outcomes.

1. fair die
2. ace-six flat die
3. two-five flat die
4. three-four flat die

Simulation

It's very easy to simulate a fair die with a random number. Recall that the *ceiling function* $\lceil x \rceil$ gives the smallest integer that is at least as large as x .

Suppose that U is uniformly distributed on the interval $(0, 1]$, so that U has the *standard uniform distribution* (a random number). Then $X = \lceil 6U \rceil$ is uniformly distributed on the set $\{1, 2, 3, 4, 5, 6\}$ and so simulates a fair six-sided die. More generally, $X = \lceil nU \rceil$ is uniformly distributed on $\{1, 2, \dots, n\}$ and so simulates a fair n -sided die.

We can also use a real fair die to simulate other types of fair dice. Recall that if X is uniformly distributed on $\{1, 2, \dots, n\}$ and $k \in \{1, 2, \dots, n-1\}$, then the conditional distribution of X given that $X \in \{1, 2, \dots, k\}$ is uniformly distributed on $\{1, 2, \dots, k\}$. Thus, suppose that we have a real, fair, n -sided die. If we ignore outcomes greater than k then we simulate a fair k -sided die. For example, suppose that we have a carefully constructed icosahedron that is a fair 20-sided die. We can simulate a fair 13-sided die by simply rolling the die and stopping as soon as we have a score between 1 and 13.

To see how to simulate a card hand, see the Introduction to Finite Sampling Models. A general method of simulating random variables is based on the quantile function.

This page titled [13.1: Introduction to Games of Chance](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.2: Poker

Basic Theory

The Poker Hand

A deck of cards naturally has the structure of a product set and thus can be modeled mathematically by

$$D = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, k\} \times \{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\} \quad (13.2.1)$$

where the first coordinate represents the *denomination* or *kind* (ace, two through 10, jack, queen, king) and where the second coordinate represents the *suit* (clubs, diamond, hearts, spades). Sometimes we represent a card as a *string* rather than an ordered pair (for example q♥).

There are many different poker games, but we will be interested in standard *draw poker*, which consists of dealing 5 cards at random from the deck D . The *order* of the cards does not matter in draw poker, so we will record the outcome of our random experiment as the random set (hand) $\mathbf{X} = \{X_1, X_2, X_3, X_4, X_5\}$ where $X_i = (Y_i, Z_i) \in D$ for each i and $X_i \neq X_j$ for $i \neq j$. Thus, the sample space consists of all possible poker hands:

$$S = \{\{x_1, x_2, x_3, x_4, x_5\} : x_i \in D \text{ for each } i \text{ and } x_i \neq x_j \text{ for all } i \neq j\} \quad (13.2.2)$$

Our basic modeling assumption (and the meaning of the term *at random*) is that all poker hands are equally likely. Thus, the random variable \mathbf{X} is uniformly distributed over the set of possible poker hands S .

$$\mathbb{P}(\mathbf{X} \in A) = \frac{\#(A)}{\#(S)} \quad (13.2.3)$$

In statistical terms, a poker hand is a random sample of size 5 drawn without replacement and without regard to order from the population D . For more on this topic, see the chapter on Finite Sampling Models.

The Value of the Hand

There are nine different types of poker hands in terms of value. We will use the numbers 0 to 8 to denote the *value* of the hand, where 0 is the type of least value (actually no value) and 8 the type of most value.

The hand value V of a poker hand is a random variable taking values 0 through 8, and is defined as follows:

0. *No Value*. The hand is of none of the other types.
1. *One Pair*. The hand has 2 cards of one kind, and one card each of three other kinds.
2. *Two Pair*. The hand has 2 cards of one kind, 2 cards of another kind, and one card of a third kind.
3. *Three of a Kind*. The hand has 3 cards of one kind and one card of each of two other kinds.
4. *Straight*. The kinds of cards in the hand form a consecutive sequence but the cards are not all in the same suit. An ace can be considered the smallest denomination or the largest denomination.
5. *Flush*. The cards are all in the same suit, but the kinds of the cards do not form a consecutive sequence.
6. *Full House*. The hand has 3 cards of one kind and 2 cards of another kind.
7. *Four of a Kind*. The hand has 4 cards of one kind, and 1 card of another kind.
8. *Straight Flush*. The cards are all in the same suit and the kinds form a consecutive sequence.

Run the poker experiment 10 times in single-step mode. For each outcome, note that the value of the random variable corresponds to the type of hand, as given above.

For some comic relief before we get to the analysis, look at two of the paintings of *Dogs Playing Poker* by CM Coolidge.

1. His Station and Four Aces
2. Waterloo

The Probability Density Function

Computing the probability density function of V is a good exercise in combinatorial probability. In the following exercises, we need the two fundamental rules of combinatorics to count the number of poker hands of a given type: the multiplication rule and the addition rule. We also need some basic combinatorial structures, particularly combinations.

The number of different poker hands is

$$\#(S) = \binom{52}{5} = 2\,598\,960 \quad (13.2.4)$$

$$\mathbb{P}(V = 1) = 1\,098\,240 / 2\,598\,960 \approx 0.422569$$

Proof

The following steps form an algorithm for generating poker hands with one pair. The number of ways of performing each step is also given.

1. Select a kind of card: 13
2. Select 2 cards of the kind in part (a): $\binom{4}{2}$
3. Select 3 kinds of cards, different than the kind in (a): $\binom{12}{3}$
4. Select a card of each of the kinds in part (c): 4^3

$$\mathbb{P}(V = 2) = 123\,552 / 2\,598\,960 \approx 0.047539$$

Proof

The following steps form an algorithm for generating poker hands with two pair. The number of ways of performing each step is also given.

1. Select two kinds of cards: $\binom{13}{2}$
2. Select two cards of each of the kinds in (a): $\binom{4}{2} \binom{4}{2}$
3. Select a kind of card different from the kinds in (a): 11
4. Select a card of the kind in (c): 4

$$\mathbb{P}(V = 3) = 54\,912 / 2\,598\,960 \approx 0.021129$$

Proof

The following steps form an algorithm for generating poker hands with three of a kind. The number of ways of performing each step is also given.

1. Select a kind of card: 13
2. Select 3 cards of the kind in (a): $\binom{4}{3}$
3. Select 2 kinds of cards, different than the kind in (a): $\binom{12}{2}$
4. Select one card of each of the kinds in (c): 4^2

$$\mathbb{P}(V = 8) = 40 / 2\,598\,960 \approx 0.000015$$

Proof

The following steps form an algorithm for generating poker hands with a straight flush. The number of ways of performing each step is also given.

1. Select the kind of the lowest card in the sequence: 10
2. Select a suit: 4

$$\mathbb{P}(V = 4) = 10\,200 / 2\,598\,960 \approx 0.003925$$

Proof

The following steps form an algorithm for generating poker hands with a straight or a straight flush. The number of ways of performing each step is also given.

1. Select the kind of the lowest card in the sequence: 10
2. Select a card of each kind in the sequence: 4^5

Finally, we need to subtract the [number of straight flushes](#) above to get the number of hands with a straight.

$$\mathbb{P}(V = 5) = 5108 / 2\,598\,960 \approx 0.001965$$

Proof

The following steps form an algorithm for generating poker hands with a flush or a straight flush. The number of ways of performing each step is also given.

1. Select a suit: 4
2. Select 5 cards of the suit in (a): $\binom{13}{5}$

Finally, we need to subtract the [number of straight flushes](#) above to get the number of hands with a flush.

$$\mathbb{P}(V = 6) = 3744 / 2\,598\,960 \approx 0.001441$$

Proof

The following steps form an algorithm for generating poker hands with a full house. The number of ways of performing each step is also given.

1. Select a kind of card: 13
2. Select 3 cards of the kind in (a): $\binom{4}{3}$
3. Select another kind of card: 12
4. Select 2 cards of the kind in (c): $\binom{4}{2}$

$$\mathbb{P}(V = 7) = 624 / 2\,598\,960 \approx 0.000240$$

Proof

The following steps form an algorithm for generating poker hands with four of a kind. The number of ways of performing each step is also given.

1. Select a kind of card: 13
2. Select 4 cards of the kind in (a): 1
3. Select another kind of card: 12
4. Select a card of the kind in (c): 4

$$\mathbb{P}(V = 0) = 1\,302\,540 / 2\,598\,960 \approx 0.501177$$

Proof

By the complement rule, $\mathbb{P}(V = 0) = 1 - \sum_{k=1}^8 \mathbb{P}(V = k)$

Note that the probability density function of V is decreasing; the more valuable the type of hand, the less likely the type of hand is to occur. Note also that *no value* and *one pair* account for more than 92% of all poker hands.

In the poker experiment, note the shape of the density graph. Note that some of the probabilities are so small that they are essentially invisible in the graph. Now run the poker hand 1000 times and compare the relative frequency function to the density function.

In the poker experiment, set the stop criterion to the value of V given below. Note the number of poker hands required.

1. $V = 3$
2. $V = 4$
3. $V = 5$

4. $V = 6$
5. $V = 7$
6. $V = 8$

Find the probability of getting a hand that is three of a kind or better.

Answer

0.0287

In the movie *The Parent Trap* (1998), both twins get straight flushes on the same poker deal. Find the probability of this event.

Answer

3.913×10^{-10}

Classify V in terms of level of measurement: *nominal*, *ordinal*, *interval*, or *ratio*. Is the expected value of V meaningful?

Answer

Ordinal. No.

A hand with a pair of aces and a pair of eights (and a fifth card of a different type) is called a *dead man's hand*. The name is in honor of Wild Bill Hickok, who held such a hand at the time of his murder in 1876. Find the probability of getting a dead man's hand.

Answer

1584/2 598 960

Drawing Cards

In *draw poker*, each player is dealt a poker hand and there is an initial round of betting. Typically, each player then gets to discard up to 3 cards and is dealt that number of cards from the remaining deck. This leads to myriad problems in conditional probability, as partial information becomes available. A complete analysis is far beyond the scope of this section, but we will consider a couple of simple examples.

Suppose that Fred's hand is $\{4 \heartsuit, 5 \heartsuit, 7 \spadesuit, q \clubsuit, 1 \diamondsuit\}$. Fred discards the $q \clubsuit$ and $1 \diamondsuit$ and draws two new cards, hoping to complete the straight. Note that Fred must get a 6 and either a 3 or an 8. Since he is missing a middle denomination (6), Fred is *drawing to an inside straight*. Find the probability that Fred is successful.

Answer

32/1081

Suppose that Wilma's hand is $\{4 \heartsuit, 5 \heartsuit, 6 \spadesuit, q \clubsuit, 1 \diamondsuit\}$. Wilma discards $q \clubsuit$ and $1 \diamondsuit$ and draws two new cards, hoping to complete the straight. Note that Wilma must get a 2 and a 3, or a 7 and an 8, or a 3 and a 7. Find the probability that Wilma is successful. Clearly, Wilma has a better chance than Fred.

Answer

48/1081

This page titled [13.2: Poker](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.3: Simple Dice Games

In this section, we will analyze several simple games played with dice—[poker dice](#), [chuck-a-luck](#), and [high-low](#). The casino game craps is more complicated and is studied in the next section.



Figure 13.3.1: The Dice Players by Georges de La Tour (c. 1651). For more on the influence of probability in painting, see the ancillary material on art.

Poker Dice

Definition

The game of *poker dice* is a bit like standard poker, but played with dice instead of cards. In poker dice, 5 fair dice are rolled. We will record the outcome of our random experiment as the (ordered) sequence of scores:

$$\mathbf{X} = (X_1, X_2, X_3, X_4, X_5) \quad (13.3.1)$$

Thus, the sample space is $S = \{1, 2, 3, 4, 5, 6\}^5$. Since the dice are fair, our basic modeling assumption is that \mathbf{X} is a sequence of independent random variables and each is uniformly distributed on $\{1, 2, 3, 4, 5, 6\}$

Equivalently, \mathbf{X} is uniformly distributed on S :

$$\mathbb{P}(\mathbf{X} \in A) = \frac{\#(A)}{\#(S)}, \quad A \subseteq S \quad (13.3.2)$$

In statistical terms, a poker dice hand is a random sample of size 5 drawn with replacement and with regard to order from the population $D = \{1, 2, 3, 4, 5, 6\}$. For more on this topic, see the chapter on Finite Sampling Models. In particular, in this chapter you will learn that the result of Exercise 1 would not be true if we recorded the outcome of the poker dice experiment as an unordered set instead of an ordered sequence.

The Value of the Hand

The value V of the poker dice hand is a random variable with support set $\{0, 1, 2, 3, 4, 5, 6\}$. The values are defined as follows:

0. *None alike*. Five distinct scores occur.
1. *One Pair*. Four distinct scores occur; one score occurs twice and the other three scores occur once each.
2. *Two Pair*. Three distinct scores occur; one score occurs twice and the other three scores occur once each.
3. *Three of a Kind*. Three distinct scores occur; one score occurs three times and the other two scores occur once each.
4. *Full House*. Two distinct scores occur; one score occurs three times and the other score occurs twice.
5. *Four of a kind*. Two distinct scores occur; one score occurs four times and the other score occurs once.
6. *Five of a kind*. Once score occurs five times.

Run the poker dice experiment 10 times in single-step mode. For each outcome, note that the value of the random variable corresponds to the type of hand, as given above.

The Probability Density Function

Computing the probability density function of V is a good exercise in combinatorial probability. In the following exercises, we will need the two fundamental rules of combinatorics to count the number of dice sequences of a given type: the multiplication rule and the addition rule. We will also need some basic combinatorial structures, particularly combinations and permutations (with types of objects that are identical).

The number of different poker dice hands is $\#(S) = 6^5 = 7776$.

$$\mathbb{P}(V = 0) = \frac{720}{7776} = 0.09259.$$

Proof

Note that the dice scores form a permutation of size 5 from $\{1, 2, 3, 4, 5\}$

$$\mathbb{P}(V = 1) = \frac{3600}{7776} \approx 0.46296.$$

Proof

The following steps form an algorithm for generating poker dice hands with one pair. The number of ways of performing each step is also given:

1. Select the score that will appear twice: 6
2. Select the 3 scores that will appear once each: $\binom{5}{3}$
3. Select a permutation of the 5 numbers in parts (a) and (b): $\binom{5}{2,1,1,1,1}$

$$\mathbb{P}(V = 2) = \frac{1800}{7776} \approx 0.23148.$$

Proof

The following steps form an algorithm for generating poker dice hands with two pair. The number of ways of performing each step is also given:

1. Select two scores that will appear twice each: $\binom{6}{2}$
2. Select the score that will appear once: 4
3. Select a permutation of the 5 numbers in parts (a) and (b): $\binom{5}{2,2,1}$

$$\mathbb{P}(V = 3) = \frac{1200}{7776} \approx 0.15432.$$

Proof

The following steps form an algorithm for generating poker dice hands with three of a kind. The number of ways of performing each step is also given:

1. Select the score that will appear 3 times: 6
2. Select the 2 scores that will appear once each: $\binom{5}{2}$
3. Select a permutation of the 5 numbers in parts (a) and (b): $\binom{5}{3,1,1}$

$$\mathbb{P}(V = 4) = \frac{300}{7776} \approx 0.03858.$$

Proof

The following steps form an algorithm for generating poker dice hands with a full house. The number of ways of performing each step is also given:

1. Select the score that will appear 3 times: 6
2. Select the score that will appear twice: 5
3. Select a permutation of the 5 numbers in parts (a) and (b): $\binom{5}{3,2}$

$$\mathbb{P}(V = 5) = \frac{150}{7776} = 0.01929.$$

Proof

The following steps form an algorithm for generating poker dice hands with four of a kind. The number of ways of performing each step is also given:

1. Select the score that will appear 4 times: 6
2. Select the score that will appear once: 5
3. Select a permutation of the 5 numbers in parts (a) and (b): $\binom{5}{4,1}$

$$\mathbb{P}(V = 6) = \frac{6}{7776} \approx 0.00077.$$

Proof

There are 6 choices for the score that will appear 5 times.

Run the poker dice experiment 1000 times and compare the relative frequency function to the density function.

Find the probability of rolling a hand that has 3 of a kind or better.

Answer

0.2130

In the poker dice experiment, set the stop criterion to the value of V given below. Note the number of hands required.

1. $V = 3$
2. $V = 4$
3. $V = 5$
4. $V = 6$

Chuck-a-Luck

Chuck-a-luck is a popular carnival game, played with three dice. According to Richard Epstein, the original name was *Sweat Cloth*, and in British pubs, the game is known as *Crown and Anchor* (because the six sides of the dice are inscribed clubs, diamonds, hearts, spades, crown and anchor). The dice are over-sized and are kept in an hourglass-shaped cage known as the *bird cage*. The dice are rolled by spinning the bird cage.

Chuck-a-luck is very simple. The gambler selects an integer from 1 to 6, and then the three dice are rolled. If exactly k dice show the gambler's number, the payoff is $k : 1$. As with poker dice, our basic mathematical assumption is that the dice are fair, and therefore the outcome vector $\mathbf{X} = (X_1, X_2, X_3)$ is uniformly distributed on the sample space $S = \{1, 2, 3, 4, 5, 6\}^3$.

Let Y denote the number of dice that show the gambler's number. Then Y has the binomial distribution with parameters $n = 3$ and $p = \frac{1}{6}$:

$$\mathbb{P}(Y = k) = \binom{3}{k} \left(\frac{1}{6}\right)^k \left(\frac{5}{6}\right)^{3-k}, \quad k \in \{0, 1, 2, 3\} \quad (13.3.3)$$

Let W denote the net winnings for a unit bet. Then

1. $W = -1$ if $Y = 0$
2. $W = Y$ if $Y > 0$

The probability density function of W is given by

1. $\mathbb{P}(W = -1) = \frac{125}{216}$
2. $\mathbb{P}(W = 1) = \frac{75}{216}$
3. $\mathbb{P}(W = 2) = \frac{15}{216}$

$$4. \mathbb{P}(W = 3) = \frac{1}{216}$$

Run the chuck-a-luck experiment 1000 times and compare the empirical density function of W to the true probability density function.

The expected value and variance of W are

$$1. \mathbb{E}(W) = -\frac{17}{216} \approx 0.0787$$

$$2. \text{var}(W) = \frac{75815}{46656} \approx 1.239$$

Run the chuck-a-luck experiment 1000 times and compare the empirical mean and standard deviation of W to the true mean and standard deviation. Suppose you had bet \$1 on each of the 1000 games. What would your net winnings be?

High-Low

In the game of *high-low*, a pair of fair dice are rolled. The outcome is

- *high* if the sum is 8, 9, 10, 11, or 12.
- *low* if the sum is 2, 3, 4, 5, or 6
- *seven* if the sum is 7

A player can bet on any of the three outcomes. The payoff for a bet of high or for a bet of low is 1 : 1. The payoff for a bet of seven is 4 : 1.

Let Z denote the outcome of a game of high-low. Find the probability density function of Z .

Answer

$$\mathbb{P}(Z = h) = \frac{15}{36}, \mathbb{P}(Z = l) = \frac{15}{36}, \mathbb{P}(Z = s) = \frac{6}{36}, \text{ where } h \text{ denotes high, } l \text{ denotes low, and } s \text{ denotes seven.}$$

Let W denote the net winnings for a unit bet. Find the expected value and variance of W for each of the three bets:

1. high
2. low
3. seven

Answer

Let W denote the net winnings on a unit bet in high-low.

1. Bet high: $\mathbb{E}(W) = -\frac{1}{6}, \text{var}(W) = \frac{35}{36}$
2. Bet low: $\mathbb{E}(W) = -\frac{1}{6}, \text{var}(W) = \frac{35}{36}$
3. Bet seven: $\mathbb{E}(W) = -\frac{1}{6}, \text{var}(W) = \frac{7}{2}$

This page titled [13.3: Simple Dice Games](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.4: Craps

The Basic Game

Craps is a popular casino game, because of its complexity and because of the rich variety of bets that can be made.




Figure 13.4.1: A typical craps table

According to Richard Epstein, craps is descended from an earlier game known as *Hazard*, that dates to the Middle Ages. The formal rules for Hazard were established by Montmort early in the 1700s. The origin of the name *craps* is shrouded in doubt, but it may have come from the English *crabs* or from the French *Crapeaud* (for toad).

From a mathematical point of view, craps is interesting because it is an example of a random experiment that takes place in stages; the evolution of the game depends critically on the outcome of the first roll. In particular, the number of rolls is a random variable.

Definitions

The rules for craps are as follows:

The player (known as the *shooter*) rolls a pair of fair dice

1. If the sum is 7 or 11 on the first throw, the shooter wins; this event is called a *natural*.
2. If the sum is 2, 3, or 12 on the first throw, the shooter loses; this event is called *craps*.
3. If the sum is 4, 5, 6, 8, 9, or 10 on the first throw, this number becomes the shooter's *point*. The shooter continues rolling the dice until either she rolls the point again (in which case she wins) or rolls a 7 (in which case she loses).

As long as the shooter wins, or loses by rolling craps, she retains the dice and continues. Once she loses by failing to make her point, the dice are passed to the next shooter.

Let us consider the game of craps mathematically. Our basic assumption, of course, is that the dice are fair and that the outcomes of the various rolls are independent. Let N denote the (random) number of rolls in the game and let (X_i, Y_i) denote the outcome of the i th roll for $i \in \{1, 2, \dots, N\}$. Finally, let $Z_i = X_i + Y_i$, the sum of the scores on the i th roll, and let V denote the event that the shooter wins.

In the craps experiment, press single step a few times and observe the outcomes. Make sure that you understand the rules of the game.

The Probability of Winning

We will compute the probability that the shooter wins in stages, based on the outcome of the first roll.

The sum of the scores Z on a given roll has the probability density function in the following table:

z	2	3	4	5	6	7	8	9	10	11	12
$\mathbb{P}(Z = z)$	$\frac{1}{36}$	$\frac{2}{36}$	$\frac{3}{36}$	$\frac{4}{36}$	$\frac{5}{36}$	$\frac{6}{36}$	$\frac{5}{36}$	$\frac{4}{36}$	$\frac{3}{36}$	$\frac{2}{36}$	$\frac{1}{36}$

The probability that the player makes her point can be computed using a simple conditioning argument. For example, suppose that the player throws 4 initially, so that 4 is the point. The player continues until she either throws 4 again or throws 7. Thus, the final roll will be an element of the following set:

$$S_4 = \{(1, 3), (2, 2), (3, 1), (1, 6), (2, 5), (3, 4), (4, 3), (5, 2), (6, 1)\} \quad (13.4.1)$$

Since the dice are fair, these outcomes are equally likely, so the probability that the player makes her 4 point is $\frac{3}{9}$. A similar argument can be used for the other points. Here are the results:

The probabilities of making the point z are given in the following table:

z	4	5	6	8	9	10
$\mathbb{P}(V \mid Z_1 = z)$	$\frac{3}{9}$	$\frac{4}{10}$	$\frac{5}{11}$	$\frac{5}{11}$	$\frac{4}{10}$	$\frac{3}{9}$

The probability that the shooter wins is $\mathbb{P}(V) = \frac{244}{495} \approx 0.49293$

Proof

This follows from the [the rules of the game](#) and the [the previous result](#), by conditioning on the first roll:

$$\mathbb{P}(V) = \sum_{z=2}^{12} \mathbb{P}(Z_1 = z) \mathbb{P}(I = 1 \mid Z_1 = z) \quad (13.4.2)$$

Note that craps is nearly a fair game. For the sake of completeness, the following result gives the probability of winning, given a “point” on the first roll.

$$\mathbb{P}(V \mid Z_1 \in \{4, 5, 6, 8, 9, 10\}) = \frac{67}{165} \approx 0.406$$

Proof

Let $A = \{4, 5, 6, 8, 9, 10\}$ From the definition of conditional probability,

$$\mathbb{P}(V \mid Z_1 \in A) = \frac{\mathbb{P}(V \cap \{Z_1 \in A\})}{\mathbb{P}(Z_1 \in A)} \quad (13.4.3)$$

For the numerator, using our results above,

$$\mathbb{P}(V \cap \{Z_1 \in A\}) = \sum_{z \in A} \mathbb{P}(V \mid Z_1 = z) \mathbb{P}(Z_1 = z) = \frac{134}{495} \quad (13.4.4)$$

Also from previous results $\mathbb{P}(Z_1 \in A) = \frac{2}{3}$.

Bets

There is a bewildering variety of bets that can be made in craps. In the exercises in this subsection, we will discuss some typical bets and compute the probability density function, mean, and standard deviation of each. (Most of these bets are illustrated in the picture of the craps table above). Note however, that some of the details of the bets and, in particular the payout odds, vary from one casino to another. Of course the expected value of any bet is inevitably negative (for the gambler), and thus the gambler is doomed to lose money in the long run. Nonetheless, as we will see, some bets are better than others.

Pass and Don't Pass

A *pass bet* is a bet that the shooter will win and pays 1 : 1.

Let W denote the winnings from a unit pass bet. Then

1. $\mathbb{P}(W = -1) = \frac{251}{495}$, $\mathbb{P}(W = 1) = \frac{244}{495}$
2. $\mathbb{E}(W) = -\frac{7}{495} \approx -0.0141$

3. $\text{sd}(W) \approx 0.9999$

In the craps experiment, select the pass bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

A *don't pass bet* is a bet that the shooter will lose, except that 12 on the first throw is excluded (that is, the shooter loses, of course, but the don't pass bet neither wins nor loses). This is the meaning of the phrase *don't pass bar double 6* on the craps table. The don't pass bet also pays 1 : 1.

Let W denote the winnings for a unit don't pass bet. Then

1. $\mathbb{P}(W = -1) = \frac{244}{495}$, $\mathbb{P}(W = 0) = \frac{1}{36}$, $\mathbb{P}(W = 1) = \frac{949}{1980}$
2. $\mathbb{E}(W) = -\frac{27}{1980} \approx -0.01363$
3. $\text{sd}(W) \approx 0.9859$

Thus, the *don't pass* bet is slightly better for the gambler than the *pass bet*.

In the craps experiment, select the don't pass bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

The *come bet* and the *don't come bet* are analogous to the pass and don't pass bets, respectively, except that they are made *after* the point has been established.

Field

A *field bet* is a bet on the outcome of the next throw. It pays 1 : 1 if 3, 4, 9, 10, or 11 is thrown, 2 : 1 if 2 or 12 is thrown, and loses otherwise.

Let W denote the winnings for a unit field bet. Then

1. $\mathbb{P}(W = -1) = \frac{5}{9}$, $\mathbb{P}(W = 1) = \frac{7}{18}$, $\mathbb{P}(W = 2) = \frac{1}{18}$
2. $\mathbb{E}(W) = -\frac{1}{18} \approx -0.0556$
3. $\text{sd}(W) \approx 1.0787$

In the craps experiment, select the field bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Seven and Eleven

A *7 bet* is a bet on the outcome of the next throw. It pays 4 : 1 if a 7 is thrown. Similarly, an *11 bet* is a bet on the outcome of the next throw, and pays 15 : 1 if an 11 is thrown. In spite of the romance of the number 7, the next exercise shows that the 7 bet is one of the worst bets you can make.

Let W denote the winnings for a unit 7 bet. Then

1. $\mathbb{P}(W = -1) = \frac{5}{6}$, $\mathbb{P}(W = 4) = \frac{1}{6}$
2. $\mathbb{E}(W) = -\frac{1}{6} \approx -0.1667$
3. $\text{sd}(W) \approx 1.8634$

In the craps experiment, select the 7 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Let W denote the winnings for a unit 11 bet. Then

1. $\mathbb{P}(W = -1) = \frac{17}{18}, \mathbb{P}(W = 15) = \frac{1}{18}$
2. $\mathbb{E}(W) = -\frac{1}{9} \approx -0.1111$
3. $\text{sd}(W) \approx 3.6650$

In the craps experiment, select the 11 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Craps

All *craps bets* are bets on the next throw. The basic *craps bet* pays 7 : 1 if 2, 3, or 12 is thrown. The *craps 2 bet* pays 30 : 1 if a 2 is thrown. Similarly, the craps 12 bet pays 30 : 1 if a 12 is thrown. Finally, the *craps 3 bet* pays 15 : 1 if a 3 is thrown.

Let W denote the winnings for a unit craps bet. Then

1. $\mathbb{P}(W = -1) = \frac{8}{9}, \mathbb{P}(W = 7) = \frac{1}{9}$
2. $\mathbb{E}(W) = -\frac{1}{9} \approx -0.1111$
3. $\text{sd}(W) \approx 5.0944$

In the craps experiment, select the craps bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Let W denote the winnings for a unit craps 2 bet or a unit craps 12 bet. Then

1. $\mathbb{P}(W = -1) = \frac{35}{36}, \mathbb{P}(W = 30) = \frac{1}{36}$
2. $\mathbb{E}(W) = -\frac{5}{36} \approx -0.1389$
3. $\text{sd}(W) = 5.0944$

In the craps experiment, select the craps 2 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

In the craps experiment, select the craps 12 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Let W denote the winnings for a unit craps 3 bet. Then

1. $\mathbb{P}(W = -1) = \frac{17}{18}, \mathbb{P}(W = 15) = \frac{1}{18}$
2. $\mathbb{E}(W) = -\frac{1}{9} \approx -0.1111$
3. $\text{sd}(W) \approx 3.6650$

In the craps experiment, select the craps 3 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Thus, of the craps bets, the basic craps bet and the craps 3 bet are best for the gambler, and the craps 2 and craps 12 are the worst.

Big Six and Big Eight

The *big 6 bet* is a bet that 6 is thrown before 7. Similarly, the *big 8 bet* is a bet that 8 is thrown before 7. Both pay even money 1 : 1.

Let W denote the winnings for a unit big 6 bet or a unit big 8 bet. Then

1. $\mathbb{P}(W = -1) = \frac{6}{11}$, $\mathbb{P}(W = 1) = \frac{5}{11}$
2. $\mathbb{E}(W) = -\frac{1}{11} \approx -0.0909$
3. $\text{sd}(W) \approx 0.9959$

In the craps experiment, select the big 6 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

In the craps experiment, select the big 8 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Hardway Bets

A *hardway bet* can be made on any of the numbers 4, 6, 8, or 10. It is a bet that the chosen number n will be thrown “the hardway” as $(n/2, n/2)$ before 7 is thrown and before the chosen number is thrown in any other combination. Hardway bets on 4 and 10 pay 7 : 1, while hardway bets on 6 and 8 pay 9 : 1.

Let W denote the winnings for a unit hardway 4 or hardway 10 bet. Then

1. $\mathbb{P}(W = -1) = \frac{8}{9}$, $\mathbb{P}(W = 7) = \frac{1}{9}$
2. $\mathbb{E}(W) = -\frac{1}{9} \approx -0.1111$
3. $\text{sd}(W) = 2.5142$

In the craps experiment, select the hardway 4 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

In the craps experiment, select the hardway 10 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Let W denote the winnings for a unit hardway 6 or hardway 8 bet. Then

1. $\mathbb{P}(W = -1) = \frac{10}{11}$, $\mathbb{P}(W = 9) = \frac{1}{11}$
2. $\mathbb{E}(W) = -\frac{1}{11} \approx -0.0909$
3. $\text{sd}(W) \approx 2.8748$

In the craps experiment, select the hardway 6 bet. Run the simulation 1000 times and compare the empirical density and moments of W to the true density and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

In the craps experiment, select the hardway 8 bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Thus, the hardway 6 and 8 bets are better than the hardway 4 and 10 bets for the gambler, in terms of expected value.

The Distribution of the Number of Rolls

Next let us compute the distribution and moments of the number of rolls N in a game of craps. This random variable is of no special interest to the casino or the players, but provides a good mathematical exercise. By definition, if the shooter wins or loses on the first roll, $N = 1$. Otherwise, the shooter continues until she either makes her point or rolls 7. In this latter case, we can use the geometric distribution on \mathbb{N}_+ which governs the trial number of the first success in a sequence of Bernoulli trials. The distribution of N is a mixture of distributions.

The probability density function of N is

$$\mathbb{P}(N = n) = \begin{cases} \frac{12}{36}, & n = 1 \\ \frac{1}{24} \left(\frac{3}{4}\right)^{n-2} + \frac{5}{81} \left(\frac{13}{18}\right)^{n-2} + \frac{55}{648} \left(\frac{25}{36}\right)^{n-2}, & n \in \{2, 3, \dots\} \end{cases} \quad (13.4.5)$$

Proof

First note that $\mathbb{P}(N = 1 \mid Z_1 = z) = 1$ for $z \in \{2, 3, 7, 11, 12\}$. Next, $\mathbb{P}(N = n \mid Z_1 = z) = p_z(1 - p_z)^{n-2}$ for $n \in \{2, 3, \dots\}$ and for the values of z and p_z given in the following table:

z	4	5	6	8	9	10
p_z	$\frac{9}{36}$	$\frac{10}{36}$	$\frac{11}{36}$	$\frac{11}{36}$	$\frac{10}{36}$	$\frac{9}{36}$

Thus the conditional distribution of $N - 1$ given $Z = z$ is geometric with probability p_z . The final result now follows by conditioning on the first roll:

$$\mathbb{P}(N = n) = \sum_{z=2}^{12} \mathbb{P}(Z_1 = z) \mathbb{P}(N = n \mid Z_1 = z) \quad (13.4.6)$$

The first few values of the probability density function of N are given in the following table:

n	1	2	3	4	5
$\mathbb{P}(N = n)$	0.33333	0.18827	0.13477	0.09657	0.06926

Find the probability that a game of craps will last at least 8 rolls.

Answer

0.09235

The mean and variance of the number of rolls are

$$1. \mathbb{E}(N) = \frac{557}{165} \approx 3.3758$$

$$2. \text{var}(N) = \frac{245\,672}{27\,225} \approx 9.02376$$

Proof

These result also can be obtained by conditioning on the first roll:

$$\mathbb{E}(N) = \mathbb{E}[\mathbb{E}(N \mid Z_1)] = \frac{557}{165} \quad (13.4.7)$$

$$\mathbb{E}(N^2) = \mathbb{E}[\mathbb{E}(N^2 \mid Z_1)] = \frac{61\,769}{3025} \quad (13.4.8)$$

This page titled [13.4: Craps](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.5: Roulette

The Roulette Wheel

According to Richard Epstein, roulette is the oldest casino game still in operation. Its invention has been variously attributed to Blaise Pascal, the Italian mathematician Don Pasquale, and several others. In any event, the roulette wheel was first introduced into Paris in 1765. Here are the characteristics of the wheel:

The (American) *roulette wheel* has 38 slots numbered 00, 0, and 1–36.

1. Slots 0, 00 are green;
2. Slots 1, 3, 5, 7, 9, 12, 14, 16, 18, 19, 21, 23, 25, 27, 30, 32, 34, 36 are red;
3. Slots 2, 4, 6, 8, 10, 11, 13, 15, 17, 20, 22, 24, 26, 28, 29, 31, 33, 35 are black.

Except for 0 and 00, the slots on the wheel alternate between red and black. The strange order of the numbers on the wheel is intended so that high and low numbers, as well as odd and even numbers, tend to alternate.




Figure 13.5.1: A typical roulette wheel and table

The roulette experiment is very simple. The wheel is spun and then a small ball is rolled in a groove, in the opposite direction as the motion of the wheel. Eventually the ball falls into one of the slots. Naturally, we assume mathematically that the wheel is fair, so that the random variable X that gives the slot number of the ball is uniformly distributed over the sample space $S = \{00, 0, 1, \dots, 36\}$. Thus, $\mathbb{P}(X = x) = \frac{1}{38}$ for each $x \in S$.

Bets

As with craps, roulette is a popular casino game because of the rich variety of bets that can be made. The [picture above](#) shows the roulette table and indicates some of the bets we will study. All bets turn out to have the same expected value (negative, of course). However, the variances differ depending on the bet.

Although all bets in roulette have the same expected value, the standard deviations vary inversely with the number of numbers selected. What are the implications of this for the gambler?

Straight Bets

A *straight bet* is a bet on a single number, and pays 35 : 1.

Let W denote the winnings on a unit straight bet. Then

1. $\mathbb{P}(W = -1) = \frac{37}{38}$, $\mathbb{P}(W = 35) = \frac{1}{38}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 5.7626$

In the roulette experiment, select the single number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Two Number Bets

A *2-number bet* (or a *split bet*) is a bet on two adjacent numbers in the roulette table. The bet pays 17 : 1.

Let W denote the winnings on a unit split bet. Then

1. $\mathbb{P}(W = -1) = \frac{18}{19}$, $\mathbb{P}(W = 17) = \frac{1}{19}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 4.0193$

In the roulette experiment, select the 2 number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Three Number Bets

A *3-number bet* (or *row bet*) is a bet on the three numbers in a vertical row on the roulette table. The bet pays 11 : 1.

Let W denote the winnings on a unit row bet. Then

1. $\mathbb{P}(W = -1) = \frac{35}{38}$, $\mathbb{P}(W = 11) = \frac{3}{38}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 3.2359$

In the roulette experiment, select the 3-number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Four Number Bets

A *4-number bet* or a *square bet* is a bet on the four numbers that form a square on the roulette table. The bet pays 8 : 1.

Let W denote the winnings on a unit 4-number bet. Then

1. $\mathbb{P}(W = -1) = \frac{17}{19}$, $\mathbb{P}(W = 8) = \frac{2}{19}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 2.7620$

In the roulette experiment, select the 4-number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Six Number Bets

A *6-number bet* or *2-row bet* is a bet on the 6 numbers in two adjacent rows of the roulette table. The bet pays 5 : 1.

Let W denote the winnings on a unit 6-number bet. Then

1. $\mathbb{P}(W = -1) = \frac{16}{19}$, $\mathbb{P}(W = 5) = \frac{3}{19}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 2.1879$

In the roulette experiment, select the 6-number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Twelve Number Bets

A *12-number bet* is a bet on 12 numbers. In particular, a *column bet* is bet on any one of the three columns of 12 numbers running horizontally along the table. Other 12-number bets are the *first 12* (1-12), the *middle 12* (13-24), and the *last 12* (25-36). A 12-number bet pays 2 : 1.

Let W denote the winnings on a unit 12-number bet. Then

1. $\mathbb{P}(W = -1) = \frac{13}{19}$, $\mathbb{P}(W = 2) = \frac{6}{19}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 1.3945$

In the roulette experiment, select the 12-number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

Eighteen Number Bets

An *18-number bet* is a bet on 18 numbers. In particular, A *color bet* is a bet either on *red* or on *black*. A *parity bet* is a bet on the *odd* numbers from 1 to 36 or the *even* numbers from 1 to 36. The *low bet* is a bet on the numbers 1-18, and the *high bet* is the bet on the numbers from 19-36. An 18-number bet pays 1 : 1.

Let W denote the winnings on a unit 18-number bet. Then

1. $\mathbb{P}(W = -1) = \frac{10}{19}$, $\mathbb{P}(W = 1) = \frac{9}{19}$
2. $\mathbb{E}(W) = -\frac{1}{19} \approx -0.0526$
3. $\text{sd}(W) \approx 0.9986$

In the roulette experiment, select the 18-number bet. Run the simulation 1000 times and compare the empirical density function and moments of W to the true probability density function and moments. Suppose that you bet \$1 on each of the 1000 games. What would your net winnings be?

This page titled [13.5: Roulette](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.6: The Monty Hall Problem

Preliminaries

Statement of the Problem

The *Monty Hall problem* involves a classical game show situation and is named after Monty Hall, the long-time host of the TV game show *Let's Make a Deal*. There are three doors labeled 1, 2, and 3. A car is behind one of the doors, while goats are behind the other two:



Figure 13.6.1: The car and the two goats

The rules are as follows:

1. The player selects a door.
2. The host selects a different door and opens it.
3. The host gives the player the option of switching from her original choice to the remaining closed door.
4. The door finally selected by the player is opened and she either wins or loses.

The Monty Hall problem became the subject of intense controversy because of several articles by Marilyn Vos Savant in the *Ask Marilyn* column of *Parade* magazine, a popular Sunday newspaper supplement. The controversy began when a reader posed the problem in the following way:

Suppose you're on a game show, and you're given a choice of three doors. Behind one door is a car; behind the others, goats. You pick a door—say No. 1—and the host, who knows what's behind the doors, opens another door—say No. 3—which has a goat. He then says to you, “Do you want to pick door No. 2?” Is it to your advantage to switch your choice?

Marilyn's response was that the contestant should switch doors, claiming that there is a $\frac{1}{3}$ chance that the car is behind door 1, while there is a $\frac{2}{3}$ chance that the car is behind door 2. In two follow-up columns, Marilyn printed a number of responses, some from academics, most of whom claimed in angry or sarcastic tones that she was wrong and that there are equal chances that the car is behind doors 1 or 2. Marilyn stood by her original answer and offered additional, but non-mathematical, arguments.

Think about the problem. Do you agree with Marilyn or with her critics, or do you think that neither solution is correct?

In the Monty Hall game, set the host strategy to *standard* (the meaning of this strategy will be explained in the below). Play the Monty Hall game 50 times with each of the following strategies. Do you want to reconsider your answer to [question above](#)?

1. Always switch
2. Never switch

In the Monty Hall game, set the host strategy to *blind* (the meaning of this strategy will be explained below). Play the Monty Hall game 50 times with each of the following strategies. Do you want to reconsider your answer to [question above](#)?

1. Always switch
2. Never switch

Modeling Assumptions

When we begin to think carefully about the Monty Hall problem, we realize that the statement of the problem by Marilyn's reader is so vague that a meaningful discussion is not possible without clarifying assumptions about the strategies of the host and player. Indeed, we will see that misunderstandings about these strategies are the cause of the controversy.

Let us try to formulate the problem mathematically. In general, the actions of the host and player can vary from game to game, but if we are to have a random experiment in the classical sense, we must assume that the same probability distributions govern the host and player on each game and that the games are independent.

There are four basic random variables for a game:

1. U : the number of the door containing the car.
2. X : the number of the first door selected by the player.
3. V : the number of the door opened by the host.
4. Y : the number of the second door selected by the player.

Each of these random variables has the possible values 1, 2, and 3. However, because of the rules of the game, the door opened by the host cannot be either of the doors selected by the player, so $V \neq X$ and $V \neq Y$. In general, we will allow the possibility $V = U$, that the host opens the door with the car behind it. Whether this is a *reasonable* action of the host is a big part of the controversy about this problem.

The Monty Hall experiment will be completely defined mathematically once the joint distribution of the basic variables is specified. This joint distribution in turn depends on the strategies of the host and player, which we will consider next.

Strategies

Host Strategies

In the Monty Hall experiment, note that the host determines the probability density function of the door containing the car, namely $\mathbb{P}(U = i)$ for $i \in \{1, 2, 3\}$. The obvious choice for the host is to randomly assign the car to one of the three doors. This leads to the uniform distribution, and unless otherwise noted, we will always assume that U has this distribution. Thus, $\mathbb{P}(U = i) = \frac{1}{3}$ for $i \in \{1, 2, 3\}$.

The host also determines the conditional density function of the door he opens, given knowledge of the door containing the car and the first door selected by the player, namely $\mathbb{P}(V = k \mid U = i, X = j)$ for $i, j \in \{1, 2, 3\}$. Recall that since the host cannot open the door chosen by the player, this probability must be 0 for $k = j$.

Thus, the distribution of U and the conditional distribution of V given U and X constitute the *host strategy*.

The Standard Strategy

In most real game shows, the host would always open a door with a goat behind it. If the player's first choice is incorrect, then the host has no choice; he cannot open the door with the car or the player's choice and must therefore open the only remaining door. On the other hand, if the player's first choice is correct, then the host can open either of the remaining doors, since goats are behind both. Thus, he might naturally pick one of these doors at random.

This strategy leads to the following conditional distribution for V given U and X :

$$\mathbb{P}(V = k \mid U = i, X = j) = \begin{cases} 1, & i \neq j, i \neq k, k \neq j \\ \frac{1}{2}, & i = j, k \neq i \\ 0, & k = i, k = j \end{cases} \quad (13.6.1)$$

This distribution, along with the uniform distribution for U , will be referred to as the *standard strategy* for the host.

In the Monty Hall game, set the host strategy to *standard*. Play the game 50 times with each of the following player strategies. Which works better?

1. Always switch
2. Never switch

The Blind Strategy

Another possible second-stage strategy is for the host to always open a door chosen at random from the two possibilities. Thus, the host might well open the door containing the car.

This strategy leads to the following conditional distribution for V given U and X :

$$\mathbb{P}(V = k \mid U = i, X = j) = \begin{cases} \frac{1}{2}, & k \neq j \\ 0, & k = j \end{cases} \quad (13.6.2)$$

This distribution, together with the uniform distribution for U , will be referred to as the *blind strategy* for the host. The blind strategy seems a bit odd. However, the confusion between the two strategies is the source of the controversy concerning this problem.

In the Monty Hall game, set the host strategy to *blind*. Play the game 50 times with each of the following player strategies. Which works better?

1. Always switch
2. Never switch

Player Strategies

The player, on the other hand, determines the probability density function of her first choice, namely $\mathbb{P}(X = j)$ for $j \in \{1, 2, 3\}$. The obvious first choice for the player is to randomly choose a door, since the player has no knowledge at this point. This leads to the uniform distribution, so $\mathbb{P}(X = j) = \frac{1}{3}$ for $j \in \{1, 2, 3\}$.

The player also determines the conditional density function of her second choice, given knowledge of her first choice and the door opened by the host, namely $\mathbb{P}(Y = l \mid X = j, V = k)$ for $i, j, k \in \{1, 2, 3\}$ with $j \neq k$. Recall that since the player cannot choose the door opened by the host, this probability must be 0 for $l = k$. The distribution of X and the conditional distribution of Y given X and V constitute the *player strategy*.

Suppose that the player switches with probability $p \in [0, 1]$. This leads to the following conditional distribution:

$$\mathbb{P}(Y = l \mid X = j, V = k) = \begin{cases} p, & j \neq k, j \neq l, k \neq l \\ 1 - p, & j \neq k, l = j \\ 0, & j = k, l = k \end{cases} \quad (13.6.3)$$

In particular, if $p = 1$, the player *always switches*, while if $p = 0$, the player *never switches*.

Mathematical Analysis

We are almost ready to analyze the Monty Hall problem mathematically. But first we must make some independence assumptions to incorporate the lack of knowledge that the host and player have about each other's actions. First, the player has no knowledge of the door containing the car, so we assume that U and X are independent. Also, the only information about the car door that the player has when she makes her second choice is the information (if any) revealed by her first choice and the host's subsequent selection. Mathematically, this means that Y is conditionally independent of U given X and V .

Distributions

The host and player strategies form the basic data for the Monty Hall problem. Because of the independence assumptions, the joint distribution of the basic random variables is completely determined by these strategies.

The joint probability density function of (U, X, V, Y) is given by

$$\mathbb{P}(U = i, X = j, V = k, Y = l) = \mathbb{P}(U = i)\mathbb{P}(X = j)\mathbb{P}(V = k | U = i, X = j)\mathbb{P}(Y = l | X = j, V = k), \quad i, j, k, l \in \{1, 2, 3\} \quad (13.6.4)$$

Proof

This follows from the independence assumptions and the multiplication rule of conditional probability.

The probability of any event defined in terms of the Monty Hall problem can be computed by summing the joint density over the appropriate values of (i, j, k, l) .

With either of the basic host strategies, V is uniformly distributed on $\{1, 2, 3\}$.

Suppose that the player switches with probability p . With either of the basic host strategies, Y is uniformly distributed on $\{1, 2, 3\}$.

In the Monty Hall experiment, set the host strategy to *standard*. For each of the following values of p , run the simulation 1000 times. Based on relative frequency, which strategy works best?

1. $p = 0$ (never switch)
2. $p = 0.3$
3. $p = 0.5$
4. $p = 0.7$
5. $p = 1$ (always switch)

In the Monty Hall experiment, set the host strategy to *blind*. For each of the following values of p , run the experiment 1000 times. Based on relative frequency, which strategy works best?

1. $p = 0$ (never switch)
2. $p = 0.3$
3. $p = 0.5$
4. $p = 0.7$
5. $p = 1$ (always switch)

The Probability of Winning

The event that the player wins a game is $\{Y = U\}$. We will compute the probability of this event with the basic host and player strategies.

Suppose that the host follows the standard strategy and that the player switches with probability p . Then the probability that the player wins is

$$\mathbb{P}(Y = U) = \frac{1+p}{3} \quad (13.6.5)$$

In particular, if the player always switches, the probability that she wins is $p = \frac{2}{3}$ and if the player never switches, the probability that she wins is $p = \frac{1}{3}$.

In the Monty Hall experiment, set the host strategy to *standard*. For each of the following values of p , run the simulation 1000 times. In each case, compare the relative frequency of winning to the probability of winning.

1. $p = 0$ (never switch)
2. $p = 0.3$
3. $p = 0.5$
4. $p = 0.7$
5. $p = 1$ (always switch)

Suppose that the host follows the blind strategy. Then for *any* player strategy, the probability that the player wins is

$$\mathbb{P}(Y = U) = \frac{1}{3} \quad (13.6.6)$$

In the Monty Hall experiment, set the host strategy to *blind*. For each of the following values of p , run the experiment 1000 times. In each case, compare the relative frequency of winning to the probability of winning.

1. $p = 0$ (never switch)
2. $p = 0.3$
3. $p = 0.5$
4. $p = 0.7$
5. $p = 1$ (always switch)

For a complete solution of the Monty Hall problem, we want to compute the conditional probability that the player wins, given that the host opens a door with a goat behind it:

$$\mathbb{P}(Y = U \mid V \neq U) = \frac{\mathbb{P}(Y = U)}{\mathbb{P}(V \neq U)} \quad (13.6.7)$$

With the basic host and player strategies, the numerator, the probability of winning, has been computed. Thus we need to consider the denominator, the probability that the host opens a door with a goat. If the host use the standard strategy, then the conditional probability of winning is the same as the unconditional probability of winning, regardless of the player strategy. In particular, we have the following result:

If the host follows the standard strategy and the player switches with probability p , then

$$\mathbb{P}(Y = U \mid V \neq U) = \frac{1+p}{3} \quad (13.6.8)$$

Proof

This follows from the [win probability above](#)

Once again, the probability increases from $\frac{1}{3}$ when $p = 0$, so that the player never switches, to $\frac{2}{3}$ when $p = 1$, so that the player always switches.

If the host follows the blind strategy, then for any player strategy, $\mathbb{P}(V \neq U) = \frac{2}{3}$ and therefore $\mathbb{P}(Y = U \mid V \neq U) = \frac{1}{2}$.

In the Monty Hall experiment, set the host strategy to *blind*. For each of the following values of p , run the experiment 500 times. In each case, compute the conditional relative frequency of winning, given that the host shows a goat, and compare with the theoretical answer above,

1. $p = 0$ (never switch)
2. $p = 0.3$
3. $p = 0.5$
4. $p = 0.7$
5. $p = 1$ (always switch)

The confusion between the conditional probability of winning for these two strategies has been the source of much controversy in the Monty Hall problem. Marilyn was probably thinking of the standard host strategy, while some of her critics were thinking of the blind strategy. This problem points out the importance of careful modeling, of the careful statement of assumptions. Marilyn is correct if the host follows the standard strategy; the critics are correct if the host follows the blind strategy; any number of other answers could be correct if the host follows other strategies.

The mathematical formulation we have used is fairly complete. However, if we just want to solve Marilyn's problem, there is a much simpler analysis (which you may have discovered yourself). Suppose that the host follows the standard strategy, and thus always opens a door with a goat. If the player's first door is incorrect (contains a goat), then the host has no choice and must open the other door with a goat. Then, if the player switches, she wins. On the other hand, if the player's first door is correct and she switches, then of course she loses. Thus, we see that if the player always switches, then *she wins if and only if her first choice is incorrect*, an event that obviously has probability $\frac{2}{3}$. If the player never switches, then she wins if and only if her first choice is correct, an event with probability $\frac{1}{3}$.

This page titled [13.6: The Monty Hall Problem](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.7: Lotteries

You realize the odds of winning [the lottery] are the same as being mauled by a polar bear and a regular bear in the same day.

—E*TRADE baby, January 2010.

Lotteries are among the simplest and most widely played of all games of chance, and unfortunately for the gambler, among the worst in terms of expected value. Lotteries come in such an incredible number of variations that it is impractical to analyze all of them. So, in this section, we will study some of the more common lottery formats.




Figure 13.7.1: A lottery ticket issued by the Continental Congress in 1776 to raise money for the American Revolutionary War.
Source: [Wikipedia](#)

The Basic Lottery

Basic Format

The basic lottery is a random experiment in which the gambling house (in many cases a government agency) selects n numbers at random, without replacement, from the integers from 1 to N . The integer parameters N and n vary from one lottery to another, and of course, n cannot be larger than N . The order in which the numbers are chosen usually does not matter, and thus in this case, the sample space S of the experiment consists of all subsets (combinations) of size n chosen from the population $\{1, 2, \dots, N\}$.

$$S = \{\mathbf{x} \subseteq \{1, 2, \dots, N\} : \#(\mathbf{x}) = n\} \quad (13.7.1)$$

Recall that

$$\#(S) = \binom{N}{n} = \frac{N!}{n!(N-n)!} \quad (13.7.2)$$

Naturally, we assume that all such combinations are equally likely, and thus, the chosen combination \mathbf{X} , the basic random variable of the experiment, is uniformly distributed on S .

$$\mathbb{P}(\mathbf{X} = \mathbf{x}) = \frac{1}{\binom{N}{n}}, \quad \mathbf{x} \in S \quad (13.7.3)$$

The player of the lottery pays a fee and gets to select m numbers, without replacement, from the integers from 1 to N . Again, order does not matter, so the player essentially chooses a combination \mathbf{y} of size m from the population $\{1, 2, \dots, N\}$. In many cases $m = n$, so that the player gets to choose the same number of numbers as the house. In general then, there are three parameters in the *basic* (N, n, m) lottery.

The player's goal, of course, is to maximize the number of matches (often called *catches* by gamblers) between her combination \mathbf{y} and the random combination \mathbf{X} chosen by the house. Essentially, the player is trying to guess the outcome of the random experiment before it is run. Thus, let $U = \#(\mathbf{X} \cap \mathbf{y})$ denote the number of catches.

The number of catches U in the (N, n, m) , lottery has probability density function given by

$$\mathbb{P}(U = k) = \frac{\binom{m}{k} \binom{N-m}{n-k}}{\binom{N}{n}}, \quad k \in \{0, 1, \dots, m\} \quad (13.7.4)$$

The distribution of U is the hypergeometric distribution with parameters N , n , and m , and is studied in detail in the chapter on Finite Sampling Models. In particular, from this section, it follows that the mean and variance of the number of catches U are

$$\mathbb{E}(U) = n \frac{m}{N} \quad (13.7.5)$$

$$\text{var}(U) = n \frac{m}{N} \left(1 - \frac{m}{N}\right) \frac{N-n}{N-1} \quad (13.7.6)$$

Note that $\mathbb{P}(U = k) = 0$ if $k > n$ or $k < n + m - N$. However, in most lotteries, $m \leq n$ and N is much larger than $n + m$. In these common cases, the density function is positive for the values of k given in [above](#).

We will refer to the special case where $m = n$ as the (N, n) lottery; this is the case in most state lotteries. In this case, the probability density function of the number of catches U is

$$\mathbb{P}(U = k) = \frac{\binom{n}{k} \binom{N-n}{n-k}}{\binom{N}{n}}, \quad k \in \{0, 1, \dots, n\} \quad (13.7.7)$$

The mean and variance of the number of catches U in this special case are

$$\mathbb{E}(U) = \frac{n^2}{N} \quad (13.7.8)$$

$$\text{var}(U) = \frac{n^2(N-n)^2}{N^2(N-1)} \quad (13.7.9)$$

Explicitly give the probability density function, mean, and standard deviation of the number of catches in the $(47, 5)$ lottery.

Answer

$$\mathbb{E}(U) = 0.5319148936 \quad \text{sd}(U) = 0.6587832083$$

k	$\mathbb{P}(U = k)$
0	0.5545644253
1	0.3648450167
2	0.0748400034
3	0.0056130003
4	0.0001369024
5	0.0000006519

Explicitly give the probability density function, mean, and standard deviation of the number of catches in the $(49, 5)$ lottery.

Answer

$$\mathbb{E}(U) = 0.5102040816 \quad \text{sd}(U) = 0.6480462207$$

k	$\mathbb{P}(U = k)$
0	0.5695196981
1	0.3559498113
2	0.0694536217
3	0.0049609730
4	0.0001153715
5	0.0000005244

Explicitly give the probability density function, mean, and standard deviation of the number of catches in the $(47, 7)$ lottery.

Answer

$$\mathbb{E}(U) = 1.042553191 \text{sd}(U) = 0.8783776109$$

k	$\mathbb{P}(U = k)$
0	0.2964400642
1	0.4272224454
2	0.2197144005
3	0.0508598149
4	0.0054983583
5	0.0002604486
6	0.0000044521
7	0.0000000159

The analysis above was based on the assumption that the player's combination \mathbf{y} is selected deterministically. Would it matter if the player chose the combination in a random way? Thus, suppose that the player's selected combination \mathbf{Y} is a random variable taking values in S . (For example, in many lotteries, players can buy tickets with combinations randomly selected by a computer; this is typically known as *Quick Pick*). Clearly, \mathbf{X} and \mathbf{Y} must be independent, since the player (and her randomizing device) can have no knowledge of the winning combination \mathbf{X} . As you might guess, such randomization makes no difference.

Let U denote the number of catches in the (N, n, m) lottery when the player's combination \mathbf{Y} is a random variable, independent of the winning combination \mathbf{X} . Then U has the same distribution as in the [deterministic case](#) above.

Proof

This follows by conditioning on the value of \mathbf{Y} :

$$\mathbb{P}(U = k) = \sum_{\mathbf{y} \in S} \mathbb{P}(U = k \mid \mathbf{Y} = \mathbf{y}) \mathbb{P}(\mathbf{Y} = \mathbf{y}) = \sum_{\mathbf{y} \in S} \mathbb{P}(U = k) \mathbb{P}(\mathbf{Y} = \mathbf{y}) = \mathbb{P}(U = k) \quad (13.7.10)$$

There are many websites that publish data on the frequency of occurrence of numbers in various state lotteries. Some gamblers evidently feel that some numbers are luckier than others.

Given the assumptions and analysis above, do you believe that some numbers are luckier than others? Does it make any mathematical sense to study historical data for a lottery?

The prize money in most state lotteries depends on the sales of the lottery tickets. Typically, about 50% of the sales money is returned as prize money; the rest goes for administrative costs and profit for the state. The total prize money is divided among the winning tickets, and the prize for a given ticket depends on the number of catches U . For all of these reasons, it is impossible to give a simple mathematical analysis of the expected value of playing a given state lottery. Note however, that since the state keeps a fixed percentage of the sales, there is essentially no risk for the state.

From a pure gambling point of view, state lotteries are bad games. In most casino games, by comparison, 90% or more of the money that comes in is returned to the players as prize money. Of course, state lotteries should be viewed as a form of voluntary taxation, not simply as games. The profits from lotteries are typically used for education, health care, and other essential services. A discussion of the value and costs of lotteries from a *political and social* point of view (as opposed to a *mathematical* one) is beyond the scope of this project.

Bonus Numbers

Many state lotteries now augment the basic (N, n) , format with a *bonus number*. The bonus number T is selected from a specified set of integers, in addition to the combination \mathbf{X} , selected as before. The player likewise picks a bonus number s , in addition to a combination \mathbf{y} . The player's prize then depends on the number of catches U between \mathbf{X} and \mathbf{y} , as before, and in addition on

whether the player's bonus number s matches the random bonus number T chosen by the house. We will let I denote the indicator variable of this latter event. Thus, our interest now is in the joint distribution of (I, U) .

In one common format, the bonus number T is selected at random from the set of integers $\{1, 2, \dots, M\}$, independently of the combination \mathbf{X} of size n chosen from $\{1, 2, \dots, N\}$. Usually $M < N$. Note that with this format, the game is essentially two independent lotteries, one in the (N, n) , format and the other in the $(M, 1)$, format.

Explicitly compute the joint probability density function of (I, U) for the $(47, 5)$ lottery with independent bonus number from 1 to 27. This format is used in the *California lottery*, among others.

Answer

Joint distribution of (I, U)

$\mathbb{P}(I = i, U = k)$	$i = 0$	1
$k = 0$	0.5340250022	0.0205394232
1	0.3513322383	0.0135127784
2	0.0720681514	0.0027718520
3	0.0054051114	0.0002078889
4	0.0001318320	0.0000050705
5	0.0000006278	0.0000000241

Explicitly compute the joint probability density function of (I, U) for the $(49, 5)$ lottery with independent bonus number from 1 to 42. This format is used in the *Powerball lottery*, among others.

Answer

Joint distribution of (I, U)

$\mathbb{P}(I = i, U = k)$	$i = 0$	1
$k = 0$	0.5559597053	0.0135599928
1	0.3474748158	0.0084749955
2	0.0677999641	0.0016536577
3	0.0048428546	0.0001181184
4	0.0001126245	0.0000027469
5	0.0000005119	0.0000000125

In another format, the bonus number T is chosen from 1 to N , and is distinct from the numbers in the combination \mathbf{X} . To model this game, we assume that T is uniformly distributed on $\{1, 2, \dots, N\}$, and given $T = t$, \mathbf{X} is uniformly distributed on the set of combinations of size n chosen from $\{1, 2, \dots, N\} \setminus \{t\}$. For this format, the joint probability density function is harder to compute.

The probability density function of (I, U) is given by

$$\mathbb{P}(I = 1, U = k) = \frac{\binom{n}{k} \binom{N-1-n}{n-k}}{N \binom{N-1}{n}}, \quad k \in \{0, 1, \dots, n\} \quad (13.7.11)$$

$$\mathbb{P}(I = 0, U = k) = (N - n + 1) \frac{\binom{n}{k} \binom{N-1-n}{n-k}}{N \binom{N-1}{n}} + n \frac{\binom{n-1}{k} \binom{N-n}{n-k}}{N \binom{N-1}{n}}, \quad k \in \{0, 1, \dots, n\} \quad (13.7.12)$$

Proof

The second equation is obtained by conditioning on whether $T \in \{y_1, y_2, \dots, y_n\}$.

Explicitly compute the joint probability density function of (I, U) for the $(47, 7)$ lottery with bonus number chosen as described above. This format is used in the *Super 7 Canada lottery*, among others.

Keno

Keno is a lottery game played in casinos. For a fixed N (usually 80) and n (usually 20), the player can play a range of basic (N, n, m) games, as described in the first subsection. Typically, m ranges from 1 to 15, and the payoff depends on m and the number of catches U . In this section, you will compute the density function, mean, and standard deviation of the random payoff, based on a unit bet, for a typical keno game with $N = 80$, $n = 20$, and $m \in \{1, 2, \dots, 15\}$. The payoff tables are based on the keno game at the *Tropicana* casino in Atlantic City, New Jersey.

Recall that the probability density function of the number of catches U [above](#), is given by

$$\mathbb{P}(U = k) = \frac{\binom{m}{k} \binom{80-m}{20-k}}{\binom{80}{20}}, \quad k \in \{0, 1, \dots, m\} \quad (13.7.13)$$

The payoff table for $m = 1$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 1$

Catches	0	1
Payoff	0	3

Answer

Pick $m = 1$, $\mathbb{E}(V) = 0.75$, $\text{sd}(V) = 1.299038106$

v	$\mathbb{P}(V = v)$
0	0.75
3	0.25

The payoff table for $m = 2$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 2$

Catches	0	1	2
Payoff	0	0	12

Answer

Pick $m = 2$, $\mathbb{E}(V) = 0.7353943525$, $\text{sd}(V) = 5.025285956$

v	$\mathbb{P}(V = v)$
12	0.0601265822

The payoff table for $m = 3$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 3$

Catches	0	1	2	3
---------	---	---	---	---

Payoff	0	0	1	43
--------	---	---	---	----

Answer

Pick $m = 3$, $\mathbb{E}(V) = 0.7353943525$, $\text{sd}(V) = 5.025285956$

v	$\mathbb{P}(V = v)$
0	0.8473709834
1	0.1387536514
43	0.0138753651

The payoff table for $m = 4$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 4$

Catches	0	1	2	3	4
Payoff	0	0	1	3	130

Answer

Pick $m = 4$, $\mathbb{E}(V) = 0.7406201394$, $\text{sd}(V) = 7.198935911$

v	$\mathbb{P}(V = v)$
0	0.7410532505
1	0.2126354658
3	0.0432478914
130	0.0030633923

The payoff table for $m = 5$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 5$

Catches	0	1	2	3	4	5
Payoff	0	0	0	1	10	800

Answer

Pick $m = 5$, $\mathbb{E}(V) = 0.7207981892$, $\text{sd}(V) = 20.33532453$

v	$\mathbb{P}(V = v)$
0	0.9033276850
1	0.0839350523
10	0.0120923380
800	0.0006449247

The payoff table for $m = 6$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 6$

Catches	0	1	2	3	4	5	6
Payoff	0	0	0	1	4	95	1500

Answer

Pick $m = 6$, $\mathbb{E}(V) = 0.7315342885$, $\text{sd}(V) = 17.83831647$

v	$\mathbb{P}(V = v)$
0	0.8384179112
1	0.1298195475
4	0.0285379178
95	0.0030956385
1500	0.0001289849

The payoff table for $m = 7$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 7$

Catches	0	1	2	3	4	5	6	7
Payoff	0	0	0	0	1	25	350	8000

Answer

Pick $m = 7$, $\mathbb{E}(V) = 0.7196008747$, $\text{sd}(V) = 40.69860455$

v	$\mathbb{P}(V = v)$
0	0.9384140492
1	0.0521909668
25	0.0086385048
350	0.0007320767
8000	0.0000244026

The payoff table for $m = 8$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 8$

Catches	0	1	2	3	4	5	6	7	8
Payoff	0	0	0	0	0	9	90	1500	25,000

Answer

Pick $m = 8$, $\mathbb{E}(V) = 0.7270517606$, $\text{sd}(V) = 55.64771986$

v	$\mathbb{P}(V = v)$
0	0.9791658999

9	0.0183025856
90	0.0023667137
1500	0.0001604552
25,000	0.0000043457

The payoff table for $m = 9$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 9$

Catches	0	1	2	3	4	5	6	7	8	9
Payoff	0	0	0	0	0	4	50	280	4000	50,000

Answer

Pick $m = 9$, $\mathbb{E}(V) = 0.7270517606$, $\text{sd}(V) = 55.64771986$

v	$\mathbb{P}(V = v)$
0	0.9791658999
9	0.0183025856
90	0.0023667137
1500	0.0001604552
25,000	0.0000043457

The payoff table for $m = 10$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 10$

Catches	0	1	2	3	4	5	6	7	8	9	10
Payoff	0	0	0	0	0	1	22	150	1000	5000	100,000

Answer

Pick $m = 10$, $\mathbb{E}(V) = 0.7228896221$, $\text{sd}(V) = 38.10367609$

v	$\mathbb{P}(V = v)$
0	0.9353401224
1	0.0514276877
22	0.0114793946
150	0.0016111431
1000	0.0001354194
5000	0.0000061206
100,000	0.0000001122

The payoff table for $m = 11$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 11$

Catches	0	1	2	3	4	5	6	7	8	9	10	11
Payoff	0	0	0	0	0	0	8	80	400	2500	25,000	100,000

Answer

Pick $m = 11$, $\mathbb{E}(V) = 0.7138083347$, $\text{sd}(V) = 32.99373346$

v	$\mathbb{P}(V = v)$
0	0.9757475913
8	0.0202037345
80	0.0036078097
400	0.0004114169
2500	0.0000283736
25,000	0.0000010580
100,000	0.0000000160

The payoff table for $m = 12$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 12$

Catches	0	1	2	3	4	5	6	7	8	9	10	11	12
Payoff	0	0	0	0	0	0	5	32	200	1000	5000	25,000	100,000

Answer

Pick $m = 12$, $\mathbb{E}(V) = 0.7167721544$, $\text{sd}(V) = 20.12030014$

v	$\mathbb{P}(V = v)$
0	0.9596431653
5	0.0322088520
32	0.0070273859
200	0.0010195984
1000	0.0000954010
5000	0.0000054280
25,000	0.0000001673
100,000	0.0000000021

The payoff table for $m = 13$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 13$

Catches	0	1	2	3	4	5	6	7	8	9	10	11	12	13
Payoff	1	0	0	0	0	0	1	20	80	600	3500	10,000	50,000	100,000

Proof

Pick $m = 13$, $\mathbb{E}(V) = 0.7216651326$ $\text{sd}(V) = 22.68311303$

v	$\mathbb{P}(V = v)$
0	0.9213238456
1	0.0638969375
20	0.0123151493
80	0.0021831401
600	0.0002598976
3500	0.0000200623
10,000	0.0000009434
50,000	0.0000000240
100,000	0.0000000002

The payoff table for $m = 14$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 14$

Catches	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Payoff	1	0	0	0	0	0	1	9	42	310	1100	8000	25,000	50,000	100,000

Answer

Pick $m = 14$, $\mathbb{E}(V) = 0.7194160496$ $\text{sd}(V) = 21.98977077$

v	$\mathbb{P}(V = v)$
0	0.898036333063
1	0.077258807301
9	0.019851285448
42	0.004181636518
310	0.000608238039
1100	0.000059737665
8000	0.000003811015
25,000	0.000000147841
50,000	0.000000003084

100,000

0.000000000026

The payoff table for $m = 15$ is given below. Compute the probability density function, mean, and standard deviation of the payoff.

Pick $m = 15$

Catches	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Payoff	1	0	0	0	0	0	0	10	25	100	300	2800	25,000	50,000	100,000	100,000

Answer

Pick $m = 15$, $\mathbb{E}(V) = 0.7144017020$ $\text{sd}(V) = 24.31901706$

v	$\mathbb{P}(V = v)$
0	0.95333046038902
1	0.00801614417729
10	0.02988971956684
25	0.00733144064847
100	0.00126716258122
300	0.00015205950975
2800	0.00001234249267
25,000	0.00000064960488
50,000	0.00000002067708
100,000	0.00000000035046
100,000	0.00000000000234

In the exercises above, you should have noticed that the expected payoff on a unit bet varies from about 0.71 to 0.75, so the expected profit (for the gambler) varies from about -0.25 to -0.29 . This is quite bad for the gambler playing a casino game, but as always, the lure of a very high payoff on a small bet for an extremely rare event overrides the expected value analysis for most players.

With $m = 15$, show that the top 4 prizes (25,000, 50,000, 100,000, 100,000) contribute only about 0.017 (less than 2 cents) to the total expected value of about 0.714.

On the other hand, the standard deviation of the payoff varies quite a bit, from about 1 to about 55.

Although the game is highly unfavorable for each m , with expected value that is nearly constant, which do you think is better for the gambler—a format with high standard deviation or one with low standard deviation?

This page titled [13.7: Lotteries](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.8: The Red and Black Game

In this section and the following three sections, we will study gambling strategies for one of the simplest gambling models. Yet in spite of the simplicity of the model, the mathematical analysis leads to some beautiful and sometimes surprising results that have importance and application well beyond gambling. Our exposition is based primarily on the classic book by Dubbins and Savage, *Inequalities for Stochastic Processes (How to Gamble if You Must)* by Lester E Dubbins and Leonard J Savage (1965).

Basic Theory

Assumptions

Here is the basic situation: The gambler starts with an initial sum of money. She bets on independent, probabilistically identical games, each with two outcomes—win or lose. If she wins a game, she receives the amount of the bet on that game; if she loses a game, she must pay the amount of the bet. Thus, the gambler plays at *even stakes*. This particular situation (IID games and even stakes) is known as *red and black*, and is named for the color bets in the casino game roulette. Other examples are the *pass* and *don't pass* bets in craps.

Let us try to formulate the gambling experiment mathematically. First, let I_n denote the outcome of the n th game for $n \in \mathbb{N}_+$, where 1 denotes a win and 0 denotes a loss. These are independent indicator random variables with the same distribution:

$$\mathbb{P}(I_j = 1) = p_j, \quad \mathbb{P}(I_j = 0) = q = 1 - p_j \quad (13.8.1)$$

where $p \in [0, 1]$ is the probability of winning an individual game. Thus, $\mathbf{I} = (I_1, I_2, \dots)$ is a sequence of Bernoulli trials.

If $p = 0$, then the gambler always loses and if $p = 1$ then the gambler always wins. These trivial cases are not interesting, so we will usually assume that $0 < p < 1$. In real gambling houses, of course, $p < \frac{1}{2}$ (that is, the games are unfair to the player), so we will be particularly interested in this case.

Random Processes

The gambler's fortune over time is the basic random process of interest: Let X_0 denote the gambler's initial fortune and X_i the gambler's fortune after i games. The gambler's strategy consists of the decisions of how much to bet on the various games and when to quit. Let Y_i denote the amount of the i th bet, and let N denote the number of games played by the gambler. If we want to, we can always assume that the games go on forever, but with the assumption that the gambler bets 0 on all games after N . With this understanding, the game outcome, fortune, and bet processes are defined for all times $i \in \mathbb{N}_+$.

The fortune process is related to the wager process as follows:

$$X_j = X_{j-1} + (2I_j - 1)Y_j, \quad j \in \mathbb{N}_+ \quad (13.8.2)$$

Strategies

The gambler's strategy can be very complicated. For example, the random variable Y_n , the gambler's bet on game n , or the event $N = n - 1$, her decision to stop after $n - 1$ games, could be based on the entire past history of the game, up to time n . Technically, this history forms a σ -algebra:

$$\mathcal{H}_n = \sigma\{X_0, Y_1, I_1, Y_2, I_2, \dots, Y_{n-1}, I_{n-1}\} \quad (13.8.3)$$

Moreover, they could have additional sources of randomness. For example a gambler playing roulette could partly base her bets on the roll of a lucky die that she keeps in her pocket. However, the gambler cannot see into the future (unfortunately from her point of view), so we can at least assume that Y_n and $\{N = n - 1\}$ are independent of $(I_1, I_2, \dots, I_{n-1})$.

At least in terms of expected value, any gambling strategy is futile if the games are unfair.

$$\mathbb{E}(X_i) = \mathbb{E}(X_{i-1}) + (2p - 1)\mathbb{E}(Y_i) \quad \text{for } i \in \mathbb{N}_+$$

Proof

This follows from the [previous result](#) and the assumption of no prescience.

Suppose that the gambler has a positive probability of making a real bet on game i , so that $\mathbb{E}(Y_i) > 0$. Then

1. $\mathbb{E}(X_i) < \mathbb{E}(X_{i-1})$ if $p < \frac{1}{2}$
2. $\mathbb{E}(X_i) > \mathbb{E}(X_{i-1})$ if $p > \frac{1}{2}$
3. $\mathbb{E}(X_i) = \mathbb{E}(X_{i-1})$ if $p = \frac{1}{2}$

Proof

This follows from the [previous result](#) on the expected value of X_i .

Thus on any game in which the gambler makes a positive bet, her expected fortune strictly decreases if the games are unfair, remains the same if the games are fair, and strictly increases if the games are favorable.

As we noted earlier, a general strategy can depend on the past history and can be randomized. However, since the underlying Bernoulli games are independent, one might guess that these complicated strategies are no better than simple strategies in which the amount of the bet and the decision to stop are based only on the gambler's current fortune. These simple strategies do indeed play a fundamental role and are referred to as *stationary, deterministic strategies*. Such a strategy can be described by a *betting function* S from the space of fortunes to the space of allowable bets, so that $S(x)$ is the amount that the gambler bets when her current fortune is x .

The Stopping Rule

From now on, we will assume that the gambler's stopping rule is a very simple and standard one: she will bet on the games until she either loses her entire fortune and is ruined or reaches a fixed target fortune a :

$$N = \min\{n \in \mathbb{N} : X_n = 0 \text{ or } X_n = a\} \quad (13.8.4)$$

Thus, any strategy (betting function) S must satisfy $s(x) \leq \min\{x, a - x\}$ for $0 \leq x \leq a$: the gambler cannot bet what she does not have, and will not bet more than is necessary to reach the target a .

If we want to, we can think of the difference between the target fortune and the initial fortune as the entire fortune of the house. With this interpretation, the player and the house play symmetric roles, but with complementary win probabilities: play continues until either the player is ruined or the house is ruined. Our main interest is in the final fortune X_N of the gambler. Note that this random variable takes just two values; 0 and a .

The mean and variance of the final fortune are given by

1. $\mathbb{E}(X_N) = a\mathbb{P}(X_N = a)$
2. $\text{var}(X_N) = a^2\mathbb{P}(X_N = a)[1 - \mathbb{P}(X_N = a)]$

Presumably, the gambler would like to maximize the probability of reaching the target fortune. Is it better to bet small amounts or large amounts, or does it not matter? How does the optimal strategy, if there is one, depend on the initial fortune, the target fortune, and the game win probability?

We are also interested in $\mathbb{E}(N)$, the expected number of games played. Perhaps a secondary goal of the gambler is to maximize the expected number of games that she gets to play. Are the two goals compatible or incompatible? That is, can the gambler maximize both her probability of reaching the target *and* the expected number of games played, or does maximizing one quantity necessarily mean minimizing the other?

In the next two sections, we will analyze and compare two strategies that are in a sense opposites:

- **Timid Play:** On each game, until she stops, the gambler makes a small constant bet, say \$1.
- **Bold Play:** On each game, until she stops, the gambler bets either her entire fortune or the amount needed to reach the target fortune, whichever is smaller.

In the final section of the chapter, we will return to the question of optimal strategies.

This page titled [13.8: The Red and Black Game](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.9: Timid Play

Basic Theory

Recall that with the strategy of *timid play* in red and black, the gambler makes a small constant bet, say \$1, on each game until she stops. Thus, on each game, the gambler's fortune either increases by 1 or decreases by 1, until the fortune reaches either 0 or the target a (which we assume is a positive integer). Thus, the fortune process (X_0, X_1, \dots) is a *random walk* on the fortune space $\{0, 1, \dots, a\}$ with 0 and a as *absorbing barriers*.

As usual, we are interested in the probability of winning and the expected number of games. The key idea in the analysis is that after each game, the fortune process simply starts over again, but with a different initial value. This is an example of the *Markov property*, named for Andrei Markov. A separate chapter on Markov Chains explores these random processes in more detail. In particular, this chapter has sections on Birth-Death Chains and Random Walks on Graphs, particular classes of Markov chains that generalize the random processes that we are studying here.



Figure 13.9.1: The transition graph for timid play

The Probability of Winning

Our analysis based on the Markov property suggests that we treat the initial fortune as a variable. Thus, we will denote the probability that the gambler reaches the target a , starting with an initial fortune x by

$$f(x) = \mathbb{P}(X_N = a \mid X_0 = x), \quad x \in \{0, 1, \dots, a\} \quad (13.9.1)$$

The function f satisfies the following difference equation and boundary conditions:

1. $f(x) = qf(x-1) + pf(x+1)$ for $x \in \{1, 2, \dots, a-1\}$
2. $f(0) = 0, f(a) = 1$

Proof

The boundary conditions are just a matter of definition. The difference equation follows from conditioning on the outcome of the first trial. She loses this trial with probability q and if she loses, then effectively she starts a new sequence of trials but with initial fortune $x-1$. She wins the first trial with probability p , and if she wins, then she effectively starts a new sequence of trials but with initial fortune $x+1$.

The difference equation is *linear* (in the unknown function f), *homogeneous* (because each term involves the unknown function f), and *second order* (because 2 is the difference between the largest and smallest fortunes in the equation). Recall that linear homogeneous difference equations can be solved by finding the roots of the characteristic equation.

The characteristic equation of the [difference equation](#) is $pr^2 - r + q = 0$, and that the roots are $r = 1$ and $r = q/p$.

If $p \neq \frac{1}{2}$, then the roots are distinct. In this case, the probability that the gambler reaches her target is

$$f(x) = \frac{(q/p)^x - 1}{(q/p)^a - 1}, \quad x \in \{0, 1, \dots, a\} \quad (13.9.2)$$

If $p = \frac{1}{2}$, the characteristic equation has a single root 1 that has multiplicity 2. In this case, the probability that the gambler reaches her target is simply the ratio of the initial fortune to the target fortune:

$$f(x) = \frac{x}{a}, \quad x \in \{0, 1, \dots, a\} \quad (13.9.3)$$

Thus, we have the distribution of the final fortune X_N in either case:

$$\mathbb{P}(X_N = 0 \mid X_0 = x) = 1 - f(x), \quad \mathbb{P}(X_N = a \mid X_0 = x) = f(x); \quad x \in \{0, 1, \dots, a\} \quad (13.9.4)$$

In the red and black experiment, choose *Timid Play*. Vary the initial fortune, target fortune, and game win probability and note how the probability of winning the game changes. For various values of the parameters, run the experiment 1000 times and compare the relative frequency of winning a game to the probability of winning a game.

As a function of x , for fixed p and a ,

1. f is increasing from 0 to a .
2. f is concave upward if $p < \frac{1}{2}$ and concave downward if $p > \frac{1}{2}$. Of course, f is linear if $p = \frac{1}{2}$.

f is continuous as a function of p , for fixed x and a .

Proof

An application of L'Hospital's Rule shows that the [probability of winning when \$p \neq \frac{1}{2}\$](#) converges to the [probability of winning when \$p = \frac{1}{2}\$](#) , as $p \rightarrow \frac{1}{2}$.

For fixed x and a , $f(x)$ increases from 0 to 1 as p increases from 0 to 1.




Figure 13.9.2: The graph of f for $p = 0.4$, $p = 0.5$, and $p = 0.6$

The Expected Number of Trials

Now let us consider the expected number of games needed with timid play, when the initial fortune is x :

$$g(x) = \mathbb{E}(N \mid X_0 = x), \quad x \in \{0, 1, \dots, a\} \quad (13.9.5)$$

The function g satisfies the following difference equation and boundary conditions:

1. $g(x) = qg(x-1) + pg(x+1) + 1$ for $x \in \{1, 2, \dots, a-1\}$
2. $g(0) = 0, g(a) = 0$

Proof

Again, the difference equation follows from conditioning on the first trial. She loses this trial with probability q and if she loses, then effectively she starts a new sequence of trials but with initial fortune $x-1$. She wins the first trial with probability p , and if she wins, then she effectively starts a new sequence of trials but with initial fortune $x+1$. In either case, one trial is over.

The difference equation in the last exercise is linear, second order, but non-homogeneous (because of the constant term 1 on the right side). The corresponding homogeneous equation is the equation satisfied by the win probability function f . Thus, only a little additional work is needed to solve the non-homogeneous equation.

If $p \neq \frac{1}{2}$, then

$$g(x) = \frac{x}{q-p} - \frac{a}{q-p} f(x), \quad x \in \{0, 1, \dots, a\} \quad (13.9.6)$$

where f is the [win probability function](#) above.

If $p = \frac{1}{2}$, then

$$g(x) = x(a-x), \quad x \in \{0, 1, \dots, a\} \quad (13.9.7)$$

Consider g as a function of the initial fortune x , for fixed values of the game win probability p and the target fortune a .

1. g at first increases and then decreases.
2. g is concave downward.

When $p = \frac{1}{2}$, the maximum value of g is $\frac{a^2}{4}$ and occurs when $x = \frac{a}{2}$. When $p \neq \frac{1}{2}$, the value of x where the maximum occurs is rather complicated.




Figure 13.9.3: The graph of g for $p = 0.4$, $p = 0.5$, and $p = 0.6$

g is continuous as a function of p , for fixed x and a .

Proof

The [expected value](#) when $p \neq \frac{1}{2}$ converges to the [expected value](#) when $p = \frac{1}{2}$, as $p \rightarrow \frac{1}{2}$.

For many parameter settings, the expected number of games is surprisingly large. For example, suppose that $p = \frac{1}{2}$ and the target fortune is 100. If the gambler's initial fortune is 1, then the expected number of games is 99, even though half of the time, the gambler will be ruined on the first game. If the initial fortune is 50, the expected number of games is 2500.

In the red and black experiment, select *Timid Play*. Vary the initial fortune, the target fortune and the game win probability and notice how the expected number of games changes. For various values of the parameters, run the experiment 1000 times and compare the sample mean number of games to the expect value.

Increasing the Bet

What happens if the gambler makes constant bets, but with an amount higher than 1? The answer to this question may give insight into what will happen with bold play.

In the red and black game, set the target fortune to 16, the initial fortune to 8, and the win probability to 0.45. Play 10 games with each of the following strategies. Which seems to work best?

1. Bet 1 on each game (timid play).
2. Bet 2 on each game.
3. Bet 4 on each game.
4. Bet 8 on each game (bold play).

We will need to embellish our notation to indicate the dependence on the target fortune. Let

$$f(x, a) = \mathbb{P}(X_N = a \mid X_0 = x), \quad x \in \{0, 1, \dots, a\}, \quad a \in \mathbb{N}_+ \quad (13.9.8)$$

Now fix p and suppose that the target fortune is $2a$ and the initial fortune is $2x$. If the gambler plays timidly (betting \$1 each time), then of course, her probability of reaching the target is $f(2x, 2a)$. On the other hand:

Suppose that the gambler bets \$2 on each game. The fortune process $(X_i/2 : i \in \mathbb{N})$ corresponds to timid play with initial fortune x and target fortune a and that therefore the probability that the gambler reaches the target is $f(x, a)$.

Thus, we need to compare the probabilities $f(2x, 2a)$ and $f(x, a)$.

The win probability functions are related as follows:

$$f(2x, 2a) = f(x, a) \frac{(q/p)^x + 1}{(q/p)^a + 1}, \quad x \in \{0, 1, \dots, a\} \quad (13.9.9)$$

In particular

1. $f(2x, 2a) < f(x, a)$ if $p < \frac{1}{2}$
2. $f(2x, 2a) = f(x, a)$ if $p = \frac{1}{2}$
3. $f(2x, 2a) > f(x, a)$ if $p > \frac{1}{2}$

Thus, it appears that increasing the bets is a good idea if the games are unfair, a bad idea if the games are favorable, and makes no difference if the games are fair.

What about the expected number of games played? It seems almost obvious that if the bets are increased, the expected number of games played should decrease, but a direct analysis using the [expected value function](#) above is harder than one might hope (try it!). We will use a different method, one that actually gives better results. Specifically, we will have the \$1 and \$2 gamblers bet on the same underlying sequence of games, so that the two fortune processes are defined on the same sample space. Then we can compare the actual random variables (the number of games played), which in turn leads to a comparison of their expected values. Recall that this general method is referred to as *coupling*.

Let X_n denote the fortune after n games for the gambler making \$1 bets (simple timid play). Then $2X_n - X_0$ is the fortune after n games for the gambler making \$2 bets (with the same initial fortune, betting on the same sequence of games). Assume again that the initial fortune is $2x$ and the target fortune $2a$ where $0 < x < a$. Let N_1 denote the number of games played by the \$1 gambler, and N_2 the number of games played by the \$2 gambler. Then

1. If the \$1 gambler falls to fortune x , the \$2 gambler is ruined (fortune 0).
2. If the \$1 gambler hits fortune $x + a$, the \$2 gambler reaches the target $2a$.
3. The \$1 gambler must hit x before hitting 0 and must hit $x + a$ before hitting $2a$.
4. $N_2 < N_1$ given $X_0 = 2x$.
5. $\mathbb{E}(N_2 \mid X_0 = 2x) < \mathbb{E}(N_1 \mid X_0 = 2x)$

Of course, the expected values agree (and are both 0) if $x = 0$ or $x = a$. This result shows that N_2 is *stochastically smaller* than N_1 when the gamblers are not playing the same sequence of games (so that the random variables are not defined on the same sample space).

Generalize the analysis in this subsection to compare timid play with the strategy of betting \$ k on each game (let the initial fortune be kx and the target fortune ka).

It appears that with unfair games, the larger the bets the better, at least in terms of the probability of reaching the target. Thus, we are naturally led to consider bold play.

This page titled [13.9: Timid Play](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

13.10: Bold Play

Basic Theory

Preliminaries

Recall that with the strategy of *bold play* in red and black, the gambler on each game bets either her entire fortune or the amount needed to reach the target fortune, whichever is smaller. As usual, we are interested in the probability that the player reaches the target and the expected number of trials. The first interesting fact is that only the ratio of the initial fortune to the target fortune matters, quite in contrast to timid play.

Suppose that the gambler plays boldly with initial fortune x and target fortune a . As usual, let $\mathbf{X} = (X_1, X_2, \dots)$ denote the fortune process for the gambler. For any $c > 0$, the random process $c\mathbf{X} = (cX_0, cX_1, \dots)$ is the fortune process for bold play with initial fortune cx and target fortune ca .

Because of this result, it is convenient to use the target fortune as the monetary unit and to allow irrational, as well as rational, initial fortunes. Thus, the fortune space is $[0, 1]$. Sometimes in our analysis we will ignore the states 0 or 1; clearly there is no harm in this because in these states, the game is over.

Recall that the betting function S is the function that gives the amount bet as a function of the current fortune. For bold play, the betting function is

$$S(x) = \min\{x, 1 - x\} = \begin{cases} x, & 0 \leq x \leq \frac{1}{2} \\ 1 - x, & \frac{1}{2} \leq x \leq 1 \end{cases} \quad (13.10.1)$$




Figure 13.10.1: The betting function for bold play

The Probability of Winning

We will denote the probability that the bold gambler reaches the target $a = 1$ starting from the initial fortune $x \in [0, 1]$ by $F(x)$. By the [scaling property](#), the probability that the bold gambler reaches some other target value $a > 0$, starting from $x \in [0, a]$ is $F(x/a)$.

The function F satisfies the following functional equation and boundary conditions:

1. $F(x) = \begin{cases} pF(2x), & 0 \leq x \leq \frac{1}{2} \\ p + qF(2x - 1), & \frac{1}{2} \leq x \leq 1 \end{cases}$
2. $F(0) = 0, F(1) = 1$

From the previous result, and a little thought, it should be clear that an important role is played by the following function:

Let d be the function defined on $[0, 1)$ by

$$d(x) = 2x - \lfloor 2x \rfloor = \begin{cases} 2x, & 0 \leq x < \frac{1}{2} \\ 2x - 1, & \frac{1}{2} \leq x < 1 \end{cases} \quad (13.10.2)$$

The function d is called the *doubling function, mod 1*, since $d(x)$ gives the fractional part of $2x$.

Note that until the last bet that ends the game (with the player ruined or victorious), the successive fortunes of the player follow iterates of the map d . Thus, bold play is intimately connected with the *dynamical system* associated with d .




Figure 13.10.2: The doubling map, modulo 1

Binary Expansions

One of the keys to our analysis is to represent the initial fortune in binary form.

The *binary expansion* of $x \in [0, 1)$ is

$$x = \sum_{i=1}^{\infty} \frac{x_i}{2^i} \quad (13.10.3)$$

where $x_i \in \{0, 1\}$ for each $i \in \mathbb{N}_+$. This representation is unique except when x is a *binary rational* (sometimes also called a *dyadic rational*), that is, a number of the form $k/2^n$ where $n \in \mathbb{N}_+$ and $k \in \{1, 3, \dots, 2^n - 1\}$; the positive integer n is called the *rank* of x . Binary rationals are discussed in more detail in the chapter on Foundations.

For a binary rational x of rank n , we will use the standard terminating representation where $x_n = 1$ and $x_i = 0$ for $i > n$. Rank can be extended to all numbers in $[0, 1)$ by defining the rank of 0 to be 0 (0 is also considered a binary rational) and by defining the rank of a binary irrational to be ∞ . We will denote the rank of x by $r(x)$.

Applied to the binary sequences, the doubling function d is the *shift operator*:

For $x \in [0, 1)$, $[d(x)]_i = x_{i+1}$.

Bold play in red and black can be elegantly described by comparing the bits of the initial fortune with the game bits.

Suppose that gambler starts with initial fortune $x \in (0, 1)$. The gambler eventually reaches the target 1 if and only if there exists a positive integer k such that $I_j = 1 - x_j$ for $j \in \{1, 2, \dots, k-1\}$ and $I_k = x_k$. That is, the gambler wins if and only if when the game bit agrees with the corresponding fortune bit for the first time, that bit is 1.

The random variable whose bits are the complements of the fortune bits will play an important role in our analysis. Thus, let

$$W = \sum_{j=1}^{\infty} \frac{1 - I_j}{2^j} \quad (13.10.4)$$

Note that W is a well defined random variable taking values in $[0, 1]$.

Suppose that the gambler starts with initial fortune $x \in (0, 1)$. Then the gambler reaches the target 1 if and only if $W < x$.

Proof

This follows from the [previous result](#).

W has a continuous distribution. That is, $\mathbb{P}(W = x) = 0$ for any $x \in [0, 1]$.

From the previous two results, it follows that F is simply the distribution function of W . In particular, F is an increasing function, and since W has a continuous distribution, F is a continuous function.

The success function F is the unique continuous solution of the [functional equation](#) above.

Proof

Induction on the rank shows that any two solutions must agree at the binary rationals. But then any two continuous solutions must agree for all $x \in [0, 1]$.

If we introduce a bit more notation, we can give nice expression for $F(x)$, and later for the expected number of games $G(x)$. Let $p_0 = p$ and $p_1 = q = 1 - p$.

The win probability function F can be expressed as follows:

$$F(x) = \sum_{n=1}^{\infty} p_{x_1} \cdots p_{x_{n-1}} p_{x_n} \quad (13.10.5)$$

Note that $p \cdot x_n$ in the last expression is correct; it's not a misprint of p_{x_n} . Thus, only terms with $x_n = 1$ are included in the sum.

F is strictly increasing on $[0, 1]$. This means that the distribution of W has support $[0, 1]$; that is, there are no subintervals of $[0, 1]$ that have positive length, but 0 probability.

In particular,

1. $F\left(\frac{1}{8}\right) = p^3$
2. $F\left(\frac{2}{8}\right) = p^2$
3. $F\left(\frac{3}{8}\right) = p^2 + p^2 q$
4. $F\left(\frac{4}{8}\right) = p$
5. $F\left(\frac{5}{8}\right) = p + p^2 q$
6. $F\left(\frac{6}{8}\right) = p + p q$
7. $F\left(\frac{7}{8}\right) = p + p q + p q^2$

If $p = \frac{1}{2}$ then $F(x) = x$ for $x \in [0, 1]$

Proof

There are two proofs. The simplest proof is to note that $x \mapsto x$ is continuous and satisfies the functional equation in [functional equation](#). Another proof can be constructed by using the representation of F as a [sum](#).

Thus, for $p = \frac{1}{2}$ (fair trials), the probability that the bold gambler reaches the target fortune a starting from the initial fortune x is x/a , just as it is for the timid gambler. Note also that the random variable W has the uniform distribution on $[0, 1]$. When $p \neq \frac{1}{2}$, the distribution of W is quite strange. To state the result succinctly, we will indicate the dependence of the of the probability measure \mathbb{P} on the parameter $p \in (0, 1)$. First we define

$$C_p = \left\{ x \in [0, 1] : \frac{1}{n} \sum_{i=1}^n (1 - x_i) \rightarrow p \text{ as } n \rightarrow \infty \right\} \quad (13.10.6)$$

Thus, C_p is the set of $x \in [0, 1]$ for which the relative frequency of 0's in the binary expansion is p .

For distinct $p, t \in (0, 1)$

1. $\mathbb{P}_p(W \in C_p) = 1$
2. $\mathbb{P}_p(W \in C_t) = 0$

Proof

Part (a) follows from the strong law of large numbers. Part (b) follows from part (a) since $C_p \cap C_t = \emptyset$.

When $p \neq \frac{1}{2}$, W does not have a probability density function (with respect to Lebesgue measure on $[0, 1]$), even though W has a continuous distribution.

Proof

The proof is by contradiction. Suppose that W has probability density function f . Then $1 = \mathbb{P}_p(W \in C_p) = \int_{C_p} f(x) dx$. But if $p \neq \frac{1}{2}$, $\int_{C_p} 1 dx = \mathbb{P}_{1/2}(W \in C_p) = 0$. That is, C_p has Lebesgue measure 0. But then $\int_{C_p} f(x) dx = 0$, a contradiction.

When $p \neq \frac{1}{2}$, F has derivative 0 at almost every point in $[0, 1]$, even though it is strictly increasing.




Figure 13.10.3: The graphs of F when $p = 0.5$, $p = 0.4$, and $p = 0.2$

In the red and black experiment, select *Bold Play*. Vary the initial fortune, target fortune, and game win probability with the scroll bars and note how the probability of winning the game changes. In particular, note that this probability depends only on x/a . Now for various values of the parameters, run the experiment 1000 times and compare the relative frequency function to the probability density function.

The Expected Number of Trials

Let $G(x) = \mathbb{E}(N \mid X_0 = x)$ for $x \in [0, 1]$, the expected number of trials starting at x . For any other target fortune $a > 0$, the expected number of trials starting at $x \in [0, a]$ is just $G(x/a)$.

G satisfies the following functional equation and boundary conditions:

1. $G(x) = \begin{cases} 1 + pG(2x), & 0 < x \leq \frac{1}{2} \\ 1 + qG(2x - 1), & \frac{1}{2} \leq x < 1 \end{cases}$
2. $G(0) = 0, G(1) = 0$

Proof

The functional equation follows from conditioning on the result of the first game.

Note, interestingly, that the functional equation is not satisfied at $x = 0$ or $x = 1$. As before, we can give an alternate analysis using the binary representation of an initial fortune $x \in (0, 1)$.

Suppose that the initial fortune of the gambler is $x \in (0, 1)$. Then $N = \min\{k \in \mathbb{N}_+ : I_k = x_k \text{ or } k = r(x)\}$.

Proof

If x is a binary rational then N takes values in the set $\{1, 2, \dots, r(x)\}$. Play continues until the game number agrees with the rank of the fortune or a game bit agrees with the corresponding fortune bit, whichever is smaller. In the first case, the penultimate fortune is $\frac{1}{2}$, the only fortune for which the next game is always final. If x is a binary irrational then N takes values in \mathbb{N}_+ . Play continues until a game bit agrees with a corresponding fortune bit.

We can give an explicit formula for the expected number of trials $G(x)$ in terms of the binary representation of x . Recall our special notation: $p_0 = p$, $p_1 = q = 1 - p$

Suppose that $x \in (0, 1)$. Then

$$G(x) = \sum_{n=0}^{r(x)-1} p_{x_1} \cdots p_{x_n} \quad (13.10.7)$$

Note that the $n = 0$ term is 1, since the product is empty. The sum has a finite number of terms if x is a binary rational, and the sum has an infinite number of terms if x is a binary irrational.

In particular,

1. $G\left(\frac{1}{8}\right) = 1 + p + p^2$
2. $G\left(\frac{2}{8}\right) = 1 + p$
3. $G\left(\frac{3}{8}\right) = 1 + p + pq$
4. $G\left(\frac{4}{8}\right) = 1$
5. $G\left(\frac{5}{8}\right) = 1 + q + pq$
6. $G\left(\frac{6}{8}\right) = 1 + q$
7. $G\left(\frac{7}{8}\right) = 1 + q + q^2$

If $p = \frac{1}{2}$ then

$$G(x) = \begin{cases} 2 - \frac{1}{2^{r(x)-1}}, & x \text{ is a binary rational} \\ 2, & x \text{ is a binary irrational} \end{cases} \quad (13.10.8)$$




Figure 13.10.4: The expected number of games in bold play with fair games

In the red and black experiment, select **Bold Play**. Vary x , a , and p with the scroll bars and note how the expected number of trials changes. In particular, note that the mean depends only on the ratio x/a . For selected values of the parameters, run the experiment 1000 times and compare the sample mean to the distribution mean.

For fixed x , G is continuous as a function of p .

However, as a function of the initial fortune x , for fixed p , the function G is very irregular.

G is discontinuous at the binary rationals in $[0, 1]$ and continuous at the binary irrationals.

This page titled 13.10: Bold Play is shared under a CC BY 2.0 license and was authored, remixed, and/or curated by Kyle Siegrist (Random Services) via source content that was edited to the style and standards of the LibreTexts platform.

13.11: Optimal Strategies

Basic Theory

Definitions

Recall that the stopping rule for red and black is to continue playing until the gambler is ruined or her fortune reaches the target fortune a . Thus, the gambler's strategy is to decide how much to bet on each game before she must stop. Suppose that we have a class of strategies that correspond to certain valid fortunes and bets; A will denote the set of fortunes and B_x will denote the set of valid bets for $x \in A$. For example, sometimes (as with timid play) we might want to restrict the fortunes to set of integers $\{0, 1, \dots, a\}$; other times (as with bold play) we might want to use the interval $[0, 1]$ as the fortune space. As for the bets, recall that the gambler cannot bet what she does not have and will not bet more than she needs in order to reach the target. Thus, a betting function S must satisfy

$$S(x) \leq \min\{x, a - x\}, \quad x \in A \quad (13.11.1)$$

Moreover, we always restrict our strategies to those for which the stopping time N is finite.

The *success function* of a strategy is the probability that the gambler reaches the target a with that strategy, as a function of the initial fortune x . A strategy with success function V is *optimal* if for any other strategy with success function U , we have $U(x) \leq V(x)$ for $x \in A$.

If there exists an optimal strategy, then the optimal success function is unique.

However, there may not exist an optimal strategy or there may be several optimal strategies. Moreover, the optimality question depends on the value of the game win probability p , in addition to the structure of fortunes and bets.

A Condition for Optimality

Our main theorem gives a condition for optimality:

A strategy S with success function V is optimal if

$$pV(x+y) + qV(x-y) \leq V(x); \quad x \in A, y \in B_x \quad (13.11.2)$$

Proof

Consider the following strategy: if the initial fortune is $x \in A$, we pick $y \in A_x$ and then bet y on the first game; thereafter we follow strategy S . Conditioning on the outcome of the first game, the success function for this new strategy is

$$U(x) = pV(x+y) + qV(x-y) \quad (13.11.3)$$

Thus, the theorem can be restated as follows: If S is optimal with respect to the class of strategies just described, then S is optimal over all strategies. Let T be an arbitrary strategy with success function U . The random variable $V(X_n)$ can be interpreted as the probability of winning if the gambler's strategy is replaced by strategy S after time n . Conditioning on the outcome of game n gives

$$\mathbb{E}[V(X_n) \mid X_0 = x] = \mathbb{E}[pV(X_{n-1} + Y_n) + qV(X_{n-1} - Y_n) \mid X_0 = x] \quad (13.11.4)$$

Using the the optimality condition gives

$$\mathbb{E}[V(X_n) \mid X_0 = x] \leq \mathbb{E}[V(X_{n-1}) \mid X_0 = x], \quad n \in \mathbb{N}_+, x \in A \quad (13.11.5)$$

It follows that $\mathbb{E}[V(X_n) \mid X_0 = x] \leq V(x)$ for $n \in \mathbb{N}_+$ and $x \in A$. Now let N denote the stopping time for strategy T . Letting $n \rightarrow \infty$ we have $\mathbb{E}[V(X_N) \mid X_0 = x] \leq V(x)$ for $x \in A$. But $\mathbb{E}[V(X_N) \mid X_0 = x] = U(x)$ for $x \in A$. Thus $U(x) \leq V(x)$ for $x \in A$.

Favorable Trials with a Minimum Bet

Suppose now that $p \geq \frac{1}{2}$ so that the trials are favorable (or at least not unfair) to the gambler. Next, suppose that all bets must be multiples of a basic unit, which we might as well assume is \$1. Of course, real gambling houses have this restriction. Thus the set of valid fortunes is $A = \{0, 1, \dots, a\}$ and the set of valid bets for $x \in A$ is $B_x = \{0, 1, \dots, \min\{x, a - x\}\}$. Our main result for this case is

Timid play is an optimal strategy.

Proof

Recall the success function f for timid play and recall the [optimality condition](#). This condition holds if $p = \frac{1}{2}$. If $p > \frac{1}{2}$, the condition for optimality is equivalent to

$$p \left(\frac{q}{p} \right)^{x+y} + q \left(\frac{q}{p} \right)^{x-y} \geq \left(\frac{q}{p} \right)^x \quad (13.11.6)$$

But this condition is equivalent to $pq(p^y - q^y)(p^{y-1} - q^{y-1}) \leq 0$ which clearly holds if $p > \frac{1}{2}$.

In the red and black game set the target fortune to 16, the initial fortune to 8, and the game win probability to 0.45. Define the strategy of your choice and play 100 games. Compare your relative frequency of wins with the probability of winning with timid play.

Favorable Trials without a Minimum Bet

We will now assume that the house allows arbitrarily small bets and that $p > \frac{1}{2}$, so that the trials are strictly favorable. In this case it is natural to take the target as the monetary unit so that the set of fortunes is $A = [0, 1]$, and the set of bets for $x \in A$ is $B_x = [0, \min\{x, 1 - x\}]$. Our main result for this case is given below. The results for timid play will play an important role in the analysis, so we will let $f(j, a)$ denote the probability of reaching an integer target a , starting at the integer $j \in [0, a]$, with unit bets.

The optimal success function is $V(x) = x$ for $x \in (0, 1]$.

Proof

Fix a rational initial fortune $x = \frac{k}{n} \in [0, 1]$. Let m be a positive integer and suppose that, starting at x , the gambler bets $\frac{1}{mn}$ on each game. This strategy is equivalent to timid play with target fortune mn , and initial fortune mk . Hence the probability of reaching the target 1 under this strategy is $f(mk, mn)$. But $f(mk, mn) \rightarrow 1$ as $m \rightarrow \infty$. It follows that $V(x) = x$ if $x \in (0, 1]$ is rational. But V is increasing so $V(x) = x$ for all $x \in (0, 1]$

Unfair Trials

We will now assume that $p \leq \frac{1}{2}$ so that the trials are unfair, or at least not favorable. As before, we will take the target fortune as the basic monetary unit and allow any valid fraction of this unit as a bet. Thus, the set of fortunes is $A = [0, 1]$, and the set of bets for $x \in A$ is $B_x = [0, \min\{x, 1 - x\}]$. Our main result for this case is

Bold play is optimal.

Proof

Let F denote the success function for bold play, and let

$$D(x, y) = F\left(\frac{x+y}{2}\right) - [pF(x) + qF(y)] \quad (13.11.7)$$

The [optimality condition](#) equivalent to $D(x, y) \leq 0$ for $0 \leq x \leq y \leq 1$. From the continuity of F , it suffices to prove this inequality when x and y are binary rationals. It's simple to see that $D(x, y) \leq 0$ when x and y have rank 0: $x = 0, y = 0$ or $x = 0, y = 1$ or $x = 1, y = 1$. Suppose now that $D(x, y) \leq 0$ when x and y have rank m or less. We have the following cases:

1. If $x \leq y \leq \frac{1}{2}$ then $D(x, y) = pD(2x, 2y)$.

2. If $\frac{1}{2} \leq x \leq y$ then $D(x, y) = qD(2x - 1, 2y - 1)$.
3. If $x \leq \frac{x+y}{2} \leq \frac{1}{2} \leq y$ and $2y - 1 \leq 2x$ then $D(x, y) = (q-p)F(2y - 1) + qD(2y - 1, 2x)$.
4. If $x \leq \frac{x+y}{2} \leq \frac{1}{2} \leq y$ and $2x \leq 2y - 1$ then $D(x, y) = (q-p)F(2x) + qD(2x, 2y - 1)$.
5. If $x \leq \frac{1}{2} \leq \frac{x+y}{2} \leq y$ and $2y - 1 \leq 2x$ then $D(x, y) = p(q-p)[1 - F(2x)] + pD(2y - 1, 2x)$.
6. If $x \leq \frac{1}{2} \leq \frac{x+y}{2} \leq y$ and $2x \leq 2y - 1$ then $D(x, y) = p(q-p)[1 - F(2y - 1)] + pD(2x, 2y - 1)$.

The induction hypothesis can now be applied to each case to finish the proof.

In the red and black game, set the target fortune to 16, the initial fortune to 8, and the game win probability to 0.45. Define the strategy of your choice and play 100 games. Compare your relative frequency of wins with the probability of winning with bold play.

Other Optimal Strategies in the Sub-Fair Case

Consider again the sub-fair case where $p \leq \frac{1}{2}$ so that the trials are not favorable to the gambler. We will show that bold play is not the only optimal strategy; amazingly, there are infinitely many optimal strategies. Recall first that the bold strategy has betting function

$$S_1(x) = \min\{x, 1 - x\} = \begin{cases} x, & 0 \leq x \leq \frac{1}{2} \\ 1 - x, & \frac{1}{2} \leq x \leq 1 \end{cases} \quad (13.11.8)$$




Figure 13.11.1: The betting function for bold play

Consider the following strategy, which we will refer to as the *second order bold strategy*:

1. With fortune $x \in (0, \frac{1}{2})$, play boldly with the object of reaching $\frac{1}{2}$ before falling to 0.
2. With fortune $x \in (\frac{1}{2}, 1)$, play boldly with the object of reaching 1 without falling below $\frac{1}{2}$.
3. With fortune $\frac{1}{2}$, play boldly and bet $\frac{1}{2}$

The second order bold strategy has betting function S_2 given by

$$S_2(x) = \begin{cases} x, & 0 \leq x < \frac{1}{4} \\ \frac{1}{2} - x, & \frac{1}{4} \leq x < \frac{1}{2} \\ \frac{1}{2}, & x = \frac{1}{2} \\ x - \frac{1}{2}, & \frac{1}{2} < x \leq \frac{3}{4} \\ 1 - x, & \frac{3}{4} < x \leq 1 \end{cases} \quad (13.11.9)$$




Figure 13.11.2: The betting function for the second order bold strategy

The second order bold strategy is optimal.

Proof

Let F_2 denote the success function associated with strategy S_2 . Suppose first that the player starts with fortune $x \in (0, \frac{1}{2})$ under strategy S_2 . Note that the player reaches the target 1 if and only if she reaches $\frac{1}{2}$ and then wins the final game. Consider the sequence of fortunes until the player reaches 0 or $\frac{1}{2}$. If we double the fortunes, then we have the fortune sequence under the ordinary bold strategy, starting at $2x$ and terminating at either 0 or 1. Thus it follows that

$$F_2(x) = pF(2x), \quad 0 < x < \frac{1}{2} \quad (13.11.10)$$

Suppose next that the player starts with fortune $x \in (\frac{1}{2}, 1)$ under strategy S_2 . Note that the player reaches the target 1 if and only if she reaches 1 without falling back to $\frac{1}{2}$ or falls back to $\frac{1}{2}$ and then wins the final game. Consider the sequence of fortunes until the player reaches $\frac{1}{2}$ or 1. If we double the fortunes and subtract 1, then we have the fortune sequence under the ordinary bold strategy, starting at $2x - 1$ and terminating at either 0 or 1. Thus it follows that

$$F_2(x) = F(2x - 1) + [1 - F(2x - 2)]p = p + qF(2x - 1), \quad \frac{1}{2} < x < 1 \quad (13.11.11)$$

But now, using the functional equation for ordinary bold play, we have $F_2(x) = F(x)$ for all $x \in (0, 1]$, and hence S_2 is optimal.

Once we understand how this construction is done, it's straightforward to define the third order bold strategy and show that it's optimal as well.




Figure 13.11.3: The betting function for the third order bold strategy

Explicitly give the third order betting function and show that the strategy is optimal.

More generally, we can define the n th order bold strategy and show that it is optimal as well.

The sequence of bold strategies can be defined recursively from the basic bold strategy S_1 as follows:

$$S_{n+1}(x) = \begin{cases} \frac{1}{2}S_n(2x), & 0 \leq x < \frac{1}{2} \\ \frac{1}{2}, & x = \frac{1}{2} \\ \frac{1}{2}S_n(2x - 1), & \frac{1}{2} < x \leq 1 \end{cases} \quad (13.11.12)$$

S_n is optimal for each n .

Even more generally, we can define an optimal strategy T in the following way: for each $x \in [0, 1]$ select $n_x \in \mathbb{N}_+$ and let $T(x) = S_{n_x}(x)$. The graph below shows a few of the graphs of the bold strategies. For an optimal strategy T , we just need to select, for each x a bet on one of the graphs.




Figure 13.11.4: The family of bold strategies

Martingales

Let's return to the unscaled formulation of red and black, where the target fortune is $a \in (0, \infty)$ and the initial fortune is $x \in (0, a)$. In the subfair case, when $p \leq \frac{1}{2}$, no strategy can do better than the optimal strategies, so that the win probability is bounded by x/a (with equality when $p = \frac{1}{2}$). Another elegant proof of this is given in the section on inequalities in the chapter on martingales.

This page titled [13.11: Optimal Strategies](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

14: The Poisson Process

The Poisson process is one of the most important random processes in probability theory. It is widely used to model random “points” in time and space, such as the times of radioactive emissions, the arrival times of customers at a service center, and the positions of flaws in a piece of material. Several important probability distributions arise naturally from the Poisson process—the Poisson distribution, the exponential distribution, and the gamma distribution. The process has a beautiful mathematical structure, and is used as a foundation for building a number of other, more complicated random processes.

- [14.1: Introduction to the Poisson Process](#)
- [14.2: The Exponential Distribution](#)
- [14.3: The Gamma Distribution](#)
- [14.4: The Poisson Distribution](#)
- [14.5: Thinning and Superposition](#)
- [14.6: Non-homogeneous Poisson Processes](#)
- [14.7: Compound Poisson Processes](#)
- [14.8: Poisson Processes on General Spaces](#)

This page titled [14: The Poisson Process](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.1: Introduction to the Poisson Process

The Poisson Model

We will consider a process in which points occur randomly in time. The phrase *points in time* is generic and could represent, for example:

- The times when a sample of radioactive material emits particles
- The times when customers arrive at a service station
- The times when file requests arrive at a server computer
- The times when accidents occur at a particular intersection
- The times when a device fails and is replaced by a new device

It turns out that under some basic assumptions that deal with independence and uniformity in time, a *single*, one-parameter probability model governs all such random processes. This is an amazing result, and because of it, the *Poisson process* (named after Simeon Poisson) is one of the most important in probability theory.

Run the Poisson experiment with the default settings in single step mode. Note the random points in time.

Random Variables

There are three collections of random variables that can be used to describe the process. First, let X_1 denote the time of the first arrival, and X_i the time between the $(i-1)$ st and i th arrival for $i \in \{2, 3, \dots\}$. Thus, $\mathbf{X} = (X_1, X_2, \dots)$ is the sequence of *inter-arrival times*. Next, let T_n denote the time of the n th arrival for $n \in \mathbb{N}_+$. It will be convenient to define $T_0 = 0$, although we do not consider this as an arrival. Thus $\mathbf{T} = (T_0, T_1, \dots)$ is the sequence of *arrival times*. Clearly \mathbf{T} is the partial sum process associated \mathbf{X} , and so in particular each sequence determines the other:

$$T_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (14.1.1)$$

$$X_n = T_n - T_{n-1}, \quad n \in \mathbb{N}_+ \quad (14.1.2)$$

Next, let N_t denote the number of arrivals in $(0, t]$ for $t \in [0, \infty)$. The random process $\mathbf{N} = (N_t : t \geq 0)$ is the *counting process*. The arrival time process \mathbf{T} and the counting process \mathbf{N} are inverses of one another in a sense, and in particular each process determines the other:

$$T_n = \min\{t \geq 0 : N_t = n\}, \quad n \in \mathbb{N} \quad (14.1.3)$$

$$N_t = \max\{n \in \mathbb{N} : T_n \leq t\}, \quad t \in [0, \infty) \quad (14.1.4)$$

Note also that $N_t \geq n$ if and only if $T_n \leq t$ for $n \in \mathbb{N}$ and $t \in [0, \infty)$ since each of these events means that there are at least n arrivals in the interval $(0, t]$.

Sometimes it will be helpful to extend the notation of the counting process. For $A \subseteq [0, \infty)$ (measurable of course), let $N(A)$ denote the number of arrivals in A :

$$N(A) = \#\{n \in \mathbb{N}_+ : T_n \in A\} = \sum_{n=1}^{\infty} \mathbf{1}(T_n \in A) \quad (14.1.5)$$

Thus, $A \mapsto N(A)$ is the counting measure associated with the random points (T_1, T_2, \dots) , so in particular it is a *random measure*. For our original counting process, note that $N_t = N(0, t]$ for $t \geq 0$. Thus, $t \mapsto N_t$ is a (random) distribution function, and $A \mapsto N(A)$ is the (random) measure associated with this distribution function.

The Basic Assumption

The assumption that we will make can be described intuitively (but imprecisely) as follows: If we fix a time t , whether constant or one of the arrival times, then the process *after* time t is independent of the process *before* time t and behaves probabilistically just like the original process. Thus, the random process has a *strong renewal property*. Making the strong renewal assumption precise will enable use to completely specify the probabilistic behavior of the process, up to a single, positive parameter.

Think about the strong renewal assumption for each of the specific applications given above.

Run the Poisson experiment with the default settings in single step mode. See if you can detect the strong renewal assumption.

As a first step, note that part of the renewal assumption, namely that the process “restarts” at each arrival time, independently of the past, implies the following result:

The sequence of inter-arrival times \mathbf{X} is an independent, identically distributed sequence

Proof

Note that X_2 is the first arrival time after $T_1 = X_1$, so X_2 must be independent of X_1 and have the same distribution. Similarly X_3 is the first arrival time after $T_2 = X_1 + X_2$, so X_3 must be independent of X_1 and X_2 and have the same distribution as X_1 . Continuing this argument, \mathbf{X} must be an independent, identically distributed sequence.

A model of random points in time in which the inter-arrival times are independent and identically distributed (so that the process “restarts” at each arrival time) is known as a *renewal process*. A separate chapter explores Renewal Processes in detail. Thus, the Poisson process is a renewal process, but a very special one, because we also require that the renewal assumption hold at *fixed times*.

Analogy with Bernoulli Trials

In some sense, the Poisson process is a continuous time version of the Bernoulli trials process. To see this, suppose that we have a Bernoulli trials process with success parameter $p \in (0, 1)$, and that we think of each success as a random point in *discrete* time. Then this process, like the Poisson process (and in fact any renewal process) is completely determined by the sequence of inter-arrival times $\mathbf{X} = (X_1, X_2, \dots)$ (in this case, the number of trials between successive successes), the sequence of arrival times $\mathbf{T} = (T_0, T_1, \dots)$ (in this case, the trial numbers of the successes), and the counting process $(N_t : t \in \mathbb{N})$ (in this case, the number of successes in the first t trials). Also like the Poisson process, the Bernoulli trials process has the strong renewal property: at each fixed time and at each arrival time, the process “starts over” independently of the past. But of course, time is discrete in the Bernoulli trials model and continuous in the Poisson model. The Bernoulli trials process can be characterized in terms of each of the three sets of random variables.

Each of the following statements characterizes the Bernoulli trials process with success parameter $p \in (0, 1)$:

1. The inter-arrival time sequence \mathbf{X} is a sequence of independent variables, and each has the geometric distributions on \mathbb{N}_+ with success parameter p .
2. The arrival time sequence \mathbf{T} has stationary, independent increments, and for $n \in \mathbb{N}_+$, T_n has the negative binomial distribution with stopping parameter n and success parameter p .
3. The counting process \mathbf{N} has stationary, independent increments, and for $t \in \mathbb{N}$, N_t has the binomial distribution with trial parameter t and success parameter p .

Run the binomial experiment with $n = 50$ and $p = 0.1$. Note the random points in discrete time.

Run the Poisson experiment with $t = 5$ and $r = 1$. Note the random points in continuous time and compare with the behavior in the previous exercise.

As we develop the theory of the Poisson process we will frequently refer back to the analogy with Bernoulli trials. In particular, we will show that if we run the Bernoulli trials at a faster and faster rate but with a smaller and smaller success probability, in just the right way, the Bernoulli trials process converges to the Poisson process.

This page titled [14.1: Introduction to the Poisson Process](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.2: The Exponential Distribution

Basic Theory

The Memoryless Property

Recall that in the basic model of the Poisson process, we have “points” that occur randomly in time. The sequence of inter-arrival times is $\mathbf{X} = (X_1, X_2, \dots)$. The strong renewal assumption states that at each arrival time and at each fixed time, the process must probabilistically restart, independent of the past. The first part of that assumption implies that \mathbf{X} is a sequence of independent, identically distributed variables. The second part of the assumption implies that if the first arrival has not occurred by time s , then the time remaining until the arrival occurs must have the same distribution as the first arrival time itself. This is known as the *memoryless property* and can be stated in terms of a general random variable as follows:

Suppose that X takes values in $[0, \infty)$. Then X has the *memoryless property* if the conditional distribution of $X - s$ given $X > s$ is the same as the distribution of X for every $s \in [0, \infty)$. Equivalently,

$$\mathbb{P}(X > t + s \mid X > s) = \mathbb{P}(X > t), \quad s, t \in [0, \infty) \quad (14.2.1)$$

The memoryless property determines the distribution of X up to a positive parameter, as we will see now.

Distribution functions

Suppose that X takes values in $[0, \infty)$ and satisfies the memoryless property.

X has a continuous distribution and there exists $r \in (0, \infty)$ such that the distribution function F of X is

$$F(t) = 1 - e^{-rt}, \quad t \in [0, \infty) \quad (14.2.2)$$

Proof

Let $F^c = 1 - F$ denote the right-tail distribution function of X (also known as the *reliability function*), so that $F^c(t) = \mathbb{P}(X > t)$ for $t \geq 0$. From the definition of conditional probability, the memoryless property is equivalent to the *law of exponents*:

$$F^c(t + s) = F^c(s)F^c(t), \quad s, t \in [0, \infty) \quad (14.2.3)$$

Let $a = F^c(1)$. Implicit in the memoryless property is $\mathbb{P}(X > t) > 0$ for $t \in [0, \infty)$, so $a > 0$. If $n \in \mathbb{N}_+$ then

$$F^c(n) = F^c\left(\sum_{i=1}^n 1\right) = \prod_{i=1}^n F^c(1) = [F^c(1)]^n = a^n \quad (14.2.4)$$

Next, if $n \in \mathbb{N}_+$ then

$$a = F^c(1) = F^c\left(\frac{n}{n}\right) = F^c\left(\sum_{i=1}^n \frac{1}{n}\right) = \prod_{i=1}^n F^c\left(\frac{1}{n}\right) = \left[F^c\left(\frac{1}{n}\right)\right]^n \quad (14.2.5)$$

so $F^c\left(\frac{1}{n}\right) = a^{1/n}$. Now suppose that $m \in \mathbb{N}$ and $n \in \mathbb{N}_+$. Then

$$F^c\left(\frac{m}{n}\right) = F^c\left(\sum_{i=1}^m \frac{1}{n}\right) = \prod_{i=1}^m F^c\left(\frac{1}{n}\right) = \left[F^c\left(\frac{1}{n}\right)\right]^m = a^{m/n} \quad (14.2.6)$$

Thus we have $F^c(q) = a^q$ for rational $q \in [0, \infty)$. For $t \in [0, \infty)$, there exists a sequence of rational numbers (q_1, q_2, \dots) with $q_n \downarrow t$ as $n \uparrow \infty$. We have $F^c(q_n) = a^{q_n}$ for each $n \in \mathbb{N}_+$. But F^c is continuous from the right, so taking limits gives $a^t = F^c(t)$. Now let $r = -\ln(a)$. Then $F^c(t) = e^{-rt}$ for $t \in [0, \infty)$.

The probability density function of X is

$$f(t) = r e^{-rt}, \quad t \in [0, \infty) \quad (14.2.7)$$

1. f is decreasing on $[0, \infty)$.
2. f is concave upward on $[0, \infty)$.
3. $f(t) \rightarrow 0$ as $t \rightarrow \infty$.

Proof

This follows since $f = F'$. The properties in parts (a)–(c) are simple.

A random variable with the [distribution function](#) above or equivalently the [probability density function](#) in the last theorem is said to have the *exponential distribution* with *rate parameter* r . The reciprocal $\frac{1}{r}$ is known as the *scale parameter* (as will be justified [below](#)). Note that the mode of the distribution is 0, regardless of the parameter r , not very helpful as a measure of center.

In the gamma experiment, set $n = 1$ so that the simulated random variable has an exponential distribution. Vary r with the scroll bar and watch how the shape of the probability density function changes. For selected values of r , run the experiment 1000 times and compare the empirical density function to the probability density function.

The quantile function of X is

$$F^{-1}(p) = \frac{-\ln(1-p)}{r}, \quad p \in [0, 1) \quad (14.2.8)$$

1. The median of X is $\frac{1}{r} \ln(2) \approx 0.6931 \frac{1}{r}$
2. The first quartile of X is $\frac{1}{r} [\ln(4) - \ln(3)] \approx 0.2877 \frac{1}{r}$
3. The third quartile X is $\frac{1}{r} \ln(4) \approx 1.3863 \frac{1}{r}$
4. The interquartile range is $\frac{1}{r} \ln(3) \approx 1.0986 \frac{1}{r}$

Proof

The formula for F^{-1} follows easily from solving $p = F^{-1}(t)$ for t in terms of p .

In the special distribution calculator, select the exponential distribution. Vary the scale parameter (which is $1/r$) and note the shape of the distribution/quantile function. For selected values of the parameter, compute a few values of the distribution function and the quantile function.

Returning to the Poisson model, we have our first formal definition:

A process of random points in time is a Poisson process with rate $r \in (0, \infty)$ if and only the interarrival times are independent, and each has the exponential distribution with rate r .

Constant Failure Rate

Suppose now that X has a continuous distribution on $[0, \infty)$ and is interpreted as the lifetime of a device. If F denotes the distribution function of X , then $F^c = 1 - F$ is the reliability function of X . If f denotes the probability density function of X then the failure rate function h is given by

$$h(t) = \frac{f(t)}{F^c(t)}, \quad t \in [0, \infty) \quad (14.2.9)$$

If X has the exponential distribution with rate $r > 0$, then from the results above, the reliability function is $F^c(t) = e^{-rt}$ and the probability density function is $f(t) = re^{-rt}$, so trivially X has constant rate r . The converse is also true.

If X has constant failure rate $r > 0$ then X has the exponential distribution with parameter r .

Proof

Recall that in general, the distribution of a lifetime variable X is determined by the failure rate function h . Specifically, if $F^c = 1 - F$ denotes the reliability function, then $(F^c)' = -f$, so $-h = (F^c)' / F^c$. Integrating and then taking exponentials gives

$$F^c(t) = \exp\left(-\int_0^t h(s) ds\right), \quad t \in [0, \infty) \quad (14.2.10)$$

In particular, if $h(t) = r$ for $t \in [0, \infty)$, then $F^c(t) = e^{-rt}$ for $t \in [0, \infty)$.

The memoryless and constant failure rate properties are the most famous characterizations of the exponential distribution, but are by no means the only ones. Indeed, entire books have been written on characterizations of this distribution.

Moments

Suppose again that X has the exponential distribution with rate parameter $r > 0$. Naturally, we want to know the mean, variance, and various other moments of X .

If $n \in \mathbb{N}$ then $\mathbb{E}(X^n) = n! / r^n$.

Proof

By the change of variables theorem for expected value,

$$\mathbb{E}(X^n) = \int_0^\infty t^n r e^{-rt} dt \quad (14.2.11)$$

Integrating by parts gives $\mathbb{E}(X^n) = \frac{n}{r} \mathbb{E}(X^{n-1})$ for $n \in \mathbb{N}^+$. Of course $\mathbb{E}(X^0) = 1$ so the result now follows by induction.

More generally, $\mathbb{E}(X^a) = \Gamma(a+1) / r^a$ for every $a \in [0, \infty)$, where Γ is the gamma function.

In particular.

1. $\mathbb{E}(X) = \frac{1}{r}$
2. $\text{var}(X) = \frac{1}{r^2}$
3. $\text{skew}(X) = 2$
4. $\text{kurt}(X) = 9$

In the context of the Poisson process, the parameter r is known as the *rate* of the process. On average, there are $1/r$ time units between arrivals, so the arrivals come at an average rate of r per unit time. The Poisson process is completely determined by the sequence of inter-arrival times, and hence is completely determined by the rate r .

Note also that the mean and standard deviation are equal for an exponential distribution, and that the median is always smaller than the mean. Recall also that skewness and kurtosis are standardized measures, and so do not depend on the parameter r (which is the reciprocal of the scale parameter).

The moment generating function of X is

$$M(s) = \mathbb{E}(e^{sX}) = \frac{r}{r-s}, \quad s \in (-\infty, r) \quad (14.2.12)$$

Proof

By the change of variables theorem

$$M(s) = \int_0^{\infty} e^{st} r e^{-rt} dt = \int_0^{\infty} r e^{(s-r)t} dt \quad (14.2.13)$$

The integral evaluates to $\frac{r}{r-s}$ if $s < r$ and to ∞ if $s \geq r$.

In the gamma experiment, set $n = 1$ so that the simulated random variable has an exponential distribution. Vary r with the scroll bar and watch how the mean \pm standard deviation bar changes. For various values of r , run the experiment 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation, respectively.

Additional Properties

The exponential distribution has a number of interesting and important mathematical properties. First, and not surprisingly, it's a member of the general exponential family.

Suppose that X has the exponential distribution with rate parameter $r \in (0, \infty)$. Then X has a one parameter general exponential distribution, with natural parameter $-r$ and natural statistic X .

Proof

This follows directly from the form of the PDF, $f(x) = r e^{-rx}$ for $x \in [0, \infty)$, and the definition of the general exponential family.

The Scaling Property

As suggested earlier, the exponential distribution is a scale family, and $1/r$ is the scale parameter.

Suppose that X has the exponential distribution with rate parameter $r > 0$ and that $c > 0$. Then cX has the exponential distribution with rate parameter r/c .

Proof

For $t \geq 0$, $\mathbb{P}(cX > t) = \mathbb{P}(X > t/c) = e^{-r(t/c)} = e^{-(r/c)t}$.

Recall that multiplying a random variable by a positive constant frequently corresponds to a change of units (minutes into hours for a lifetime variable, for example). Thus, the exponential distribution is preserved under such changes of units. In the context of the Poisson process, this has to be the case, since the memoryless property, which led to the exponential distribution in the first place, clearly does not depend on the time units.

In fact, the exponential distribution with rate parameter 1 is referred to as the *standard exponential distribution*. From the previous result, if Z has the standard exponential distribution and $r > 0$, then $X = \frac{1}{r}Z$ has the exponential distribution with rate parameter r . Conversely, if X has the exponential distribution with rate $r > 0$ then $Z = rX$ has the standard exponential distribution.

Similarly, the Poisson process with rate parameter 1 is referred to as the *standard Poisson process*. If Z_i is the i th inter-arrival time for the standard Poisson process for $i \in \mathbb{N}_+$, then letting $X_i = \frac{1}{r}Z_i$ for $i \in \mathbb{N}_+$ gives the inter-arrival times for the Poisson process with rate r . Conversely if X_i is the i th inter-arrival time of the Poisson process with rate $r > 0$ for $i \in \mathbb{N}_+$, then $Z_i = rX_i$ for $i \in \mathbb{N}_+$ gives the inter-arrival times for the standard Poisson process.

Relation to the Geometric Distribution

In many respects, the geometric distribution is a discrete version of the exponential distribution. In particular, recall that the geometric distribution on \mathbb{N}_+ is the only distribution on \mathbb{N}_+ with the memoryless and constant rate properties. So it is not surprising that the two distributions are also connected through various transformations and limits.

Suppose that X has the exponential distribution with rate parameter $r > 0$. Then

1. $\lfloor X \rfloor$ has the geometric distributions on \mathbb{N} with success parameter $1 - e^{-r}$.
2. $\lceil X \rceil$ has the geometric distributions on \mathbb{N}_+ with success parameter $1 - e^{-r}$.

Proof

1. For $n \in \mathbb{N}$ note that $\mathbb{P}(\lfloor X \rfloor = n) = \mathbb{P}(n \leq X < n+1) = F(n+1) - F(n)$. Substituting into the distribution function and simplifying gives $\mathbb{P}(\lfloor X \rfloor = n) = (e^{-r})^n (1 - e^{-r})$.
2. For $n \in \mathbb{N}_+$ note that $\mathbb{P}(\lceil X \rceil = n) = \mathbb{P}(n-1 < X \leq n) = F(n) - F(n-1)$. Substituting into the distribution function and simplifying gives $\mathbb{P}(\lceil X \rceil = n) = (e^{-r})^{n-1} (1 - e^{-r})$.

The following connection between the two distributions is interesting by itself, but will also be very important in the section on splitting Poisson processes. In words, a random, geometrically distributed sum of independent, identically distributed exponential variables is itself exponential.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent variables, each with the exponential distribution with rate r . Suppose that U has the geometric distribution on \mathbb{N}_+ with success parameter p and is independent of \mathbf{X} . Then $Y = \sum_{i=1}^U X_i$ has the exponential distribution with rate rp .

Proof

Recall that the moment generating function of Y is $P \circ M$ where M is the common moment generating function of the terms in the sum, and P is the probability generating function of the number of terms U . But $M(s) = r/(r-s)$ for $s < r$ and $P(s) = ps/[1 - (1-p)s]$ for $s < 1/(1-p)$. Thus,

$$(P \circ M)(s) = \frac{pr/(r-s)}{1 - (1-p)r/(r-s)} = \frac{pr}{pr-s}, \quad s < pr \quad (14.2.14)$$

It follows that Y has the exponential distribution with parameter pr

The next result explores the connection between the Bernoulli trials process and the Poisson process that was begun in the Introduction.

For $n \in \mathbb{N}_+$, suppose that U_n has the geometric distribution on \mathbb{N}_+ with success parameter p_n , where $np_n \rightarrow r > 0$ as $n \rightarrow \infty$. Then the distribution of U_n/n converges to the exponential distribution with parameter r as $n \rightarrow \infty$.

Proof

Let F_n denote the CDF of U_n/n . Then for $x \in [0, \infty)$

$$F_n(x) = \mathbb{P}\left(\frac{U_n}{n} \leq x\right) = \mathbb{P}(U_n \leq nx) = \mathbb{P}(U_n \leq \lfloor nx \rfloor) = 1 - (1 - p_n)^{\lfloor nx \rfloor} \quad (14.2.15)$$

But by a famous limit from calculus, $(1 - p_n)^n = \left(1 - \frac{np_n}{n}\right)^n \rightarrow e^{-r}$ as $n \rightarrow \infty$, and hence $(1 - p_n)^{nx} \rightarrow e^{-rx}$ as $n \rightarrow \infty$. But by definition, $\lfloor nx \rfloor \leq nx < \lfloor nx \rfloor + 1$ or equivalently, $nx - 1 < \lfloor nx \rfloor \leq nx$ so it follows that $(1 - p_n)^{\lfloor nx \rfloor} \rightarrow e^{-rx}$ as $n \rightarrow \infty$. Hence $F_n(x) \rightarrow 1 - e^{-rx}$ as $n \rightarrow \infty$, which is the CDF of the exponential distribution.

To understand this result more clearly, suppose that we have a sequence of Bernoulli trials processes. In process n , we run the trials at a rate of n per unit time, with probability of success p_n . Thus, the actual time of the first success in process n is U_n/n . The last result shows that if $np_n \rightarrow r > 0$ as $n \rightarrow \infty$, then the sequence of Bernoulli trials processes converges to the Poisson process with rate parameter r as $n \rightarrow \infty$. We will return to this point in subsequent sections.

Orderings and Order Statistics

Suppose that X and Y have exponential distributions with parameters a and b , respectively, and are independent. Then

$$\mathbb{P}(X < Y) = \frac{a}{a+b} \quad (14.2.16)$$

Proof

This result can be proved in a straightforward way by integrating the joint PDF of (X, Y) over $\{(x, y) : 0 < x < y < \infty\}$. A more elegant proof uses conditioning and the [moment generating function](#) above:

$$\mathbb{P}(Y > X) = \mathbb{E}[\mathbb{P}(Y > X | X)] = \mathbb{E}(e^{-bX}) = \frac{a}{a+b} \quad (14.2.17)$$

The following theorem gives an important random version of the memoryless property.

Suppose that X and Y are independent variables taking values in $[0, \infty)$ and that Y has the exponential distribution with rate parameter $r > 0$. Then X and $Y - X$ are conditionally independent given $X < Y$, and the conditional distribution of $Y - X$ is also exponential with parameter r .

Proof

Suppose that $A \subseteq [0, \infty)$ (measurable of course) and $t \geq 0$. Then

$$\mathbb{P}(X \in A, Y - X \geq t | X < Y) = \frac{\mathbb{P}(X \in A, Y - X \geq t)}{\mathbb{P}(X < Y)} \quad (14.2.18)$$

But conditioning on X we can write the numerator as

$$\mathbb{P}(X \in A, Y - X > t) = \mathbb{E}[\mathbb{P}(X \in A, Y - X > t | X)] = \mathbb{E}[\mathbb{P}(Y > X + t | X), X \in A] = \mathbb{E}[e^{-r(t+X)}, X \in A] = e^{-rt} \mathbb{E}(e^{-rX}, X \in A) \quad (14.2.19)$$

Similarly, conditioning on X gives $\mathbb{P}(X < Y) = \mathbb{E}(e^{-rX})$. Thus

$$\mathbb{P}(X \in A, Y - X > t | X < Y) = e^{-rt} \frac{\mathbb{E}(e^{-rX}, X \in A)}{\mathbb{E}(e^{-rX})} \quad (14.2.20)$$

Letting $A = [0, \infty)$ we have $\mathbb{P}(Y > t) = e^{-rt}$ so given $X < Y$, the variable $Y - X$ has the exponential distribution with parameter r . Letting $t = 0$, we see that given $X < Y$, variable X has the distribution

$$A \mapsto \frac{\mathbb{E}(e^{-rX}, X \in A)}{\mathbb{E}(e^{-rX})} \quad (14.2.21)$$

Finally, because of the factoring, X and $Y - X$ are conditionally independent given $X < Y$.

For our next discussion, suppose that $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is a sequence of independent random variables, and that X_i has the exponential distribution with rate parameter $r_i > 0$ for each $i \in \{1, 2, \dots, n\}$.

Let $U = \min\{X_1, X_2, \dots, X_n\}$. Then U has the exponential distribution with parameter $\sum_{i=1}^n r_i$.

Proof

Recall that in general, $\{U > t\} = \{X_1 > t, X_2 > t, \dots, X_n > t\}$ and therefore by independence, $F^c(t) = F_1^c(t)F_2^c(t) \cdots F_n^c(t)$ for $t \geq 0$, where F^c is the reliability function of U and F_i^c is the reliability function of X_i for each i . When X_i has the exponential distribution with rate r_i for each i , we have $F^c(t) = \exp[-(\sum_{i=1}^n r_i)t]$ for $t \geq 0$.

In the context of reliability, if a series system has independent components, each with an exponentially distributed lifetime, then the lifetime of the system is also exponentially distributed, and the failure rate of the system is the sum of the component failure rates. In the context of random processes, if we have n independent Poisson processes, then the new process obtained by combining the random points in time is also Poisson, and the rate of the new process is the sum of the rates of the individual processes (we will return to this point later).

Let $V = \max\{X_1, X_2, \dots, X_n\}$. Then V has distribution function F given by

$$F(t) = \prod_{i=1}^n (1 - e^{-r_i t}), \quad t \in [0, \infty) \quad (14.2.22)$$

Proof

Recall that in general, $\{V \leq t\} = \{X_1 \leq t, X_2 \leq t, \dots, X_n \leq t\}$ and therefore by independence, $F(t) = F_1(t)F_2(t) \cdots F_n(t)$ for $t \geq 0$, where F is the distribution function of V and F_i is the distribution function of X_i for each i .

Consider the special case where $r_i = r \in (0, \infty)$ for each $i \in \mathbb{N}_+$. In statistical terms, \mathbf{X} is a random sample of size n from the exponential distribution with parameter r . From the last couple of theorems, the minimum U has the exponential distribution with rate nr while the maximum V has distribution function $F(t) = (1 - e^{-rt})^n$ for $t \in [0, \infty)$. Recall that U and V are the first and last order statistics, respectively.

In the order statistic experiment, select the exponential distribution.

1. Set $k = 1$ (this gives the minimum U). Vary n with the scroll bar and note the shape of the probability density function. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.
2. Vary n with the scroll bar, set $k = n$ each time (this gives the maximum V), and note the shape of the probability density function. For selected values of n , run the simulation 1000 times and compare the empirical density function to the true probability density function.

Curiously, the distribution of the maximum of independent, identically distributed exponential variables is also the distribution of the sum of independent exponential variables, with rates that grow linearly with the index.

Suppose that $r_i = ir$ for each $i \in \{1, 2, \dots, n\}$ where $r \in (0, \infty)$. Then $Y = \sum_{i=1}^n X_i$ has distribution function F given by

$$F(t) = (1 - e^{-rt})^n, \quad t \in [0, \infty) \quad (14.2.23)$$

Proof

By assumption, X_k has PDF f_k given by $f_k(t) = kre^{-krt}$ for $t \in [0, \infty)$. We want to show that $Y_n = \sum_{i=1}^n X_i$ has PDF g_n given by

$$g_n(t) = nre^{-rt}(1 - e^{-rt})^{n-1}, \quad t \in [0, \infty) \quad (14.2.24)$$

The PDF of a sum of independent variables is the convolution of the individual PDFs, so we want to show that

$$f_1 * f_2 * \cdots * f_n = g_n, \quad n \in \mathbb{N}_+ \quad (14.2.25)$$

The proof is by induction on n . Trivially $f_1 = g_1$, so suppose the result holds for a given $n \in \mathbb{N}_+$. Then

$$\begin{aligned} g_n * f_{n+1}(t) &= \int_0^t g_n(s)f_{n+1}(t-s)ds = \int_0^t nre^{-rs}(1 - e^{-rs})^{n-1}(n+1)re^{-(n+1)(t-s)}ds \\ &= r(n+1)e^{-r(n+1)t} \int_0^t n(1 - e^{-rs})^{n-1}re^{rs}ds \end{aligned}$$

Now substitute $u = e^{-rs}$ so that $du = re^{-rs}ds$ or equivalently $rd s = du/u$. After some algebra,

$$\begin{aligned} g_n * f_{n+1}(t) &= r(n+1)e^{-r(n+1)t} \int_1^{e^{-rt}} n(u-1)^{n-1}du \\ &= r(n+1)e^{-r(n+1)t}(e^{-rt} - 1)^n = r(n+1)e^{-rt}(1 - e^{-rt})^n = g_{n+1}(t) \end{aligned}$$

This result has an application to the *Yule process*, named for George Yule. The Yule process, which has some parallels with the Poisson process, is studied in the chapter on Markov processes. We can now generalize the [order probability](#) above:

For $i \in \{1, 2, \dots, n\}$,

$$\mathbb{P}(X_i < X_j \text{ for all } j \neq i) = \frac{r_i}{\sum_{j=1}^n r_j} \quad (14.2.26)$$

Proof

First, note that $X_i < X_j$ for all $i \neq j$ if and only if $X_i < \min\{X_j : j \neq i\}$. But the minimum on the right is independent of X_i and, by [result on minimums](#) above, has the exponential distribution with parameter $\sum_{j \neq i} r_j$. The result now follows from [order probability](#) for two events above.

Suppose that for each i , X_i is the time until an event of interest occurs (the arrival of a customer, the failure of a device, etc.) and that these times are independent and exponentially distributed. Then the *first* time U that one of the events occurs is also exponentially distributed, and the probability that the first event to occur is event i is proportional to the rate r_i .

The probability of a total ordering is

$$\mathbb{P}(X_1 < X_2 < \cdots < X_n) = \prod_{i=1}^n \frac{r_i}{\sum_{j=i}^n r_j} \quad (14.2.27)$$

Proof

Let $A = \{X_1 < X_j \text{ for all } j \in \{2, 3, \dots, n\}\}$. then

$$\mathbb{P}(X_1 < X_2 < \cdots < X_n) = \mathbb{P}(A, X_2 < X_3 < \cdots < X_n) = \mathbb{P}(A)\mathbb{P}(X_2 < X_3 < \cdots < X_n | A) \quad (14.2.28)$$

But $\mathbb{P}(A) = \frac{r_1}{\sum_{i=1}^n r_i}$ from the previous result, and $\{X_2 < X_3 < \cdots < X_n\}$ is independent of A . Thus we have

$$\mathbb{P}(X_1 < X_2 < \dots < X_n) = \frac{r_1}{\sum_{i=1}^n r_i} \mathbb{P}(X_2 < X_3 < \dots < X_n) \quad (14.2.29)$$

so the result follows by induction.

Of course, the probabilities of other orderings can be computed by permuting the parameters appropriately in the formula on the right.

The result on [minimums](#) and the [order probability](#) result above are very important in the theory of continuous-time Markov chains. But for that application and others, it's convenient to extend the exponential distribution to two degenerate cases: point mass at 0 and point mass at ∞ (so the first is the distribution of a random variable that takes the value 0 with probability 1, and the second the distribution of a random variable that takes the value ∞ with probability 1). In terms of the rate parameter r and the distribution function F , point mass at 0 corresponds to $r = \infty$ so that $F(t) = 1$ for $0 < t < \infty$. Point mass at ∞ corresponds to $r = 0$ so that $F(t) = 0$ for $0 < t < \infty$. The memoryless property, as expressed in terms of the reliability function F^c , still holds for these degenerate cases on $(0, \infty)$:

$$F^c(s)F^c(t) = F^c(s+t), \quad s, t \in (0, \infty) \quad (14.2.30)$$

We also need to extend some of results above for a finite number of variables to a countably infinite number of variables. So for the remainder of this discussion, suppose that $\{X_i : i \in I\}$ is a countable collection of independent random variables, and that X_i has the exponential distribution with parameter $r_i \in (0, \infty)$ for each $i \in I$.

Let $U = \inf\{X_i : i \in I\}$. Then U has the exponential distribution with parameter $\sum_{i \in I} r_i$

Proof

The proof is almost the same as the one [above](#) for a finite collection. Note that $\{U \geq t\} = \{X_i \geq t \text{ for all } i \in I\}$ and so

$$\mathbb{P}(U \geq t) = \prod_{i \in I} \mathbb{P}(X_i \geq t) = \prod_{i \in I} e^{-r_i t} = \exp\left[-\left(\sum_{i \in I} r_i\right)t\right] \quad (14.2.31)$$

If $\sum_{i \in I} r_i < \infty$ then U has a proper exponential distribution with the sum as the parameter. If $\sum_{i \in I} r_i = \infty$ then $P(U \geq t) = 0$ for all $t \in (0, \infty)$ so $P(U = 0) = 1$.

For $i \in \mathbb{N}_+$,

$$\mathbb{P}(X_i < X_j \text{ for all } j \in I - \{i\}) = \frac{r_i}{\sum_{j \in I} r_j} \quad (14.2.32)$$

Proof

First note that since the variables have continuous distributions and I is countable,

$$\mathbb{P}(X_i < X_j \text{ for all } j \in I - \{i\}) = \mathbb{P}(X_i \leq X_j \text{ for all } j \in I - \{i\}) \quad (14.2.33)$$

Next note that $X_i \leq X_j$ for all $j \in I - \{i\}$ if and only if $X_i \leq U_i$ where $U_i = \inf\{X_j : j \in I - \{i\}\}$. But U_i is independent of X_i and, by [previous result](#), has the exponential distribution with parameter $s_i = \sum_{j \in I - \{i\}} r_j$. If $s_i = \infty$, then U_i is 0 with probability 1, and so $P(X_i \leq U_i) = 0 = r_i/s_i$. If $s_i < \infty$, then X_i and U_i have proper exponential distributions, and so the result now follows from [order probability](#) for two variables above.

We need one last result in this setting: a condition that ensures that the sum of an infinite collection of exponential variables is finite with probability one.

Let $Y = \sum_{i \in I} X_i$ and $\mu = \sum_{i \in I} 1/r_i$. Then $\mu = \mathbb{E}(Y)$ and $\mathbb{P}(Y < \infty) = 1$ if and only if $\mu < \infty$.

Proof

The result is trivial if I is finite, so assume that $I = \mathbb{N}_+$. Recall that $\mathbb{E}(X_i) = 1/r_i$ and hence $\mu = \mathbb{E}(Y)$. Trivially if $\mu < \infty$ then $\mathbb{P}(Y < \infty) = 1$. Conversely, suppose that $\mathbb{P}(Y < \infty) = 1$. Then $\mathbb{P}(e^{-Y} > 0) = 1$ and hence $\mathbb{E}(e^{-Y}) > 0$. Using independence and the [moment generating function](#) above,

$$\mathbb{E}(e^{-Y}) = \mathbb{E}\left(\prod_{i=1}^{\infty} e^{-X_i}\right) = \prod_{i=1}^{\infty} \mathbb{E}(e^{-X_i}) = \prod_{i=1}^{\infty} \frac{r_i}{r_i + 1} > 0 \quad (14.2.34)$$

Next recall that if $p_i \in (0, 1)$ for $i \in \mathbb{N}_+$ then

$$\prod_{i=1}^{\infty} p_i > 0 \text{ if and only if } \sum_{i=1}^{\infty} (1 - p_i) < \infty \quad (14.2.35)$$

Hence it follows that

$$\sum_{i=1}^{\infty} \left(1 - \frac{r_i}{r_i + 1}\right) = \sum_{i=1}^{\infty} \frac{1}{r_i + 1} < \infty \quad (14.2.36)$$

In particular, this means that $1/(r_i + 1) \rightarrow 0$ as $i \rightarrow \infty$ and hence $r_i \rightarrow \infty$ as $i \rightarrow \infty$. But then

$$\frac{1/(r_i + 1)}{1/r_i} = \frac{r_i}{r_i + 1} \rightarrow 1 \text{ as } i \rightarrow \infty \quad (14.2.37)$$

By the comparison test for infinite series, it follows that

$$\mu = \sum_{i=1}^{\infty} \frac{1}{r_i} < \infty \quad (14.2.38)$$

Computational Exercises

Show directly that the exponential probability density function is a valid probability density function.

Solution

Clearly $f(t) = re^{-rt} > 0$ for $t \in [0, \infty)$. Simple integration that

$$\int_0^{\infty} re^{-rt} dt = 1 \quad (14.2.39)$$

Suppose that the length of a telephone call (in minutes) is exponentially distributed with rate parameter $r = 0.2$. Find each of the following:

1. The probability that the call lasts between 2 and 7 minutes.
2. The median, the first and third quartiles, and the interquartile range of the call length.

Answer

Let X denote the call length.

1. $\mathbb{P}(2 < X < 7) = 0.4237$
2. $q_1 = 1.4384$ $q_2 = 3.4657$ $q_3 = 6.9315$ $q_3 - q_1 = 5.4931$

Suppose that the lifetime of a certain electronic component (in hours) is exponentially distributed with rate parameter $r = 0.001$. Find each of the following:

1. The probability that the component lasts at least 2000 hours.
2. The median, the first and third quartiles, and the interquartile range of the lifetime.

Answer

Let T denote the lifetime

1. $\mathbb{P}(T \geq 2000) = 0.1353$
2. $q_1 = 287.682$ $q_2 = 693.147$ $q_3 = 1386.294$ $q_3 - q_1 = 1098.612$

Suppose that the time between requests to a web server (in seconds) is exponentially distributed with rate parameter $r = 2$. Find each of the following:

1. The mean and standard deviation of the time between requests.
2. The probability that the time between requests is less than 0.5 seconds.
3. The median, the first and third quartiles, and the interquartile range of the time between requests.

Answer

Let T denote the time between requests.

1. $\mathbb{E}(T) = 0.5$, $\text{sd}(T) = 0.5$
2. $\mathbb{P}(T < 0.5) = 0.6321$
3. $q_1 = 0.1438$ $q_2 = 0.3466$ $q_3 = 0.6931$ $q_3 - q_1 = 0.5493$

Suppose that the lifetime X of a fuse (in 100 hour units) is exponentially distributed with $\mathbb{P}(X > 10) = 0.8$. Find each of the following:

1. The rate parameter.
2. The mean and standard deviation.
3. The median, the first and third quartiles, and the interquartile range of the lifetime.

Answer

Let X denote the lifetime.

1. $r = 0.02231$
2. $\mathbb{E}(X) = 44.814$, $\text{sd}(X) = 44.814$
3. $q_1 = 12.8922$ $q_2 = 31.0628$ $q_3 = 62.1257$ $q_3 - q_1 = 49.2334$

The position X of the first defect on a digital tape (in cm) has the exponential distribution with mean 100. Find each of the following:

1. The rate parameter.
2. The probability that $X < 200$ given $X > 150$.
3. The standard deviation.
4. The median, the first and third quartiles, and the interquartile range of the position.

Answer

Let X denote the position of the first defect.

1. $r = 0.01$
2. $\mathbb{P}(X < 200 \mid X > 150) = 0.3935$
3. $\text{sd}(X) = 100$
4. $q_1 = 28.7682$ $q_2 = 69.3147$ $q_3 = 138.6294$ $q_3 - q_1 = 109.8612$

Suppose that X , Y , Z are independent, exponentially distributed random variables with respective parameters a , b , $c \in (0, \infty)$. Find the probability of each of the 6 orderings of the variables.

Proof

1. $\mathbb{P}(X < Y < Z) = \frac{a}{a+b+c} \frac{b}{b+c}$
2. $\mathbb{P}(X < Z < Y) = \frac{a}{a+b+c} \frac{c}{b+c}$
3. $\mathbb{P}(Y < X < Z) = \frac{b}{a+b+c} \frac{a}{a+c}$
4. $\mathbb{P}(Y < Z < X) = \frac{b}{a+b+c} \frac{c}{a+c}$
5. $\mathbb{P}(Z < X < Y) = \frac{c}{a+b+c} \frac{a}{a+b}$
6. $\mathbb{P}(Z < Y < X) = \frac{c}{a+b+c} \frac{b}{a+b}$

This page titled [14.2: The Exponential Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.3: The Gamma Distribution

Basic Theory

We now know that the sequence of inter-arrival times $\mathbf{X} = (X_1, X_2, \dots)$ in the Poisson process is a sequence of independent random variables, each having the exponential distribution with rate parameter r , for some $r > 0$. No other distribution gives the strong renewal assumption that we want: the property that the process probabilistically restarts, independently of the past, at each arrival time and at each fixed time.

The n th arrival time is simply the sum of the first n inter-arrival times:

$$T_n = \sum_{i=0}^n X_i, \quad n \in \mathbb{N} \quad (14.3.1)$$

Thus, the sequence of arrival times $\mathbf{T} = (T_0, T_1, \dots)$ is the partial sum process associated with the sequence of inter-arrival times $\mathbf{X} = (X_1, X_2, \dots)$.

Distribution Functions

Recall that the common probability density function of the inter-arrival times is

$$f(t) = re^{-rt}, \quad 0 \leq t < \infty \quad (14.3.2)$$

Our first goal is to describe the distribution of the n th arrival T_n .

For $n \in \mathbb{N}_+$, T_n has a continuous distribution with probability density function f_n given by

$$f_n(t) = r^n \frac{t^{n-1}}{(n-1)!} e^{-rt}, \quad 0 \leq t < \infty \quad (14.3.3)$$

1. f_n increases and then decreases, with mode at $(n-1)/r$.
2. f_1 is concave upward. f_2 is concave downward and then upward, with inflection point at $t = 2/r$. For $n \geq 2$, f_n is concave upward, then downward, then upward again with inflection points at $t = [(n-1) \pm \sqrt{n-1}]/r$.

Proof

Since T_n is the sum of n independent variables, each with PDF f , the PDF of T_n is the convolution power of f of order n . That is, $f_n = f^{*n}$. A simple induction argument shows that f_n has the form given above. For example,

$$f_2(t) = \int_0^t f(s)f(t-s) ds = \int_0^t re^{-rs} re^{-r(t-s)} ds = \int_0^t r^2 e^{-rt} ds = r^2 t e^{-rt}, \quad 0 \leq t < \infty \quad (14.3.4)$$

Parts (a) and (b) follow from standard calculus.

The distribution with this probability density function is known as the *gamma distribution* with *shape parameter* n and *rate parameter* r . It is also known as the *Erlang distribution*, named for the Danish mathematician Agner Erlang. Again, $1/r$ is the scale parameter, and that term will be justified below. The term *shape parameter* for n clearly makes sense in light of parts (a) and (b) of the last result. The term *rate parameter* for r is inherited from the inter-arrival times, and more generally from the underlying Poisson process itself: the random points are arriving at an average rate of r per unit time. A more general version of the gamma distribution, allowing non-integer shape parameters, is studied in the chapter on Special Distributions. Note that since the arrival times are continuous, the probability of an arrival at any given instant of time is 0.

In the gamma experiment, vary r and n with the scroll bars and watch how the shape of the probability density function changes. For various values of the parameters, run the experiment 1000 times and compare the empirical density function to the true probability density function.

The distribution function and the quantile function of the gamma distribution do not have simple, closed-form expressions. However, it's easy to write the distribution function as a sum.

For $n \in \mathbb{N}_+$, T_n has distribution function F_n given by

$$F_n(t) = 1 - \sum_{k=0}^{n-1} e^{-rt} \frac{(rt)^k}{k!}, \quad t \in [0, \infty) \quad (14.3.5)$$

Proof

Note that

$$F_n(t) = \int_0^t f_n(s) ds = \int_0^n r^n \frac{s^{n-1}}{(n-1)!} e^{-rs} \quad (14.3.6)$$

The result follows by repeated integration by part.

Open the special distribution calculator, select the gamma distribution, and select CDF view. Vary the parameters and note the shape of the distribution and quantile functions. For selected values of the parameters, compute the quantiles.

Moments

The mean, variance, and moment generating function of T_n can be found easily from the representation as a sum of independent exponential variables.

The mean and variance of T_n are.

1. $\mathbb{E}(T_n) = n/r$
2. $\text{var}(T_n) = n/r^2$

Proof

Recall that the exponential distribution with rate parameter r has mean $1/r$ and variance $1/r^2$.

1. The expected value of a sum is the sum of the expected values, so $\mathbb{E}(T_n) = n/r$.
2. The variance of a sum of *independent* variables is the sum of the variances, so $\text{var}(T_n) = n/r^2$.

For $k \in \mathbb{N}$, the moment of order k of T_n is

$$\mathbb{E}(T_n^k) = \frac{(k+n-1)!}{(n-1)!} \frac{1}{r^k} \quad (14.3.7)$$

Proof

Using the standard change of variables theorem,

$$\mathbb{E}(T_n^k) = \int_0^\infty t^k f_n(t) dt = \frac{r^{n-1}}{(n-1)!} \int_0^\infty t^{k+n-1} r e^{-rt} dt \quad (14.3.8)$$

But the integral on the right is the moment of order $k+n-1$ for the exponential distribution, which we showed in the last section is $(k+n-1)!/r^{k+n-1}$. Simplifying gives the result.

More generally, the moment of order $k > 0$ (not necessarily an integer) is

$$\mathbb{E}(T_n^k) = \frac{\Gamma(k+n)}{\Gamma(n)} \frac{1}{r^k} \quad (14.3.9)$$

where Γ is the gamma function.

In the gamma experiment, vary r and n with the scroll bars and watch how the size and location of the mean \pm standard deviation bar changes. For various values of r and n , run the experiment 1000 times and compare the empirical moments to the true moments.

Our next result gives the skewness and kurtosis of the gamma distribution.

The skewness and kurtosis of T_n are

1. $\text{skew}(X) = \frac{2}{\sqrt{n}}$
2. $\text{kurt}(X) = 3 + \frac{6}{n}$

Proof

These results follow from the [moment results](#) above and the computational formulas for skewness and kurtosis.

In particular, note that the gamma distribution is positively skewed but $\text{skew}(X) \rightarrow 0$ and as $n \rightarrow \infty$. Recall also that the *excess kurtosis* is $\text{kurt}(T_n) - 3 = \frac{6}{n} \rightarrow 0$ as $n \rightarrow \infty$. This result is related to the convergence of the gamma distribution to the normal, discussed below. Finally, note that the skewness and kurtosis do not depend on the rate parameter r . This is because, as we show below, $1/r$ is a scale parameter.

The moment generating function of T_n is

$$M_n(s) = \mathbb{E}(e^{sT_n}) = \left(\frac{r}{r-s}\right)^n, \quad -\infty < s < r \quad (14.3.10)$$

Proof

Recall that the MGF of a sum of independent variables is the product of the corresponding MGFs. We showed in the last section that the exponential distribution with parameter r has MGF $s \mapsto r/(r-s)$ for $-\infty < s < r$.

The moment generating function can also be used to derive the moments of the gamma distribution given above—recall that $M_n^{(k)}(0) = \mathbb{E}(T_n^k)$.

Estimating the Rate

In many practical situations, the rate r of the process is unknown and must be estimated based on data from the process. We start with a natural estimate of the scale parameter $1/r$. Note that

$$M_n = \frac{T_n}{n} = \frac{1}{n} \sum_{i=1}^n X_i \quad (14.3.11)$$

is the sample mean of the first n inter-arrival times (X_1, X_2, \dots, X_n) . In statistical terms, this sequence is a random sample of size n from the exponential distribution with rate r .

M_n satisfies the following properties:

1. $\mathbb{E}(M_n) = \frac{1}{r}$
2. $\text{var}(M_n) = \frac{1}{nr^2}$
3. $M_n \rightarrow \frac{1}{r}$ as $n \rightarrow \infty$ with probability 1

Proof

Parts (a) and (b) follow from the expected value of T_n and standard properties. Part (c) is the strong law of large numbers.

In statistical terms, part (a) means that M_n is an *unbiased* estimator of $1/r$ and hence the variance in part (b) is the *mean square error*. Part (b) means that M_n is a *consistent* estimator of $1/r$ since $\text{var}(M_n) \downarrow 0$ as $n \rightarrow \infty$. Part (c) is a stronger form of consistency. In general, the sample mean of a random sample from a distribution is an unbiased and consistent estimator of the distribution mean. On the other hand, a natural estimator of r itself is $1/M_n = n/T_n$. However, this estimator is positively biased.

$$\mathbb{E}(n/T_n) \geq r.$$

Proof

This follows immediately from Jensen's inequality since $x \mapsto 1/x$ is concave upward on $(0, \infty)$.

Properties and Connections

Scaling

As noted above, the gamma distribution is a scale family.

Suppose that T has the gamma distribution with rate parameter $r \in (0, \infty)$ and shape parameter $n \in \mathbb{N}_+$. If $c \in (0, \infty)$ then cT has the gamma distribution with rate parameter r/c and shape parameter n .

Proof

The moment generating function of cT is

$$\mathbb{E}[e^{s(cT)}] = \mathbb{E}[e^{(cs)T}] = \left(\frac{r}{r - cs} \right)^n = \left(\frac{r/c}{r/c - s} \right)^n, \quad s < \frac{r}{c} \quad (14.3.12)$$

The scaling property also follows from the fact that the gamma distribution governs the arrival times in the Poisson process. A time change in a Poisson process clearly does not change the strong renewal property, and hence results in a new Poisson process.

General Exponential Family

The gamma distribution is also a member of the general exponential family of distributions.

Suppose that T has the gamma distribution with shape parameter $n \in \mathbb{N}_+$ and rate parameter $r \in (0, \infty)$. Then T has a two parameter general exponential distribution with natural parameters $n - 1$ and $-r$, and natural statistics $\ln(T)$ and T .

Proof

This follows from the form of the PDF and the definition of the general exponential family:

$$f(t) = r^n \frac{t^{n-1}}{(n-1)!} e^{-rt} = \frac{r^n}{(n-1)!} \exp[(n-1)\ln(t) - rt], \quad t \in (0, \infty) \quad (14.3.13)$$

Increments

A number of important properties flow from the fact that the sequence of arrival times $\mathbf{T} = (T_0, T_1, \dots)$ is the partial sum process associated with the sequence of independent, identically distributed inter-arrival times $\mathbf{X} = (X_1, X_2, \dots)$.

The arrival time sequence \mathbf{T} has stationary, independent increments:

1. If $m < n$ then $T_n - T_m$ has the same distribution as T_{n-m} , namely the gamma distribution with shape parameter $n - m$ and rate parameter r .
2. If $n_1 < n_2 < n_3 < \dots$ then $(T_{n_1}, T_{n_2} - T_{n_1}, T_{n_3} - T_{n_2}, \dots)$ is an independent sequence.

Proof

The stationary and independent increments properties hold for any partial sum process associated with an independent, identically distributed sequence.

Of course, the stationary and independent increments properties are related to the fundamental “renewal” assumption that we started with. If we fix $n \in \mathbb{N}_+$, then $(T_n - T_n, T_{n+1} - T_n, T_{n+2} - T_n, \dots)$ is independent of (T_1, T_2, \dots, T_n) and has the same distribution as (T_0, T_1, T_2, \dots) . That is, if we “restart the clock” at time T_n , then the process in the future looks just like the original process (in a probabilistic sense) and is independent of the past. Thus, we have our second characterization of the Poisson process.

A process of random points in time is a Poisson process with rate $r \in (0, \infty)$ if and only if the arrival time sequence \mathbf{T} has stationary, independent increments, and for $n \in \mathbb{N}_+$, T_n has the gamma distribution with shape parameter n and rate parameter r .

Sums

The gamma distribution is closed with respect to sums of independent variables, as long as the rate parameter is fixed.

Suppose that V has the gamma distribution with shape parameter $m \in \mathbb{N}_+$ and rate parameter $r > 0$, W has the gamma distribution with shape parameter $n \in \mathbb{N}_+$ and rate parameter r , and that V and W are independent. Then $V + W$ has the gamma distribution with shape parameter $m + n$ and rate parameter r .

Proof

There are at least three different proofs of this fundamental result. Perhaps the best is a probabilistic proof based on the Poisson process. We start with an IID sequence \mathbf{X} of independent exponentially distributed variables, each with rate parameter r . Then we can associate V with T_m and W with $T_{m+n} - T_m$ so that $V + W$ becomes T_{m+n} . The result now follows from the previous theorem.

Another simple proof uses moment generating functions. Recall again that the MGF of $V + W$ is the product of the MGFs of V and of W . A third, analytic proof uses convolution. Recall again that the PDF of $V + W$ is the convolution of the PDFs of V and of W .

Normal Approximation

In the gamma experiment, vary r and n with the scroll bars and watch how the shape of the probability density function changes. Now set $n = 10$ and for various values of r run the experiment 1000 times and compare the empirical density function to the true probability density function.

Even though you are restricted to relatively small values of n in the app, note that the probability density function of the n th arrival time becomes more bell shaped as n increases (for r fixed). This is yet another application of the central limit theorem, since T_n is the sum of n independent, identically distributed random variables (the inter-arrival times).

The distribution of the random variable Z_n below converges to the standard normal distribution as $n \rightarrow \infty$:

$$Z_n = \frac{rT_n - n}{\sqrt{n}} \quad (14.3.14)$$

Proof

Z_n is the standard score associated with T_n , so the result follows from the central limit theorem.

Connection to Bernoulli Trials

We return to the analogy between the Bernoulli trials process and the Poisson process that started in the Introduction and continued in the last section on the Exponential Distribution. If we think of the successes in a sequence of Bernoulli trials as random points in discrete time, then the process has the same strong renewal property as the Poisson process, but restricted to discrete time. That is, at each fixed time and at each arrival time, the process “starts over”, independently of the past. In Bernoulli trials, the time of the n th arrival has the negative binomial distribution with parameters n and p (the success probability), while in the Poisson process, as we now know, the time of the n th arrival has the gamma distribution with parameters n and r (the rate). Because of this strong analogy, we expect a relationship between these two processes. In fact, we have the same type of limit as with the geometric and exponential distributions.

Fix $n \in \mathbb{N}_+$ and suppose that for each $m \in \mathbb{N}_+$ $T_{m,n}$ has the negative binomial distribution with parameters n and $p_m \in (0, 1)$, where $mp_m \rightarrow r \in (0, \infty)$ as $m \rightarrow \infty$. Then the distribution of $T_{m,n}/m$ converges to the gamma distribution with parameters n and r as $m \rightarrow \infty$.

Proof

Suppose that X_m has the geometric distribution on \mathbb{N}_+ with success parameter p_m . We know from our convergence result in the last section that the distribution of X_m/m converges to the exponential distribution with rate parameter r as $m \rightarrow \infty$. It follows that if M_m denotes the moment generating function of X_m/m , then $M_m(s) \rightarrow r/(r-s)$ as $m \rightarrow \infty$ for $s < r$. But then M_m^n is the MGF of $T_{m,n}/m$ and clearly

$$M_m^n(s) \rightarrow \left(\frac{r}{r-s} \right)^n \quad (14.3.15)$$

as $m \rightarrow \infty$ for $s < r$. The expression on the right is the MGF of the gamma distribution with shape parameter n and rate parameter r .

Computational Exercises

Suppose that customers arrive at a service station according to the Poisson model, at a rate of $r = 3$ per hour. Relative to a given starting time, find the probability that the second customer arrives sometime after 1 hour.

Answer

0.1991

Defects in a type of wire follow the Poisson model, with rate 1 per 100 meter. Find the probability that the 5th defect is located between 450 and 550 meters.

Answer

0.1746

Suppose that requests to a web server follow the Poisson model with rate $r = 5$. Relative to a given starting time, compute the mean and standard deviation of the time of the 10th request.

Answer

2, 0.6325

Suppose that Y has a gamma distribution with mean 40 and standard deviation 20. Find the shape parameter n and the rate parameter r .

Answer

$r = 1/10, n = 4$

Suppose that accidents at an intersection occur according to the Poisson model, at a rate of 8 per year. Compute the normal approximation to the event that the 10th accident (relative to a given starting time) occurs within 2 years.

Answer

0.5752

In the gamma experiment, set $n = 5$ and $r = 2$. Run the experiment 1000 times and compute the following:

1. $\mathbb{P}(1.5 \leq T_t \leq 3)$
2. The relative frequency of the event $\{1.5 \leq T_5 \leq 3\}$
3. The normal approximation to $\mathbb{P}(1.5 \leq T_5 \leq 3)$

Answer

1. 0.5302
3. 0.4871

Suppose that requests to a web server follow the Poisson model. Starting at 12:00 noon on a certain day, the requests are logged. The 100th request comes at 12:15. Estimate the rate of the process.

Answer

$r = 6.67$ hits per minute

This page titled [14.3: The Gamma Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.4: The Poisson Distribution

Basic Theory

Recall that in the Poisson model, $\mathbf{X} = (X_1, X_2, \dots)$ denotes the sequence of inter-arrival times, and $\mathbf{T} = (T_0, T_1, T_2, \dots)$ denotes the sequence of arrival times. Thus \mathbf{T} is the partial sum process associated with \mathbf{X} :

$$T_n = \sum_{i=0}^n X_i, \quad n \in \mathbb{N} \quad (14.4.1)$$

Based on the strong renewal assumption, that the process restarts at each fixed time and each arrival time, independently of the past, we now know that \mathbf{X} is a sequence of independent random variables, each with the exponential distribution with rate parameter r , for some $r \in (0, \infty)$. We also know that \mathbf{T} has stationary, independent increments, and that for $n \in \mathbb{N}_+$, T_n has the gamma distribution with rate parameter r and scale parameter n . Both of the statements characterize the Poisson process with rate r .

Recall that for $t \geq 0$, N_t denotes the number of arrivals in the interval $(0, t]$, so that $N_t = \max\{n \in \mathbb{N} : T_n \leq t\}$. We refer to $\mathbf{N} = (N_t : t \geq 0)$ as the *counting process*. In this section we will show that N_t has a *Poisson distribution*, named for Simeon Poisson, one of the most important distributions in probability theory. Our exposition will alternate between properties of the distribution and properties of the counting process. The two are intimately intertwined. It's not too much of an exaggeration to say that wherever there is a *Poisson distribution*, there is a *Poisson process* lurking in the background.

Probability density function.

Recall that the probability density function of the n th arrival time T_n is

$$f_n(t) = r^n \frac{t^{n-1}}{(n-1)!} e^{-rt}, \quad 0 \leq t < \infty \quad (14.4.2)$$

We can find the distribution of N_t because of the inverse relation between \mathbf{N} and \mathbf{T} . In particular, recall that

$$N_t \geq n \iff T_n \leq t, \quad t \in (0, \infty), \quad n \in \mathbb{N}_+ \quad (14.4.3)$$

since both events mean that there are at least n arrivals in $(0, t]$.

For $t \in [0, \infty)$, the probability density function of N_t is given by

$$\mathbb{P}(N_t = n) = e^{-rt} \frac{(rt)^n}{n!}, \quad n \in \mathbb{N} \quad (14.4.4)$$

Proof

Using the inverse relationship noted above, and integration by parts, we have

$$\mathbb{P}(N_t \geq n) = \mathbb{P}(T_n \leq t) = \int_0^t f_n(s) ds = 1 - \sum_{k=0}^{n-1} e^{-rt} \frac{(rt)^k}{k!}, \quad n \in \mathbb{N} \quad (14.4.5)$$

For $n \in \mathbb{N}$ we have $\mathbb{P}(N_t = n) = \mathbb{P}(N_t \geq n) - \mathbb{P}(N_t \geq n+1)$. Simplifying gives the result.

Note that the distribution of N_t depends on the parameters r and t only through the product rt . The distribution is called the *Poisson distribution* with parameter rt .

In the Poisson experiment, vary r and t with the scroll bars and note the shape of the probability density function. For various values of r and t , run the experiment 1000 times and compare the relative frequency function to the probability density function.

In general, a random variable N taking values in \mathbb{N} is said to have the Poisson distribution with parameter $a \in (0, \infty)$ if it has the probability density function

$$g(n) = e^{-a} \frac{a^n}{n!}, \quad n \in \mathbb{N} \quad (14.4.6)$$

1. $g(n-1) < g(n)$ if and only if $n < a$.
2. If $a \notin \mathbb{N}_+$, there is a single mode at $\lfloor a \rfloor$.
3. If $a \in \mathbb{N}_+$, there are consecutive modes at $a-1$ and a .

Proof

Part (a) follows from simple algebra, and similarly, $g(n-1) = g(n)$ if and only if $n = a$ (and thus $a \in \mathbb{N}_+$). Parts (b) and (c) then follow.

The Poisson distribution does not have simple closed-form distribution or quantile functions. Trivially, we can write the distribution function as a sum of the probability density function.

The Poisson distribution with parameter a has distribution function G given by

$$G(n) = \sum_{k=0}^n e^{-a} \frac{a^k}{k!}, \quad n \in \mathbb{N} \quad (14.4.7)$$

Open the special distribution calculator, select the Poisson distribution, and select CDF view. Vary the parameter and note the shape of the distribution and quantile functions. For various values of the parameter, compute the quantiles.

Sometimes it's convenient to allow the parameter a to be 0. This *degenerate* Poisson distribution is simply point mass at 0. That is, with the usual conventions regarding nonnegative integer powers of 0, the probability density function g above reduces to $g(0) = 1$ and $g(n) = 0$ for $n \in \mathbb{N}_+$.

Moments

Suppose that N has the Poisson distribution with parameter $a > 0$. Naturally we want to know the mean, variance, skewness and kurtosis, and the probability generating function of N . The easiest moments to compute are the factorial moments. For this result, recall the falling power notation for the number of permutations of size k chosen from a population of size n :

$$n^{(k)} = n(n-1) \cdots [n-(k+1)] \quad (14.4.8)$$

The factorial moment of N of order $k \in \mathbb{N}$ is $\mathbb{E}[N^{(k)}] = a^k$.

Proof

Using the standard change of variables formula for expected value,

$$\mathbb{E}[N^{(k)}] = \sum_{n=0}^{\infty} n^{(k)} e^{-a} \frac{a^n}{n!} = e^{-a} a^k \sum_{n=k}^{\infty} \frac{a^{n-k}}{(n-k)!} = e^{-a} a^k e^a = a^k \quad (14.4.9)$$

The mean and variance of N are the parameter a .

1. $\mathbb{E}(N) = a$
2. $\text{var}(N) = a$

Proof

1. This follows directly from the first factorial moment: $\mathbb{E}(N) = \mathbb{E}[N^{(1)}] = a$.
2. Note that $\mathbb{E}(N^2) = \mathbb{E}[N^{(2)}] + \mathbb{E}(N) = a^2 + a$.

Open the special distribution simulator and select the Poisson distribution. Vary the parameter and note the location and size of the mean \pm standard deviation bar. For selected values of the parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the distribution mean and standard deviation.

The skewness and kurtosis of N are

1. $\text{skew}(N) = 1/\sqrt{a}$
2. $\text{kurt}(N) = 3 + 1/a$

Proof

These results follow from the computational formulas for skewness and kurtosis and the results for factorial moments above. Specifically,

1. $\mathbb{E}(N^3) = a^3 + 3a^2 + a$ and $\text{skew}(N) = [\mathbb{E}(N^3) - 3a^2 - a^3]/a^{3/2}$
2. $\mathbb{E}(N^4) = a^4 + 6a^3 + 7a^2 + a$ and $\text{kurt}(N) = [\mathbb{E}(N^4) - 4a\mathbb{E}(N^3) + 6a^3 + 3a^4]/a^2$

Note that the Poisson distribution is positively skewed, but $\text{skew}(N) \rightarrow 0$ as $n \rightarrow \infty$. Recall also that the *excess kurtosis* is $\text{kurt}(N) - 3 = +1/a \rightarrow 0$ as $n \rightarrow \infty$. This limit is related to the convergence of the Poisson distribution to the normal, discussed [below](#).

Open the special distribution simulator and select the Poisson distribution. Vary the parameter and note the shape of the probability density function in the context of the results on skewness and kurtosis above.

The probability generating function P of N is given by

$$P(s) = \mathbb{E}(s^N) = e^{a(s-1)}, \quad s \in \mathbb{R} \quad (14.4.10)$$

Proof

Using the change of variables formula again,

$$\mathbb{E}(s^N) = \sum_{n=0}^{\infty} s^n e^{-a} \frac{a^n}{n!} = e^{-a} \sum_{n=0}^{\infty} \frac{(as)^n}{n!} = e^{-a} e^{as}, \quad s \in \mathbb{R} \quad (14.4.11)$$

Returning to the Poisson counting process $\mathbf{N} = (N_t : t \geq 0)$ with rate parameter r , it follows that $\mathbb{E}(N_t) = rt$ and $\text{var}(N_t) = rt$ for $t \geq 0$. Once again, we see that r can be interpreted as the average arrival *rate*. In an interval of length t , we expect about rt arrivals.

In the Poisson experiment, vary r and t with the scroll bars and note the location and size of the mean \pm standard deviation bar. For various values of r and t , run the experiment 1000 times and compare the sample mean and standard deviation to the distribution mean and standard deviation, respectively.

Estimating the Rate

Suppose again that we have a Poisson process with rate $r \in (0, \infty)$. In many practical situations, the rate r is unknown and must be estimated based on observing data. For fixed $t > 0$, a natural estimator of the rate r is N_t/t .

The mean and variance of N_t/t are

1. $\mathbb{E}(N_t/t) = r$
2. $\text{var}(N_t/t) = r/t$

Proof

These results follow easily from $\mathbb{E}(N_t) = \text{var}(N_t) = rt$ and basic properties of expected value and variance.

Part (a) means that the estimator is unbiased. Since this is the case, the variance in part (b) gives the mean square error. Since $\text{var}(N_t)$ decreases to 0 as $t \rightarrow \infty$, the estimator is consistent.

Additional Properties and Connections

Increments and Characterizations

Let's explore the basic renewal assumption of the Poisson model in terms of the counting process $\mathbf{N} = (N_t : t \geq 0)$. Recall that N_t is the number of arrivals in the interval $(0, t]$, so it follows that if $s, t \in [0, \infty)$ with $s < t$, then $N_t - N_s$ is the number of arrivals in the interval $(s, t]$. Of course, the arrival times have continuous distributions, so the probability that an arrival occurs at a specific point t is 0. Thus, it does not really matter if we write the interval above as $(s, t]$, (s, t) , $[s, t)$ or $[s, t]$.

The process \mathbf{N} has stationary, independent increments.

1. If $s, t \in [0, \infty)$ with $s < t$ then $N_t - N_s$ has the same distribution as N_{t-s} , namely Poisson with parameter $r(t-s)$.
2. If $t_1, t_2, t_3, \dots \in [0, \infty)$ with $t_1 < t_2 < t_3 < \dots$ then $(N_{t_1}, N_{t_2} - N_{t_1}, N_{t_3} - N_{t_2}, \dots)$ is an independent sequence.

Statements about the increments of the counting process can be expressed more elegantly in terms of our more general counting process. Recall that for $A \subseteq [0, \infty)$ (measurable of course), $N(A)$ denotes the number of random points in A :

$$N(A) = \# \{n \in \mathbb{N}_+ : T_n \in A\} \quad (14.4.12)$$

and so in particular, $N_t = N(0, t]$. Thus, note that $t \mapsto N_t$ is a (random) distribution function and $A \mapsto N(A)$ is the (random) measure associated with this distribution function. Recall also that λ denotes the standard length (Lebesgue) measure on $[0, \infty)$. Here is our third characterization of the Poisson process.

A process of random points in time is a Poisson process with rate $r \in (0, \infty)$ if and only if the following properties hold:

1. If $A \subseteq [0, \infty)$ is measurable then $N(A)$ has the Poisson distribution with parameter $r\lambda(A)$.
2. If $\{A_i : i \in I\}$ is a countable, disjoint collection of measurable sets in $[0, \infty)$ then $\{N(A_i) : i \in I\}$ is a set of independent variables.

From a modeling point of view, the assumptions of stationary, independent increments are ones that might be reasonably made. But the assumption that the increments have *Poisson* distributions does not seem as clear. Our next characterization of the process is more primitive in a sense, because it just imposes some limiting assumptions (in addition to stationary, independent increments).

A process of random points in time is a Poisson process with rate $r \in (0, \infty)$ if and only if the following properties hold:

1. If $A, B \subseteq [0, \infty)$ are measurable and $\lambda(A) = \lambda(B)$, then $N(A)$ and $N(B)$ have the same distribution.
2. If $\{A_i : i \in I\}$ is a countable, disjoint collection of measurable sets in $[0, \infty)$ then $\{N(A_i) : i \in I\}$ is a set of independent variables.
3. If $A_n \subseteq [0, \infty)$ is measurable and $\lambda(A_n) > 0$ for $n \in \mathbb{N}_+$, and if $\lambda(A_n) \rightarrow 0$ as $n \rightarrow \infty$ then

$$\frac{\mathbb{P}[N(A_n) = 1]}{\lambda(A_n)} \rightarrow r \text{ as } n \rightarrow \infty \quad (14.4.13)$$

$$\frac{\mathbb{P}[N(A_n) > 1]}{\lambda(A_n)} \rightarrow 0 \text{ as } n \rightarrow \infty \quad (14.4.14)$$

Proof

As usual, let $N_t = N(0, t]$, the number of arrivals in $(0, t]$, and in addition let $P_n(t) = \mathbb{P}(N_t = n)$ for $t \geq 0$ and $n \in \mathbb{N}$. Note first that P_0 satisfies the following differential equation and initial condition:

$$P'_0(t) = -rP_0(t), \quad t > 0; \quad P_0(0) = 1 \quad (14.4.15)$$

Hence $P_0(t) = e^{-rt}$ for $t \geq 0$. Next for $n \in \mathbb{N}_+$, P_n satisfies the following differential equation and initial condition

$$P'_n(t) = -rP_n(t) + rP_{n-1}(t), \quad t > 0; \quad P_n(0) = 0 \quad (14.4.16)$$

Hence $P_n(t) = e^{-rt}(rt)^n/n!$ for $t \geq 0$ and therefore N_t has the Poisson distribution with parameter rt .

Of course, part (a) is the stationary assumption and part (b) the independence assumption. The first limit in (c) is sometimes called the *rate property* and the second limit the *sparseness property*. In a “small” time interval of length dt , the probability of a single random point is approximately $r dt$, and the probability of two or more random points is negligible.

Sums

Suppose that N and M are independent random variables, and that N has the Poisson distribution with parameter $a \in (0, \infty)$ and M has the Poisson distribution with parameter $b \in (0, \infty)$. Then $N + M$ has the Poisson distribution with parameter $a + b$.

Proof from the Poisson process

There are several ways to prove this result, but the one that gives the most insight is a probabilistic proof based on the Poisson process. Thus suppose that $\mathbf{N} = (N_t : t \geq 0)$ is a Poisson counting process with rate 1. We can associate N with N_a and M with $N_{a+b} - N_a$, since these have the correct distributions and are independent. But then $N + M$ is N_{a+b} .

Proof from probability generating functions

From our result [above](#), M has PGF $P(s) = e^{a(s-1)}$ for $s \in \mathbb{R}$, and N has PGF $Q(s) = e^{b(s-1)}$ for $s \in \mathbb{R}$. Hence $M + N$ has PGF $P(s)Q(s) = e^{(a+b)(s-1)}$ for $s \in \mathbb{R}$. But this is the PGF of the Poisson distribution with parameter $a + b$.

Proof from convolution

From our results above, M has PDF $g(n) = e^{-a}a^n/n!$ for $n \in \mathbb{N}$, and N has PDF $h(n) = e^{-b}b^n/n!$ for $n \in \mathbb{N}$. Hence the PDF of $M + N$ is the convolution $g * h$. For $n \in \mathbb{N}$,

$$(g * h)(n) = \sum_{k=0}^n g(k)h(n-k) = \sum_{k=0}^n e^{-a} \frac{a^k}{k!} e^{-b} \frac{b^{n-k}}{(n-k)!} = e^{-(a+b)} \frac{1}{n!} \sum_{k=0}^n \frac{n!}{k!(n-k)!} a^k b^{n-k} \quad (14.4.17)$$

By the binomial theorem, the last sum is $(a + b)^n$.

From the last theorem, it follows that the Poisson distribution is infinitely divisible. That is, a Poisson distributed variable can be written as the sum of an arbitrary number of independent, identically distributed (in fact also Poisson) variables.

Suppose that N has the Poisson distribution with parameter $a \in (0, \infty)$. Then for $n \in \mathbb{N}_+$, N has the same distribution as $\sum_{i=1}^n N_i$ where (N_1, N_2, \dots, N_n) are independent, and each has the Poisson distribution with parameter a/n .

Normal Approximation

Because of the representation as a sum of independent, identically distributed variables, it's not surprising that the Poisson distribution can be approximated by the normal.

Suppose that N_t has the Poisson distribution with parameter $t > 0$. Then the distribution of the variable below converges to the standard normal distribution as $t \rightarrow \infty$.

$$Z_t = \frac{N_t - t}{\sqrt{t}} \quad (14.4.18)$$

Proof

As usual, we can assume that $(N_t : t \geq 0)$ is the Poisson counting process with rate 1. Note that Z_t is simply the standard score associated with N_t . For $n \in \mathbb{N}_+$, N_n is the sum of n independent variables, each with the Poisson distribution with parameter 1. Thus, from the central limit theorem, the distribution of Z_n converges to the standard normal distribution as $n \rightarrow \infty$. For general $t \in [0, \infty)$, it's possible to write $Z_t = Z_n + W_t$ where $n = \lfloor t \rfloor$ and $W_t \rightarrow 0$ as $t \rightarrow \infty$ (in probability and hence in distribution).

Thus, if N has the Poisson distribution with parameter a , and a is “large”, then the distribution of N is approximately normal with mean a and standard deviation \sqrt{a} . When using the normal approximation, we should remember to use the continuity correction, since the Poisson is a discrete distribution.

In the Poisson experiment, set $r = t = 1$. Increase t and note how the graph of the probability density function becomes more bell-shaped.

General Exponential

The Poisson distribution is a member of the general exponential family of distributions. This fact is important in various statistical procedures.

Suppose that N has the Poisson distribution with parameter $a \in (0, \infty)$. This distribution is a one-parameter exponential family with natural parameter $\ln(a)$ and natural statistic N .

Proof

This follows from the form of the Poisson PDF:

$$g(n) = e^{-a} \frac{a^n}{n!} = \frac{e^{-a}}{n!} \exp[n \ln(a)], \quad n \in \mathbb{N} \quad (14.4.19)$$

The Uniform Distribution

The Poisson process has some basic connections to the uniform distribution. Consider again the Poisson process with rate $r > 0$. As usual, $\mathbf{T} = (T_0, T_1, \dots)$ denotes the arrival time sequence and $\mathbf{N} = (N_t : t \geq 0)$ the counting process.

For $t > 0$, the conditional distribution of T_1 given $N_t = 1$ is uniform on the interval $(0, t]$.

Proof

Given $N_t = 1$ (one arrival in $(0, t]$) the arrival time T_1 takes values in $(0, t]$. From independent and stationary increments properties,

$$\mathbb{P}(T_1 \leq s \mid N_t = 1) = \mathbb{P}(N_s = 1, N_t - N_s = 0 \mid N_t = 1) = \frac{\mathbb{P}(N_s = 1, N_t - N_s = 0)}{\mathbb{P}(N_t = 1)} = \frac{\mathbb{P}(N_s = 1)\mathbb{P}(N_t - N_s = 0)}{\mathbb{P}(N_t = 1)} \quad (14.4.20)$$

Hence using the Poisson distribution,

$$\mathbb{P}(T_1 \leq s \mid N_t = 1) = \frac{e^{-rs} s e^{-r(t-s)}}{e^{-rt} t} = \frac{s}{t}, \quad 0 < s \leq t \quad (14.4.21)$$

More generally, for $t > 0$ and $n \in \mathbb{N}_+$, the conditional distribution of (T_1, T_2, \dots, T_n) given $N_t = n$ is the same as the distribution of the order statistics of a random sample of size n from the uniform distribution on the interval $(0, t]$.

Heuristic proof

Suppose that $0 < t_1 < t_2 < \dots < t_n < t$. On the event $N_t = n$, the probability density of (T_1, T_2, \dots, T_n) at (t_1, t_2, \dots, t_n) is the probability density function of independent inter-arrival times $(t_1, t_2 - t_1, \dots, t_n - t_{n-1})$ times the probability of no arrivals in the interval (t_n, t) . Hence given $N_t = n$, the conditional density of (T_1, T_2, \dots, T_n) at (t_1, t_2, \dots, t_n) is

$$\frac{r e^{-rt_1} r e^{-r(t_2 - t_1)} \dots r e^{-r(t_n - t_{n-1})} e^{-r(t - t_n)}}{e^{-rt} (rt)^n / n!} = \frac{n!}{t^n} \quad (14.4.22)$$

But this is the PDF of the order statistics from a sample of size n from the uniform distribution on $[0, t]$.

Note that the conditional distribution in the last result is independent of the rate r . This means that, in a sense, the Poisson model gives the most “random” distribution of points in time.

The Binomial Distribution

The Poisson distribution has important connections to the binomial distribution. First we consider a conditional distribution based on the number of arrivals of a Poisson process in a given interval, as we did in the last subsection.

Suppose that $(N_t : t \in [0, \infty))$ is a Poisson counting process with rate $r \in (0, \infty)$. If $s, t \in (0, \infty)$ with $s < t$, and $n \in \mathbb{N}_+$, then the conditional distribution of N_s given $N_t = n$ is binomial with trial parameter n and success parameter $p = s/t$.

Proof

Note that given $N_t = n$, the number of arrivals N_s in $(0, s]$ takes values in $\{0, 1, \dots, n\}$. Again, the stationary and independent increments properties are critical for the proof.

$$\mathbb{P}(N_s = k \mid N_t = n) = \frac{\mathbb{P}(N_s = k, N_t = n)}{\mathbb{P}(N_t = n)} = \frac{\mathbb{P}(N_s = k, N_t - N_s = n - k)}{\mathbb{P}(N_t = n)} = \frac{\mathbb{P}(N_s = k)\mathbb{P}(N_t - N_s = n - k)}{\mathbb{P}(N_t = n)} \quad (14.4.23)$$

Substituting into the Poisson PDFs gives

$$\mathbb{P}(N_s = k \mid N_t = n) = \frac{(e^{-rs}(rs)^k/k!)(e^{-r(t-s)}[(r(t-s))^{n-k}/(n-k)!])}{e^{-rt}(rt)^n/n!} = \frac{n!}{k!(n-k)!} \left(\frac{s}{t}\right)^k \left(1 - \frac{s}{t}\right)^{n-k} \quad (14.4.24)$$

Note again that the conditional distribution in the last result does not depend on the rate r . Given $N_t = n$, each of the n arrivals, independently of the others, falls into the interval $(0, s]$ with probability s/t and into the interval $(s, t]$ with probability $1 - s/t = (t - s)/t$. Here is essentially the same result, outside of the context of the Poisson process.

Suppose that M has the Poisson distribution with parameter $a \in (0, \infty)$, N has the Poisson distribution with parameter $b \in (0, \infty)$, and that M and N are independent. Then the conditional distribution of M given $M + N = n$ is binomial with parameters n and $p = a/(a + b)$.

Proof

The proof is essentially the same as the previous theorem, with minor modifications. First recall from the result [above](#) that $M + N$ has the Poisson distribution with parameter $a + b$. For $n, k \in \mathbb{N}$ with $k \leq n$,

$$\mathbb{P}(M = k \mid M + N = n) = \frac{\mathbb{P}(M = k, M + N = n)}{\mathbb{P}(M + N = n)} = \frac{\mathbb{P}(M = k, N = n - k)}{\mathbb{P}(N + M = n)} = \frac{\mathbb{P}(M = k)\mathbb{P}(N = n - k)}{\mathbb{P}(M + N = n)} \quad (14.4.25)$$

Substituting into the Poisson PDFs gives

$$\mathbb{P}(M = k \mid M + N = n) = \frac{(e^{-a}a^k/k!)(e^{-b}b^{n-k}/(n-k)!)}{e^{-(a+b)}(a+b)^n/n!} = \frac{n!}{k!(n-k)!} \left(\frac{a}{a+b}\right)^k \left(1 - \frac{a}{a+b}\right)^{n-k} \quad (14.4.26)$$

More importantly, the Poisson distribution is the limit of the binomial distribution in a certain sense. As we will see, this convergence result is related to the analogy between the Bernoulli trials process and the Poisson process that we discussed in the Introduction, the section on the inter-arrival times, and the section on the arrival times.

Suppose that $p_n \in (0, 1)$ for $n \in \mathbb{N}_+$ and that $np_n \rightarrow a \in (0, \infty)$ as $n \rightarrow \infty$. Then the binomial distribution with parameters n and p_n converges to the Poisson distribution with parameter a as $n \rightarrow \infty$. That is, for fixed $k \in \mathbb{N}$,

$$\binom{n}{k} p_n^k (1 - p_n)^{n-k} \rightarrow e^{-a} \frac{a^k}{k!} \text{ as } n \rightarrow \infty \quad (14.4.27)$$

Direct proof

The binomial PDF with parameters n and p_n at $k \in \{0, 1, \dots, n\}$ can be written as

$$\frac{1}{k!} [np_n] [(n-1)p_n] \cdots [(n-k+1)p_n] \left(1 - \frac{np_n}{n}\right)^{n-k} \quad (14.4.28)$$

But $(n-j)p_n \rightarrow a$ as $n \rightarrow \infty$ for fixed j . Also, using a basic theorem from calculus, $(1 - np_n/n)^{n-k} \rightarrow e^{-a}$ as $n \rightarrow \infty$.

Proof from generating functions

An easier proof uses probability generating functions. For $s \in \mathbb{R}$, using the same basic limit from calculus,

$$[(1 - p_n) + p_n s]^n \rightarrow e^{a(s-1)} \text{ as } n \rightarrow \infty \quad (14.4.29)$$

The left side is the PGF of the binomial distribution with parameters n and p_n , while the right side is the PGF of the Poisson distribution with parameter a .

The mean and variance of the binomial distribution converge to the mean and variance of the limiting Poisson distribution, respectively.

1. $np_n \rightarrow a$ as $n \rightarrow \infty$
2. $np_n(1 - p_n) \rightarrow a$ as $n \rightarrow \infty$

Of course the convergence of the means is precisely our basic assumption, and is further evidence that this is the essential assumption. But for a deeper look, let's return to the analogy between the Bernoulli trials process and the Poisson process. Recall that both have the strong renewal property that at each fixed time, and at each arrival time, the process stochastically starts over, independently of the past. The difference, of course, is that time is discrete in the Bernoulli trials process and continuous in the Poisson process. The convergence result is a special case of the more general fact that if we run Bernoulli trials at a faster and faster rate but with a smaller and smaller success probability, in just the right way, the Bernoulli trials process converges to the Poisson process. Specifically, suppose that we have a sequence of Bernoulli trials processes. In process n we perform the trials at a rate of n per unit time, with success probability p_n . Our basic assumption is that $np_n \rightarrow r$ as $n \rightarrow \infty$ where $r > 0$. Now let $Y_{n,t}$ denote the number of successes in the time interval $(0, t]$ for Bernoulli trials process n , and let N_t denote the number of arrivals in this interval for the Poisson process with rate r . Then $Y_{n,t}$ has the binomial distribution with parameters $[nt]$ and p_n , and of course N_t has the Poisson distribution with parameter rt .

For $t > 0$, the distribution of $Y_{n,t}$ converges to the distribution of N_t as $n \rightarrow \infty$.

Proof

Note that $nt - 1 < \lfloor nt \rfloor \leq nt$ and hence $np_n t - p_n < \lfloor nt \rfloor p_n \leq np_n t$. Since $np_n \rightarrow r$ and $p_n \rightarrow 0$ as $n \rightarrow \infty$, it follows from the squeeze theorem for limits that $\lfloor nt \rfloor p_n \rightarrow rt$ as $n \rightarrow \infty$. Thus, the result follows from our previous [convergence theorem](#).

Compare the Poisson experiment and the binomial timeline experiment.

1. Open the Poisson experiment and set $r = 1$ and $t = 5$. Run the experiment a few times and note the general behavior of the random points in time. Note also the shape and location of the probability density function and the mean \pm standard deviation bar.
2. Now open the binomial timeline experiment and set $n = 100$ and $p = 0.05$. Run the experiment a few times and note the general behavior of the random points in time. Note also the shape and location of the probability density function and the mean \pm standard deviation bar.

From a practical point of view, the convergence of the binomial distribution to the Poisson means that if the number of trials n is “large” and the probability of success p “small”, so that np^2 is small, then the binomial distribution with parameters n and p is well approximated by the Poisson distribution with parameter $r = np$. This is often a useful result, because the Poisson distribution has fewer parameters than the binomial distribution (and often in real problems, the parameters may only be known approximately). Specifically, in the approximating Poisson distribution, we do not need to know the number of trials n and the probability of success p *individually*, but only in the *product* np . The condition that np^2 be small means that the variance of the binomial distribution, namely $np(1-p) = np - np^2$ is approximately $r = np$, the variance of the approximating Poisson distribution.

Recall that the binomial distribution can also be approximated by the normal distribution, by virtue of the central limit theorem. The normal approximation works well when np and $n(1-p)$ are large; the rule of thumb is that both should be at least 5. The Poisson approximation works well, as we have already noted, when n is large and np^2 small.

Computational Exercises

Suppose that requests to a web server follow the Poisson model with rate $r = 5$ per minute. Find the probability that there will be at least 8 requests in a 2 minute period.

Answer

0.7798

Defects in a certain type of wire follow the Poisson model with rate 1.5 per meter. Find the probability that there will be no more than 4 defects in a 2 meter piece of the wire.

Answer

0.8153

Suppose that customers arrive at a service station according to the Poisson model, at a rate of $r = 4$. Find the mean and standard deviation of the number of customers in an 8 hour period.

Answer

32, 5.657

In the Poisson experiment, set $r = 5$ and $t = 4$. Run the experiment 1000 times and compute the following:

1. $\mathbb{P}(15 \leq N_4 \leq 22)$
2. The relative frequency of the event $\{15 \leq N_4 \leq 22\}$.
3. The normal approximation to $\mathbb{P}(15 \leq N_4 \leq 22)$.

Answer

1. 0.6157
3. 0.6025

Suppose that requests to a web server follow the Poisson model with rate $r = 5$ per minute. Compute the normal approximation to the probability that there will be at least 280 requests in a 1 hour period.

Answer

0.8818

Suppose that requests to a web server follow the Poisson model, and that 1 request comes in a five minute period. Find the probability that the request came during the first 3 minutes of the period.

Answer

0.6

Suppose that requests to a web server follow the Poisson model, and that 10 requests come during a 5 minute period. Find the probability that at least 4 requests came during the first 3 minutes of the period.

Answer

0.9452

In the Poisson experiment, set $r = 3$ and $t = 5$. Run the experiment 100 times.

1. For each run, compute the estimate of r based on N_t .
2. Over the 100 runs, compute the average of the squares of the errors.
3. Compare the result in (b) with $\text{var}(N_t)$.

Suppose that requests to a web server follow the Poisson model with unknown rate r per minute. In a one hour period, the server receives 342 requests. Estimate r .

Answer

$r = 5.7$ per minute

In the binomial experiment, set $n = 30$ and $p = 0.1$, and run the simulation 1000 times. Compute and compare each of the following:

1. $\mathbb{P}(Y_{30} \leq 4)$
2. The relative frequency of the event $\{Y_{30} \leq 4\}$
3. The Poisson approximation to $\mathbb{P}(Y_{30} \leq 4)$

Answer

1. 0.8245
3. 0.8153

Suppose that we have 100 memory chips, each of which is defective with probability 0.05, independently of the others. Approximate the probability that there are at least 3 defectives in the batch.

Answer

0.7350

In the binomial timeline experiment, set $n = 40$ and $p = 0.1$ and run the simulation 1000 times. Compute and compare each of the following:

1. $\mathbb{P}(Y_{40} > 5)$
2. The relative frequency of the event $\{Y_{40} > 5\}$
3. The Poisson approximation to $\mathbb{P}(Y_{40} > 5)$
4. The normal approximation to $\mathbb{P}(Y_{40} > 5)$

Answer

1. 0.2063
3. 0.2149
4. 0.2146

In the binomial timeline experiment, set $n = 100$ and $p = 0.1$ and run the simulation 1000 times. Compute and compare each of the following:

1. $\mathbb{P}(8 < Y_{100} < 15)$
2. The relative frequency of the event $\{8 < Y_{100} < 15\}$
3. The Poisson approximation to $\mathbb{P}(8 < Y_{100} < 15)$
4. The normal approximation to $\mathbb{P}(8 < Y_{100} < 15)$

Answer

1. 0.6066
3. 0.5837
4. 0.6247

A text file contains 1000 words. Assume that each word, independently of the others, is misspelled with probability p .

1. If $p = 0.015$, approximate the probability that the file contains at least 20 misspelled words.
2. If $p = 0.001$, approximate the probability that the file contains at least 3 misspelled words.

Answer

The true distribution of the number of misspelled words is binomial, with $n = 1000$ and p .

1. The normal approximation (with $\mu = np = 15$ and $\sigma^2 = np(1 - p) = 14.775$) is 0.120858. The Poisson approximation (with parameter $np = 15$) is 0.124781. The true binomial probability is 0.123095.
2. The Poisson approximation (with parameter $np = 1$) is 0.0803014. The true binomial probability is 0.0802093.

This page titled [14.4: The Poisson Distribution](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.5: Thinning and Superposition

Thinning

Thinning or *splitting* a Poisson process refers to classifying each random point, independently, into one of a finite number of different types. The random points of a given type also form Poisson processes, and these processes are independent. Our exposition will concentrate on the case of just two types, but this case has all of the essential ideas.

The Two-Type Process

We start with a Poisson process with rate $r > 0$. Recall that this statement really means three interrelated stochastic processes: the sequence of inter-arrival times $\mathbf{X} = (X_1, X_2, \dots)$, the sequence of arrival times $\mathbf{T} = (T_0, T_1, \dots)$, and the counting process $\mathbf{N} = (N_t : t \geq 0)$. Suppose now that each arrival, independently of the others, is one of two types: type 1 with probability p and type 0 with probability $1 - p$, where $p \in (0, 1)$ is a parameter. Here are some common examples:

- The arrivals are radioactive emissions and each emitted particle is either *detected* (type 1) or *missed* (type 0) by a counter.
- The arrivals are customers at a service station and each customer is classified as either *male* (type 1) or *female* (type 0).

We want to consider the type 1 and type 0 random points separately. For this reason, the new random process is usually referred to as *thinning* or *splitting* the original Poisson process. In some applications, the type 1 points are *accepted* while the type 0 points are *rejected*. The main result of this section is that the type 1 and type 0 points form separate Poisson processes, with rates rp and $r(1 - p)$ respectively, and are independent. We will explore this important result from several points of view.

Bernoulli Trials

In the previous sections, we have explored the analogy between the Bernoulli trials process and the Poisson process. Both have the strong renewal property that at each fixed time and at each arrival time, the process stochastically “restarts”, independently of the past. The difference, of course, is that time is discrete in the Bernoulli trials process and continuous in the Poisson process. In this section, we have both processes simultaneously, and given our previous explorations, it's perhaps not surprising that this leads to some interesting mathematics.

Thus, in addition to the processes \mathbf{X} , \mathbf{T} , and \mathbf{N} , we have a sequence of Bernoulli trials $\mathbf{I} = (I_1, I_2, \dots)$ with success parameter p . Indicator variable I_j specifies the type of the j th arrival. Moreover, because of our assumptions, \mathbf{I} is independent of \mathbf{X} , \mathbf{T} , and \mathbf{N} . Recall that V_k , the trial number of the k th success has the negative binomial distribution with parameters k and p for $k \in \mathbb{N}_+$. We take $V_0 = 0$ by convention. Also, U_k , the number of trials needed to go from the $(k - 1)$ st success to the k th success has the geometric distribution with success parameter p for $k \in \mathbb{N}_+$. Moreover, $\mathbf{U} = (U_1, U_2, \dots)$ is independent and $\mathbf{V} = (V_0, V_1, \dots)$ is the partial sum process associated with \mathbf{U} :

$$V_k = \sum_{i=1}^k U_i, \quad k \in \mathbb{N} \quad (14.5.1)$$

$$U_k = V_k - V_{k-1}, \quad k \in \mathbb{N}_+ \quad (14.5.2)$$

As noted above, the Bernoulli trials process can be thought of as random points in *discrete* time, namely the trial numbers of the successes. With this understanding, \mathbf{U} is the sequence of inter-arrival times and \mathbf{V} is the sequence of arrival times.

The Inter-arrival Times

Now consider just the type 1 points in our Poisson process. The time between the arrivals of $(k - 1)$ st and k th type 1 point is

$$Y_k = \sum_{i=V_{k-1}+1}^{V_k} X_i, \quad k \in \mathbb{N}_+ \quad (14.5.3)$$

Note that Y_k has U_k terms. The next result shows that the type 1 points form a Poisson process with rate pr .

$\mathbf{Y} = (Y_1, Y_2, \dots)$ is a sequence of independent variables and each has the exponential distribution with rate parameter pr

Proof

From the renewal properties of the Poisson process and the Bernoulli trials process, the inter-arrival times are independent and identically distributed. Each inter-arrival time is the sum of a random number of independent terms; each term has the exponential distribution with rate r , and the number of terms has the geometric distribution on \mathbb{N}_+ with parameter p . Moreover, the number of terms is independent of the terms themselves. We showed in the section on the exponential distribution that a random sum of this form has the exponential distribution with parameter rp .

Similarly, if $\mathbf{Z} = (Z_1, Z_2, \dots)$ is the sequence of interarrival times for the type 0 points, then \mathbf{Z} is a sequence of independent variables, and each has the exponential distribution with rate $(1-p)r$. Moreover, \mathbf{Y} and \mathbf{Z} are independent.

Counting Processes

For $t \geq 0$, let M_t denote the number of type 1 arrivals in $(0, t]$ and W_t the number of type 0 arrivals in $(0, t]$. Thus, $\mathbf{M} = (M_t : t \geq 0)$ and $\mathbf{W} = (W_t : t \geq 0)$ are the counting processes for the type 1 arrivals and for the type 0 arrivals. The next result follows from the previous subsection, but a direct proof is interesting.

For $t \geq 0$, M_t has the Poisson distribution with parameter prt , W_t has the Poisson distribution with parameter $(1-p)r$, and M_t and W_t are independent.

Proof

The important observation is that the conditional distribution of M_t given $N_t = n$ is binomial with parameters n and p . Thus for $j \in \mathbb{N}$ and $k \in \mathbb{N}$,

$$\mathbb{P}(M_t = j, W_t = k) = \mathbb{P}(M_t = j, N_t = j+k) = \mathbb{P}(N_t = j+k) \mathbb{P}(M_t = j | N_t = j+k) \quad (14.5.4)$$

$$= e^{-rt} \frac{(rt)^{j+k}}{(j+k)!} \frac{(j+k)!}{j!k!} p^j (1-p)^k \quad (14.5.5)$$

$$= e^{-prt} \frac{(prt)^j}{j!} e^{-(1-p)rt} \frac{[(1-p)rt]^k}{k!} \quad (14.5.6)$$

In the two-type Poisson experiment vary r , p , and t with the scroll bars and note the shape of the probability density functions. For various values of the parameters, run the experiment 1000 times and compare the relative frequency functions to the probability density functions.

Estimating the Number of Arrivals

Suppose that the type 1 arrivals are observable, but not the type 0 arrivals. This setting is natural, for example, if the arrivals are radioactive emissions, and the type 1 arrivals are emissions that are detected by a counter, while the type 0 arrivals are emissions that are missed. Suppose that for a given $t > 0$, we would like to estimate the total number arrivals N_t after observing the number of type 1 arrivals M_t .

The conditional distribution of N_t given $M_t = k$ is the same as the distribution of $k + W_t$.

$$\mathbb{P}(N_t = n | M_t = k) = e^{-(1-p)rt} \frac{[(1-p)rt]^{n-k}}{(n-k)!}, \quad n \in \{k, k+1, \dots\} \quad (14.5.7)$$

Proof

Recall from the basic [splitting result above](#) that M_t and W_t are independent. Thus, for $n \in \mathbb{N}$,

$$\mathbb{P}(N_t = n | M_t = k) = \frac{\mathbb{P}(N_t = n, M_t = k)}{\mathbb{P}(M_t = k)} = \frac{\mathbb{P}(M_t = k, W_t = n-k)}{\mathbb{P}(M_t = k)} \quad (14.5.8)$$

$$= \frac{\mathbb{P}(M_t = k) \mathbb{P}(W_t = n-k)}{\mathbb{P}(M_t = k)} = \mathbb{P}(W_t = n-k) \quad (14.5.9)$$

The form of the probability density function follows since W_t as the Poisson distribution with parameter $(1-p)r$.

$$\mathbb{E}(N_t | M_t = k) = k + (1-p)r \quad .$$

Proof

This follows easily from our previous theorem since $\mathbb{E}(N_t | M_t = k) = \mathbb{E}(k + W_t) = k + (1 - p)r$.

Thus, if the overall rate r of the process and the probability p that an arrival is type 1 are known, then it follows from the general theory of conditional expectation that the best estimator of N_t based on M_t , in the least squares sense, is

$$\mathbb{E}(N_t | M_t) = M_t + (1 - p)r \quad (14.5.10)$$

The mean square error is $\mathbb{E}([N_t - \mathbb{E}(N_t | M_t)]^2) = (1 - p)rt$.

Proof

Note that $N_t - \mathbb{E}(N_t | M_t) = W_t - (1 - p)r$. Thus the mean square error is just $\text{var}(W_t) = (1 - p)rt$.

The Multi-Type Process

As you might guess, the results above generalize from 2 types to k types for general $k \in \mathbb{N}_+$. Once again, we start with a Poisson process with rate $r > 0$. Suppose that each arrival, independently of the others, is type i with probability p_i for $i \in \{0, 1, \dots, k-1\}$. Of course we must have $p_i \geq 0$ for each i and $\sum_{i=0}^{k-1} p_i = 1$. Then for each i , the type i points form a Poisson process with rate $p_i r$, and these processes are independent.

Superposition

Complementary to splitting or thinning a Poisson process is *superposition*: if we combine the random points in time from independent Poisson processes, then we have a new Poisson processes. The rate of the new process is the sum of the rates of the processes that were combined. Once again, our exposition will concentrate on the superposition of two processes. This case contains all of the essential ideas.

Two Processes

Suppose that we have two independent Poisson processes. We will denote the sequence of inter-arrival times, the sequence of arrival times, and the counting variables for the process $i \in \{1, 2\}$ by $\mathbf{X}^i = (X_1^i, X_2^i, \dots)$, $\mathbf{T}^i = (T_1^i, T_2^i, \dots)$, and $\mathbf{N}^i = (N_t^i : t \in [0, \infty))$, and we assume that process i has rate $r_i \in (0, \infty)$. The new process that we want to consider is obtained by simply combining the random points. That is, the new random points are $\{T_n^1 : n \in \mathbb{N}_+\} \cup \{T_n^2 : n \in \mathbb{N}_+\}$, but of course then ordered in time. We will denote the sequence of inter-arrival times, the sequence of arrival times, and the counting variables for the new process by $\mathbf{X} = (X_1, X_2, \dots)$, $\mathbf{T} = (T_1, T_2, \dots)$, and $\mathbf{N} = (N_t : t \in [0, \infty))$. Clearly if A is an interval in $[0, \infty)$ then

$$N(A) = N^1(A) + N^2(A) \quad (14.5.11)$$

the number of combined points in A is simply the sum of the number of point in A for processes 1 and 2. It's also worth noting that

$$X_1 = \min\{X_1^1, X_1^2\} \quad (14.5.12)$$

the first arrival for the combined process is the smaller of the first arrival times for processes 1 and 2. The other inter-arrival times, and hence also the arrival times, for the combined process are harder to state.

The combined process is a Poisson process with rate $r_1 + r_2$. Moreover,

Proof

As noted above, if A is a subinterval of $[0, \infty)$ then $N(A) = N^1(A) + N^2(A)$. The first term has the Poisson distribution with parameter $r_1 \lambda(A)$, the second term has the Poisson distribution with parameter $r_2 \lambda(A)$, and the terms are independent. Hence $N(A)$ has the Poisson distribution with parameter $r_1 \lambda(A) + r_2 \lambda(A) = (r_1 + r_2) \lambda(A)$. Thus the counting process has stationary, Poisson distributed increments. Next, if (A_1, A_2, \dots, A_n) is a sequence of disjoint subintervals of $[0, \infty)$ then

$$(N(A_1), N(A_2), \dots, N(A_n)) = (N^1(A_1) + N^2(A_1), N^1(A_2) + N^2(A_2), \dots, N^1(A_n) + N^2(A_n)) \quad (14.5.13)$$

is an independent sequence, so the counting process has independent increments.

Computational Exercises

In the two-type Poisson experiment, set $r = 2$, $t = 3$, and $p = 0.7$. Run the experiment 1000 times, Compute the appropriate relative frequency functions and investigate empirically the independence of the number of type 1 points and the number of type 0 points.

Suppose that customers arrive at a service station according to the Poisson model, with rate $r = 20$ per hour. Moreover, each customer, independently, is female with probability 0.6 and male with probability 0.4. Find the probability that in a 2 hour period, there will be at least 20 women and at least 15 men.

Answer

0.5814

In the two-type Poisson experiment, set $r = 3$, $t = 4$, and $p = 0.8$. Run the experiment 100 times.

1. Compute the estimate of N_t based on M_t for each run.
2. Over the 100 runs, compute average of the sum of the squares of the errors.
3. Compare the result in (b) with the result in Exercise 8.

Suppose that a piece of radioactive material emits particles according to the Poisson model at a rate of $r = 100$ per second. Moreover, assume that a counter detects each emitted particle, independently, with probability 0.9. Suppose that the number of detected particles in a 5 second period is 465.

1. Estimate the number of particles emitted.
2. Compute the mean square error of the estimate.

Answer

1. 515
2. 50

This page titled [14.5: Thinning and Superposition](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.6: Non-homogeneous Poisson Processes

Basic Theory

A *non-homogeneous Poisson process* is similar to an ordinary Poisson process, except that the average rate of arrivals is allowed to vary with time. Many applications that generate random points in time are modeled more faithfully with such non-homogeneous processes. The mathematical cost of this generalization, however, is that we lose the property of stationary increments.

Non-homogeneous Poisson processes are best described in measure-theoretic terms. Thus, you may need to review the sections on measure theory in the chapters on Foundations, Probability Measures, and Distributions. Our basic measure space in this section is $[0, \infty)$ with the σ -algebra of Borel measurable subsets (named for Émile Borel). As usual, λ denotes Lebesgue measure on this space, named for Henri Lebesgue. Recall that the Borel σ -algebra is the one generated by the intervals, and λ is the generalization of length on intervals.

Definition and Basic Properties

Of all of our various characterizations of the ordinary Poisson process, in terms of the inter-arrival times, the arrival times, and the counting process, the characterizations involving the counting process leads to the most natural generalization to non-homogeneous processes. Thus, consider a process that generates random points in time, and as usual, let N_t denote the number of random points in the interval $(0, t]$ for $t \geq 0$, so that $\mathbf{N} = \{N_t : t \geq 0\}$ is the counting process. More generally, $N(A)$ denotes the number of random points in a measurable $A \subseteq [0, \infty)$, so N is our random counting measure. As before, $t \mapsto N_t$ is a (random) distribution function and $A \mapsto N(A)$ is the (random) measure associated with this distribution function.

Suppose now that $r : [0, \infty) \rightarrow [0, \infty)$ is measurable, and define $m : [0, \infty) \rightarrow [0, \infty)$ by

$$m(t) = \int_{(0,t]} r(s) d\lambda(s) \quad (14.6.1)$$

From properties of the integral, m is increasing and right-continuous on $[0, \infty)$ and hence is distribution function. The positive measure on $[0, \infty)$ associated with m (which we will also denote by m) is defined on a measurable $A \subseteq [0, \infty)$ by

$$m(A) = \int_A r(s) d\lambda(s) \quad (14.6.2)$$

Thus, $m(t) = m(0, t]$, and for $s, t \in [0, \infty)$ with $s < t$, $m(s, t] = m(t) - m(s)$. Finally, note that the measure m is absolutely continuous with respect to λ , and r is the density function. Note the parallels between the *random* distribution function and measure N and the *deterministic* distribution function and measure m . With the setup involving r and m complete, we are ready for our first definition.

A process that produces random points in time is a *non-homogeneous Poisson process* with rate function r if the counting process N satisfies the following properties:

1. If $\{A_i : i \in I\}$ is a countable, disjoint collection of measurable subsets of $[0, \infty)$ then $\{N(A_i) : i \in I\}$ is a collection of independent random variables.
2. If $A \subseteq [0, \infty)$ is measurable then $N(A)$ has the Poisson distribution with parameter $m(A)$.

Property (a) is our usual property of independent increments, while property (b) is a natural generalization of the property of Poisson distributed increments. Clearly, if r is a positive constant, then $m(t) = rt$ for $t \in [0, \infty)$ and as a measure, m is proportional to Lebesgue measure λ . In this case, the non-homogeneous process reduces to an ordinary, homogeneous Poisson process with rate r . However, if r is not constant, then m is not linear, and as a measure, is not proportional to Lebesgue measure. In this case, the process does not have stationary increments with respect to λ , but does of course, have stationary increments with respect to m . That is, if A, B are measurable subsets of $[0, \infty)$ and $\lambda(A) = \lambda(B)$ then $N(A)$ and $N(B)$ will not in general have the same distribution, but of course they will have the same distribution if $m(A) = m(B)$.

In particular, recall that the parameter of the Poisson distribution is both the mean and the variance, so $\mathbb{E}[N(A)] = \text{var}[N(A)] = m(A)$ for measurable $A \subseteq [0, \infty)$, and in particular, $\mathbb{E}(N_t) = \text{var}(N_t) = m(t)$ for $t \in [0, \infty)$. The

function m is usually called the *mean function*. Since $m'(t) = r(t)$ (if r is continuous at t), it makes sense to refer to r as the rate function. Locally, at t , the arrivals are occurring at an average rate of $r(t)$ per unit time.

As before, from a modeling point of view, the property of independent increments can reasonably be evaluated. But we need something more primitive to replace the property of Poisson increments. Here is the main theorem.

A process that produces random points in time is a non-homogeneous Poisson process with rate function r if and only if the counting process \mathbf{N} satisfies the following properties:

1. If $\{A_i : i \in I\}$ is a countable, disjoint collection of measurable subsets of $[0, \infty)$ then $\{N(A_i) : i \in I\}$ is a set of independent variables.
2. For $t \in [0, \infty)$,

$$\frac{\mathbb{P}[N(t, t+h) = 1]}{h} \rightarrow r(t) \text{ as } h \downarrow 0 \quad (14.6.3)$$

$$\frac{\mathbb{P}[N(t, t+h) > 1]}{h} \rightarrow 0 \text{ as } h \downarrow 0 \quad (14.6.4)$$

So if h is “small” the probability of a single arrival in $[t, t+h)$ is approximately $r(t)h$, while the probability of more than 1 arrival in this interval is negligible.

Arrival Times and Time Change

Suppose that we have a non-homogeneous Poisson process with rate function r , as defined above. As usual, let T_n denote the time of the n th arrival for $n \in \mathbb{N}$. As with the ordinary Poisson process, we have an inverse relation between the counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ and the arrival time sequence $\mathbf{T} = \{T_n : n \in \mathbb{N}\}$, namely $T_n = \min\{t \in [0, \infty) : N_t = n\}$, $N_t = \#\{n \in \mathbb{N} : T_n \leq t\}$, and $\{T_n \leq t\} = \{N_t \geq n\}$, since both events mean at least n random points in $(0, t]$. The last relationship allows us to get the distribution of T_n .

For $n \in \mathbb{N}_+$, T_n has probability density function f_n given by

$$f_n(t) = \frac{m^{n-1}(t)}{(n-1)!} r(t) e^{-m(t)}, \quad t \in [0, \infty) \quad (14.6.5)$$

Proof

Using the inverse relationship above and the Poisson distribution of N_t , the distribution function of T_n is

$$\mathbb{P}(T_n \leq t) = \mathbb{P}(N_t \geq n) = \sum_{k=n}^{\infty} e^{-m(t)} \frac{m^k(t)}{k!}, \quad t \in [0, \infty) \quad (14.6.6)$$

Differentiating with respect to t gives

$$f_n(t) = \sum_{k=n}^{\infty} \left[-m'(t) e^{-m(t)} \frac{m^k(t)}{k!} + e^{-m(t)} \frac{k m^{k-1}(t) m'(t)}{k!} \right] = r(t) e^{-m(t)} \sum_{k=n}^{\infty} \left[\frac{m^{k-1}(t)}{(k-1)!} - \frac{m^k(t)}{k!} \right] \quad (14.6.7)$$

The last sum collapses to $m^{n-1}(t)/(n-1)!$.

In particular, T_1 has probability density function f_1 given by

$$f_1(t) = r(t) e^{-m(t)}, \quad t \in [0, \infty) \quad (14.6.8)$$

Recall that in reliability terms, r is the *failure rate function*, and that the *reliability function* is the right distribution function:

$$F_1^c(t) = \mathbb{P}(T_1 > t) = e^{-m(t)}, \quad t \in [0, \infty) \quad (14.6.9)$$

In general, the functional form of f_n is clearly similar to the probability density function of the gamma distribution, and indeed, T_n can be transformed into a random variable with a gamma distribution. This amounts to a *time change* which will give us additional insight into the non-homogeneous Poisson process.

Let $U_n = m(T_n)$ for $n \in \mathbb{N}_+$. Then U_n has the gamma distribution with shape parameter n and rate parameter 1

Proof

Let g_n denote the PDF of U_n . Since m is strictly increasing and differentiable, we can use the standard change of variables formula. So letting $u = m(t)$, the relationship is

$$g_n(u) = f_n(t) \frac{dt}{du} \quad (14.6.10)$$

Simplifying gives $g_n(u) = u^{n-1} e^{-u} / (n-1)!$ for $u \in [0, \infty)$.

Thus, the time change $u = m(t)$ transforms the non-homogeneous Poisson process into a standard (rate 1) Poisson process. Here is an equivalent way to look at the time change result.

For $u \in [0, \infty)$, let $M_u = N_t$ where $t = m^{-1}(u)$. Then $\{M_u : u \in [0, \infty)\}$ is the counting process for a standard, rate 1 Poisson process.

Proof

1. Suppose that (u_1, u_2, \dots) as a sequence of points in $[0, \infty)$ with $0 \leq u_1 < u_2 < \dots$. Since m^{-1} is strictly increasing, we have $0 \leq t_1 < t_2 < \dots$, where of course $t_i = m^{-1}(u_i)$. By assumption, the sequence of random variables $(N_{t_1}, N_{t_2} - N_{t_1}, \dots)$ is independent, but this is also the sequence $(M_{u_1}, M_{u_2} - M_{u_1}, \dots)$.
2. Suppose that $u, v \in [0, \infty)$ with $u < v$, and let $s = m^{-1}(u)$ and $t = m^{-1}(v)$. Then $s < t$ and so $M_v - M_u = N_t - N_s$ has the Poisson distribution with parameter $m(t) - m(s) = v - u$.

Equivalently, we can transform a standard (rate 1) Poisson process into a non-homogeneous Poisson process with a time change.

Suppose that $M = \{M_u : u \in [0, \infty)\}$ is the counting process for a standard Poisson process, and let $N_t = M_{m(t)}$ for $t \in [0, \infty)$. Then $\{N_t : t \in [0, \infty)\}$ is the counting process for a non-homogeneous Poisson process with mean function m (and rate function r).

Proof

1. Let (t_1, t_2, \dots) be a sequence of points in $[0, \infty)$ with $0 \leq t_1 < t_2 < \dots$. Since m is strictly increasing, we have $0 \leq m(t_1) < m(t_2) < \dots$. Hence $(M_{m(t_1)}, M_{m(t_2)} - M_{m(t_1)}, \dots)$ is a sequence of independent variables. But this sequence is simply $(N_{t_1}, N_{t_2} - N_{t_1}, \dots)$.
2. If $s, t \in [0, \infty)$ with $s < t$. Then $N_t - N_s = M_{m(t)} - M_{m(s)}$ has the Poisson distribution with parameter $m(t) - m(s)$.

This page titled [14.6: Non-homogeneous Poisson Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.7: Compound Poisson Processes

In a *compound Poisson process*, each arrival in an ordinary Poisson process comes with an associated real-valued random variable that represents the *value* of the arrival in a sense. These variables are independent and identically distributed, and are independent of the underlying Poisson process. Our interest centers on the sum of the random variables for all the arrivals up to a fixed time t , which thus is a Poisson-distributed random sum of random variables. Distributions of this type are said to be *compound Poisson distributions*, and are important in their own right, particularly since some surprising parametric distributions turn out to be compound Poisson.

Basic Theory

Definition

Suppose we have a Poisson process with rate $r \in (0, \infty)$. As usual, we will denote the sequence of inter-arrival times by $\mathbf{X} = (X_1, X_2, \dots)$, the sequence of arrival times by $\mathbf{T} = (T_0, T_1, T_2, \dots)$, and the counting process by $\mathbf{N} = \{N_t : t \in [0, \infty)\}$. To review some of the most important facts briefly, recall that \mathbf{X} is a sequence of independent random variables, each having the exponential distribution on $[0, \infty)$ with rate r . The sequence \mathbf{T} is the partial sum sequence associated with \mathbf{X} , and has stationary independent increments. For $n \in \mathbb{N}_+$, the n th arrival time T_n has the gamma distribution with parameters n and r . The process \mathbf{N} is the inverse of \mathbf{T} , in a certain sense, and also has stationary independent increments. For $t \in (0, \infty)$, the number of arrivals N_t in $(0, t]$ has the Poisson distribution with parameter rt .

Suppose now that each arrival has an associated real-valued random variable that represents the *value* of the arrival in a certain sense. Here are some typical examples:

- The arrivals are customers at a store. Each customer spends a random amount of money.
- The arrivals are visits to a website. Each visitor spends a random amount of time at the site.
- The arrivals are failure times of a complex system. Each failure requires a random repair time.
- The arrivals are earthquakes at a particular location. Each earthquake has a random severity, a measure of the energy released.

For $n \in \mathbb{N}_+$, let U_n denote the value of the n th arrival. We assume that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent, identically distributed, real-valued random variables, and that \mathbf{U} is independent of the underlying Poisson process. The common distribution may be discrete or continuous, but in either case, we let f denote the common probability density function. We will let $\mu = \mathbb{E}(U_n)$ denote the common mean, $\sigma^2 = \text{var}(U_n)$ the common variance, and G the common moment generating function, so that $G(s) = \mathbb{E}[\exp(sU_n)]$ for s in some interval I about 0. Here is our main definition:

The *compound Poisson process* associated with the given Poisson process \mathbf{N} and the sequence \mathbf{U} is the stochastic process $\mathbf{V} = \{V_t : t \in [0, \infty)\}$ where

$$V_t = \sum_{n=1}^{N_t} U_n \quad (14.7.1)$$

Thus, V_t is the total value for all of the arrivals in $(0, t]$. For the examples above

- V_t is the total income to the store up to time t .
- V_t is the total time spent at the site by the customers who arrived up to time t .
- V_t is the total repair time for the failures up to time t .
- V_t is the total energy released up to time t .

Recall that a sum over an empty index set is 0, so $V_0 = 0$.

Properties

Note that for fixed t , V_t is a random sum of independent, identically distributed random variables, a topic that we have studied before. In this sense, we have a special case, since the number of terms N_t has the Poisson distribution with parameter rt . But we also have a new wrinkle, since the process is indexed by the continuous time parameter t , and so we can study its properties as a stochastic process. Our first result is a pair of properties shared by the underlying Poisson process.

\mathbf{V} has stationary, independent increments:

1. If $s, t \in [0, \infty)$ with $s < t$, then $V_t - V_s$ has the same distribution as V_{t-s} .
2. If (t_1, t_2, \dots, t_n) is a sequence of points in $[0, \infty)$ with $t_1 < t_2 < \dots < t_n$ then $(V_{t_1}, V_{t_2} - V_{t_1}, \dots, V_{t_n} - V_{t_{n-1}})$ is a sequence of independent variables.

Proof

1. For $0 \leq s < t$,

$$V_t - V_s = \sum_{i=1}^{N_t} U_i - \sum_{i=1}^{N_s} U_i = \sum_{i=N_s+1}^{N_t} U_i \quad (14.7.2)$$

The number of terms in the last sum is $N_t - N_s$, which has the same distribution as N_{t-s} . Since the variables in the sequence \mathbf{U} are identically distributed, it follows that $V_t - V_s$ has the same distribution as V_{t-s} .

2. Suppose that $0 \leq t_1 < t_2 < \dots < t_n$ and let $t_0 = 0$. Then for $i \in \{1, 2, \dots, n\}$ as in (a)

$$V_{t_i} - V_{t_{i-1}} = \sum_{j=N_{t_{i-1}}+1}^{N_{t_i}} U_j \quad (14.7.3)$$

The number of terms in this sum is $N_{t_i} - N_{t_{i-1}}$. Since \mathbf{N} has independent increments, and the variables in \mathbf{U} are independent, and since the indices between $N_{t_{i-1}+1}$ and N_{t_i} are disjoint over $i \in \{1, 2, \dots, n\}$, it follows that the random variables $V_{t_i} - V_{t_{i-1}}$ are independent over $i \in \{1, 2, \dots, n\}$.

Next we consider various moments of the compound process.

For $t \in [0, \infty)$, the mean and variance of V_t are

1. $\mathbb{E}(V_t) = \mu rt$
2. $\text{var}(V_t) = (\mu^2 + \sigma^2)rt$

Proof

Again, these are special cases of general results for random sums of IID variables, but we give separate proofs for completeness. The basic tool is conditional expected value and conditional variance. Recall also that $\mathbb{E}(N_t) = \text{var}(N_t) = rt$.

1. Note that $\mathbb{E}(V_t) = \mathbb{E}[\mathbb{E}(V_t | N_t)] = \mathbb{E}(\mu N_t) = \mu rt$.
2. Similarly, note that $\text{var}(V_t | N_t) = \sigma^2 N_t$ and hence
 $\text{var}(V_t) = \mathbb{E}[\text{var}(V_t | N_t)] + \text{var}[\mathbb{E}(V_t | N_t)] = \mathbb{E}(\sigma^2 N_t) + \text{var}(\mu N_t) = \sigma^2 rt + \mu^2 rt$.

For $t \in [0, \infty)$, the moment generating function of V_t is given by

$$\mathbb{E}[\exp(sV_t)] = \exp(rt[G(s) - 1]), \quad s \in I \quad (14.7.4)$$

Proof

Again, this is a special case of the more general result for random sums of IID variables, but we give a another proof for completeness. As with the last theorem, the key is to condition on N_t and recall that the MGF of a sum of independent variables is the product of the MGFs. Thus

$$\mathbb{E}[\exp(sV_t)] = \mathbb{E}[\mathbb{E}[\exp(sV_t | N_t)]] = \mathbb{E}[G^{N_t}(s)] = P_t[G(s)] \quad (14.7.5)$$

where P_t is the probability generating function of N_t . But we know from our study of the Poisson distribution that $P_t(x) = \exp[rt(x - 1)]$ for $x \in \mathbb{R}$.

By exactly the same argument, the same relationship holds for characteristic functions and, in the case that the variables in \mathbf{U} take values in \mathbb{N} , for probability generating functions.. That is, if the variables in \mathbf{U} have generating function G , then the generating function H of V_t is given by

$$H(s) = \exp(rt[G(s) - 1]) \quad (14.7.6)$$

for s in the domain of G , where *generating function* can be any of the three types we have discussed: probability, moment, or characteristic.

Examples and Special Cases

The Discrete Case

First we note that Thinning a Poisson process can be thought of as a special case of a compound Poisson process. Thus, suppose that $\mathbf{U} = (U_1, U_2, \dots)$ is a Bernoulli trials sequence with success parameter $p \in (0, 1)$, and as above, that \mathbf{U} is independent of the Poisson process \mathbf{N} . In the usual language of thinning, the arrivals are of two types (1 and 0), and U_i is the type of the i th arrival. Thus the compound process \mathbf{V} constructed above is the thinned process, so that V_t is the number of type 1 points up to time t . We know that \mathbf{V} is also a Poisson process, with rate rp .

The results above for thinning generalize to the case where the values of the arrivals have a discrete distribution. Thus, suppose U_i takes values in a countable set $S \subseteq \mathbb{R}$, and as before, let f denote the common probability density function so that $f(u) = \mathbb{P}(U_i = u)$ for $u \in S$ and $i \in \mathbb{N}_+$. For $u \in S$, let N_t^u denote the number of arrivals up to time t that have the value u , and let $\mathbf{N}^u = \{N_t^u : t \in [0, \infty)\}$ denote the corresponding stochastic process. Armed with this setup, here is the result:

The compound Poisson process \mathbf{V} associated with \mathbf{N} and \mathbf{U} can be written in the form

$$V_t = \sum_{u \in S} u N_t^u, \quad t \in [0, \infty) \quad (14.7.7)$$

The processes $\{\mathbf{N}^u : u \in S\}$ are independent Poisson processes, and \mathbf{N}^u has rate $rf(u)$ for $u \in S$.

Proof

Note that $U_i = \sum_{u \in S} u \mathbf{1}(U_i = u)$ and hence

$$V_t = \sum_{i=1}^{N_t} U_i = \sum_{i=1}^{N_t} \sum_{u \in S} u \mathbf{1}(U_i = u) = \sum_{u \in S} u \sum_{i=1}^{N_t} \mathbf{1}(U_i = u) = \sum_{u \in S} u N_t^u \quad (14.7.8)$$

The fact that $\{\mathbf{N}^u : u \in S\}$ are independent Poisson processes, and that \mathbf{N}^u has rate $rf(u)$ for $u \in S$ follows from our result on thinning.

Compound Poisson Distributions

A compound Poisson random variable can be defined outside of the context of a Poisson process. Here is the formal definition:

Suppose that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent, identically distributed random variables, and that N is independent of \mathbf{U} and has the Poisson distribution with parameter $\lambda \in (0, \infty)$. Then $V = \sum_{i=1}^N U_i$ has a *compound Poisson distribution*.

But in fact, compound Poisson variables usually *do* arise in the context of an underlying Poisson process. In any event, the results on the mean and variance [above](#) and the generating function [above](#) hold with rt replaced by λ . Compound Poisson distributions are infinitely divisible. A famous theorem of William Feller gives a partial converse: an infinitely divisible distribution on \mathbb{N} must be compound Poisson.

The negative binomial distribution on \mathbb{N} is infinitely divisible, and hence must be compound Poisson. Here is the construction:

Let $p, k \in (0, \infty)$. Suppose that $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent variables, each having the logarithmic series distribution with shape parameter $1 - p$. Suppose also that N is independent of \mathbf{U} and has the Poisson distribution with parameter $-k \ln(p)$. Then $V = \sum_{i=1}^N U_i$ has the negative binomial distribution on \mathbb{N} with parameters k and p .

Proof

As noted [above](#), the probability generating function of V is $P(t) = \exp(\lambda[Q(t) - 1])$ where λ is the parameter of the Poisson variable N and $Q(t)$ is the common PGF of the terms in the sum. Using the PGF of the logarithmic series distribution, and the particular values of the parameters, we have

$$P(t) = \exp \left[-k \ln(p) \left(\frac{\ln[1 - (1 - p)t]}{\ln(p)} - 1 \right) \right], \quad |t| < \frac{1}{1 - p} \quad (14.7.9)$$

Using properties of logarithms and simple algebra, this reduces to

$$P(t) = \left(\frac{p}{1 - (1-p)t} \right)^k, \quad |t| < \frac{1}{1-p} \quad (14.7.10)$$

which is the PGF of the negative binomial distribution with parameters k and p .

As a special case ($k = 1$), it follows that the geometric distribution on \mathbb{N} is also compound Poisson.

This page titled [14.7: Compound Poisson Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

14.8: Poisson Processes on General Spaces

Basic Theory

The Process

So far, we have studied the Poisson process as a model for random points in *time*. However there is also a Poisson model for random points in *space*. Some specific examples of such “random points” are

- Defects in a sheet of material.
- Raisins in a cake.
- Stars in the sky.

The Poisson process for random points in space can be defined in a very general setting. All that is really needed is a measure space (S, \mathcal{S}, μ) . Thus, S is a set (the underlying space for our random points), \mathcal{S} is a σ -algebra of subsets of S (as always, the allowable sets), and μ is a positive measure on (S, \mathcal{S}) (a measure of the size of sets). The most important special case is when S is a (Lebesgue) measurable subset of \mathbb{R}^d for some $d \in \mathbb{N}_+$, \mathcal{S} is the σ -algebra of measurable subsets of S , and $\mu = \lambda_d$ is d -dimensional Lebesgue measure. Specializing further, recall the lower dimensional spaces:

1. When $d = 1$, $S \subseteq \mathbb{R}$ and λ_1 is length measure.
2. When $d = 2$, $S \subseteq \mathbb{R}^2$ and λ_2 is area measure.
3. When $d = 3$, $S \subseteq \mathbb{R}^3$ and λ_3 is volume measure.

Of course, the characterizations of the Poisson process on $[0, \infty)$, in terms of the inter-arrival times and the characterization in terms of the arrival times do not generalize because they depend critically on the *order relation* on $[0, \infty)$. However the characterization in terms of the counting process generalizes perfectly to our new setting. Thus, consider a process that produces random points in S , and as usual, let $N(A)$ denote the number of random points in $A \in \mathcal{S}$. Thus N is a random, counting measure on (S, \mathcal{S})

The random measure N is a *Poisson process* or a *Poisson random measure* on S with density parameter $r > 0$ if the following axioms are satisfied:

1. If $A \in \mathcal{S}$ then $N(A)$ has the Poisson distribution with parameter $r\mu(A)$.
2. If $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} then $\{N(A_i) : i \in I\}$ is a set of independent random variables.

To draw parallels with the Poisson process on $[0, \infty)$, note that axiom (a) is the generalization of stationary, Poisson-distributed increments, and axiom (b) is the generalization of independent increments. By convention, if $\mu(A) = 0$ then $N(A) = 0$ with probability 1, and if $\mu(A) = \infty$ then $N(A) = \infty$ with probability 1. (These distributions are considered degenerate members of the Poisson family.) On the other hand, note that if $0 < \mu(A) < \infty$ then $N(A)$ has support \mathbb{N} .

In the two-dimensional Poisson process, vary the width w and the rate r . Note the location and shape of the probability density function of N . For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the true probability density function.

For $A \subseteq D$

1. $\mathbb{E}[N(A)] = r\mu(A)$
2. $\text{var}[N(A)] = r\mu(A)$

Proof

These results follow of course from our previous study of the Poisson distribution. Recall that the parameter of the Poisson distribution is both the mean and the variance.

In particular, r can be interpreted as the expected *density* of the random points (that is, the expected number of points in a region of unit size), justifying the name of the parameter.

In the two-dimensional Poisson process, vary the width w and the density parameter r . Note the size and location of the mean \pm standard deviation bar of N . For various values of the parameters, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

The Distribution of the Random Points

As before, the Poisson model defines the most random way to distribute points in space, in a certain sense. Assume that we have a Poisson process N on (S, \mathcal{S}, μ) with density parameter $r \in (0, \infty)$.

Given that $A \in \mathcal{S}$ contains exactly one random point, the position X of the point is uniformly distributed on A .

Proof

For $B \in \mathcal{S}$ with $B \subseteq A$,

$$\mathbb{P}[N(B) = 1 \mid N(A) = 1] = \frac{\mathbb{P}[N(B) = 1, N(A) = 1]}{\mathbb{P}[N(A) = 1]} = \frac{\mathbb{P}[N(B) = 1, N(A \setminus B) = 0]}{\mathbb{P}[N(A) = 1]} = \frac{\mathbb{P}[N(B) = 1] \mathbb{P}[N(A \setminus B) = 0]}{\mathbb{P}[N(A) = 1]} \quad (14.8.1)$$

Using the Poisson distributions we have

$$\mathbb{P}[N(B) = 1 \mid N(A) = 1] = \frac{\exp[-r\mu(B)] [r\mu(B)] \exp[-r\mu(A \setminus B)]}{\exp[-r\mu(A)] [r\mu(A)]} = \frac{\mu(B)}{\mu(A)} \quad (14.8.2)$$

As a function of B , this is the uniform distribution on A (with respect to μ).

More generally, if A contains n points, then the positions of the points are independent and each is uniformly distributed in A .

Suppose that $A, B \in \mathcal{S}$ and $B \subseteq A$. For $n \in \mathbb{N}_+$, the conditional distribution of $N(B)$ given $N(A) = n$ is the binomial distribution with trial parameter n and success parameter $p = \mu(B)/\mu(A)$.

Proof

For $k \in \{0, 1, \dots, n\}$,

$$\mathbb{P}[N(B) = k | N(A) = n] = \frac{\mathbb{P}[N(B) = k, N(A) = n]}{\mathbb{P}[N(A) = n]} = \frac{\mathbb{P}[N(B) = k, N(A \setminus B) = n - k]}{\mathbb{P}[N(A) = n]} = \frac{\mathbb{P}[N(B) = k] \mathbb{P}[N(A \setminus B) = n - k]}{\mathbb{P}[N(A) = n]} \quad (14.8.3)$$

Using the Poisson distributions,

$$\mathbb{P}[N(B) = k | N(A) = n] = \frac{\exp[-r\mu(B)] \left([r\mu(B)]^k / k! \right) \exp[-r\mu(A \setminus B)] \left([r\mu(A \setminus B)]^{n-k} / (n-k)! \right)}{\exp[-r\mu(A)] [r\mu(A)]^n / n!} \quad (14.8.4)$$

Canceling factors and letting $p = \mu(B)/\mu(A)$, we have

$$\mathbb{P}[N(B) = k | N(A) = n] = \frac{n!}{k!(n-k)!} p^k (1-p)^{n-k} \quad (14.8.5)$$

Thus, given $N(A) = n$, each of the n random points falls into B , independently, with probability $p = \mu(B)/\mu(A)$, regardless of the density parameter r .

More generally, suppose that $A \in \mathcal{S}$ and that A is partitioned into k subsets (B_1, B_2, \dots, B_k) in \mathcal{S} . Then the conditional distribution of $(N(B_1), N(B_2), \dots, N(B_k))$ given $N(A) = n$ is the multinomial distribution with parameters n and (p_1, p_2, \dots, p_k) , where $p_i = \mu(B_i)/\mu(A)$ for $i \in \{1, 2, \dots, k\}$.

Thinning and Combining

Suppose that N is a Poisson random process on (S, \mathcal{S}, μ) with density parameter $r \in [0, \infty)$. Thinning (or splitting) this process works just like thinning the Poisson process on $[0, \infty)$. Specifically, suppose that the each random point, independently of the others is either type 1 with probability p or type 0 with probability $1 - p$, where $p \in (0, 1)$ is a new parameter. Let N_1 and N_0 denote the random counting measures associated with the type 1 and type 0 points, respectively. That is, $N_i(A)$ is the number of type i random points in A , for $A \in \mathcal{S}$ and $i \in \{0, 1\}$.

N_0 and N_1 are independent Poisson processes on (S, \mathcal{S}, μ) with density parameters pr and $(1-p)r$, respectively.

Proof

The proof is like the one for the Poisson process on $[0, \infty)$. For $j, k \in \mathbb{N}$,

$$\mathbb{P}[N_0(A) = j, N_1(A) = k] = \mathbb{P}[N_1(A) = k, N(A) = j + k] = \mathbb{P}[N(A) = j + k] \mathbb{P}[N_1(A) = k | N_0(A) = j + k] \quad (14.8.6)$$

But given $N(A) = n$, the number of type 1 points $N_1(A)$ has the binomial distribution with parameters n and p . Hence letting $t = \mu(A)$ to simplify the notation, we have

$$\mathbb{P}[N_0(A) = j, N_1(A) = k] = e^{-rt} \frac{(rt)^{j+k}}{(j+k)!} \frac{(j+k)!}{j!k!} p^k (1-p)^j = e^{-prt} \frac{(prt)^k}{k!} e^{-(1-p)rt} \frac{[(1-p)rt]^j}{j!} \quad (14.8.7)$$

It follows from the factorization theorem that $N_0(A)$ has the Poisson distribution with parameter $p\mu(A)$, $N_1(A)$ has the Poisson distribution with parameter $(1-p)\mu(A)$, and $N_0(A)$ and $N_1(A)$ are independent. Next suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} . Then $\{N_0(A_i) : i \in I\}$ and $\{N_1(A_i) : i \in I\}$ are each independent sets of random variables, and the two sets are independent of each other.

This result extends naturally to $k \in \mathbb{N}_+$ types. As in the standard case, combining independent Poisson processes produces a new Poisson process, and the density parameters add.

Suppose that N_0 and N_1 are independent Poisson processes on (S, \mathcal{S}, μ) , with density parameters r_0 and r_1 , respectively. Then the process obtained by combining the random points is also a Poisson process on (S, \mathcal{S}, μ) with density parameter $r_0 + r_1$.

Proof

The new random measure, of course, is simply $N = N_1 + N_2$. Thus for $A \in \mathcal{S}$, $N(A) = N_1(A) + N_2(A)$. But $N_i(A)$ has the Poisson distribution with parameter $r_i\mu(A)$ for $i \in \{1, 2\}$, and the variables are independent, so $N(A)$ has the Poisson distribution with parameter $r_0\mu(A) + r_1\mu(A) = (r_0 + r_1)\mu(A)$. Next suppose that $\{A_i : i \in I\}$ is a countable, disjoint collection of sets in \mathcal{S} . Then $\{N(A_i) : i \in I\} = \{N_0(A_i) + N_1(A_i) : i \in I\}$ is a set of independent random variables.

Applications and Special Cases

Non-homogeneous Poisson Processes

A non-homogeneous Poisson process on $[0, \infty)$ can be thought of simply as a Poisson process on $[0, \infty)$ with respect to a measure that is not the standard Lebesgue measure λ_1 on $[0, \infty)$. Thus suppose that $r : [0, \infty) \rightarrow (0, \infty)$ is piece-wise continuous with $\int_0^\infty r(t) dt = \infty$, and let

$$m(t) = \int_0^t r(s) ds, \quad t \in [0, \infty) \quad (14.8.8)$$

Consider the non-homogeneous Poisson process with rate function r (and hence mean function m). Recall that the Lebesgue-Stieltjes measure on $[0, \infty)$ associated with m (which we also denote by m) is defined by the condition

$$m(a, b] = m(b) - m(a), \quad a, b \in [0, \infty), \quad a < b \quad (14.8.9)$$

Equivalently, m is the measure that is absolutely continuous with respect to λ_1 , with density function r . That is, if A is a measurable subset of $[0, \infty)$ then

$$m(A) = \int_A r(t) dt \quad (14.8.10)$$

The non-homogeneous Poisson process on $[0, \infty)$ with rate function r is the Poisson process on $[0, \infty)$ with respect to the measure m .

Proof

This follows directly from the definitions. If N denotes the counting process associated with the non-homogeneous Poisson process, then N has stationary increments, and for $s, t \in [0, \infty)$ with $s < t$, $N(s, t]$ has the Poisson distribution with parameter $m(t) - m(s) = m(s, t]$.

Nearest Points in \mathbb{R}^d

In this subsection, we consider a rather specialized topic, but one that is fun and interesting. Consider the Poisson process on $(\mathbb{R}^d, \mathcal{R}_d, \lambda_d)$ with density parameter $r > 0$, where as usual, \mathcal{R}_d is the σ -algebra of Lebesgue measurable subsets of \mathbb{R}^d , and λ_d is d -dimensional Lebesgue measure. We use the usual Euclidean norm on \mathbb{R}^d :

$$\|\mathbf{x}\|_d = (x_1^2 + x_2^2 + \dots + x_d^2)^{1/2}, \quad \mathbf{x} = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d \quad (14.8.11)$$

For $t > 0$, let $B_t = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_d \leq t\}$ denote the ball of radius t centered at the origin. Recall that $\lambda_d(B_t) = c_d t^d$ where

$$c_d = \frac{\pi^{d/2}}{\Gamma(d/2 + 1)} \quad (14.8.12)$$

is the measure of the unit ball in \mathbb{R}^d , and where Γ is the gamma function. Of course, $c_1 = 2$, $c_2 = \pi$, $c_3 = 4\pi/3$.

For $t \geq 0$, let $M_t = N(B_t)$, the number of random points in the ball B_t , or equivalently, the number of random points within distance t of the origin. From our formula for the measure of B_t above, it follows that M_t has the Poisson distribution with parameter $rc_d t^d$.

Now let $Z_0 = 0$ and for $n \in \mathbb{N}_+$ let Z_n denote the distance of the n th closest random point to the origin. Note that Z_n is analogous to the n th arrival time for the Poisson process on $[0, \infty)$. Clearly the processes $\mathbf{M} = (M_t : t \geq 0)$ and $\mathbf{Z} = (Z_0, Z_1, \dots)$ are inverses of each other in the sense that $Z_n \leq t$ if and only if $M_t \geq n$. Both of these events mean that there are at least n random points within distance t of the origin.

Distributions

- $c_d Z_n^d$ has the gamma distribution with shape parameter n and rate parameter r .
- Z_n has probability density function g_n given by

$$g_n(z) = \frac{d(c_d r)^n z^{nd-1}}{(n-1)!} \exp(-rc_d z^d), \quad 0 \leq z < \infty \quad (14.8.13)$$

Proof

Let $T_n = c_d Z_n^d$.

- From the inverse relationship above,

$$\mathbb{P}(T_n \leq t) = \mathbb{P}\left[Z_n \leq (t/c_d)^{1/d}\right] = \mathbb{P}\left\{M\left[(t/c_d)^{1/d}\right] \geq n\right\} \quad (14.8.14)$$

But $M\left[(t/c_d)^{1/d}\right]$ has the Poisson distribution with parameter $rc_d\left[(t/c_d)^{1/d}\right]^d = rt$ so

$$\mathbb{P}(T_n \leq t) = \sum_{k=n}^{\infty} e^{-rt} \frac{(rt)^k}{k!} \quad (14.8.15)$$

which we know is the gamma CDF with parameters n and r

- Let f_n denote the gamma PDF with parameters n and r and let $t = c_d z^d$. From the standard change of variables formula,

$$g_n(z) = f_n(t) \frac{dt}{dz} \quad (14.8.16)$$

Substituting and simplifying gives the result.

$c_d Z_n^d - c_d Z_{n-1}^d$ are independent for $n \in \mathbb{N}_+$ and each has the exponential distribution with rate parameter r .

Computational Exercises

Suppose that defects in a sheet of material follow the Poisson model with an average of 1 defect per 2 square meters. Consider a 5 square meter sheet of material.

1. Find the probability that there will be at least 3 defects.
2. Find the mean and standard deviation of the number of defects.

Answer

1. 0.4562
2. 2.5, 1.581

Suppose that raisins in a cake follow the Poisson model with an average of 2 raisins per cubic inch. Consider a slab of cake that measures 3 by 4 by 1 inches.

1. Find the probability that there will be at no more than 20 raisins.
2. Find the mean and standard deviation of the number of raisins.

Answer

1. 0.2426
2. 24, 4.899

Suppose that the occurrence of trees in a forest of a certain type that exceed a certain critical size follows the Poisson model. In a one-half square mile region of the forest there are 40 trees that exceed the specified size.

1. Estimate the density parameter.
2. Using the estimated density parameter, find the probability of finding at least 100 trees that exceed the specified size in a square mile region of the forest

Answer

1. $r = 80$ per square mile
2. 0.0171

Suppose that defects in a type of material follow the Poisson model. It is known that a square sheet with side length 2 meters contains one defect. Find the probability that the defect is in a circular region of the material with radius $\frac{1}{4}$ meter.

Answer

0.0491

Suppose that raisins in a cake follow the Poisson model. A 6 cubic inch piece of the cake with 20 raisins is divided into 3 equal parts. Find the probability that each piece has at least 6 raisins.

Answer

0.2146

Suppose that defects in a sheet of material follow the Poisson model, with an average of 5 defects per square meter. Each defect, independently of the others is *mild* with probability 0.5, *moderate* with probability 0.3, or *severe* with probability 0.2. Consider a circular piece of the material with radius 1 meter.

1. Give the mean and standard deviation of the number of defects of each type in the piece.
2. Find the probability that there will be at least 2 defects of each type in the piece.

Answer

1. Mild: 7.854, 2.802; Moderate: 4.712, 2.171; Severe: 3.142, 1.772
2. 0.7762

This page titled [14.8: Poisson Processes on General Spaces](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

15: Renewal Processes

A *renewal process* is an idealized stochastic model for “events” that occur randomly in time (generically called renewals or arrivals). The basic mathematical assumption is that the times between the successive arrivals are independent and identically distributed. Renewal processes have a very rich and interesting mathematical structure and can be used as a foundation for building more realistic models. Moreover, renewal processes are often found embedded in other stochastic processes, most notably Markov chains.

[15.1: Introduction](#)

[15.2: Renewal Equations](#)

[15.3: Renewal Limit Theorems](#)

[15.4: Delayed Renewal Processes](#)

[15.5: Alternating Renewal Processes](#)

[15.6: Renewal Reward Processes](#)

This page titled [15: Renewal Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

15.1: Introduction

A *renewal process* is an idealized stochastic model for “events” that occur randomly in time. These temporal events are generically referred to as *renewals* or *arrivals*. Here are some typical interpretations and applications.

- The arrivals are “customers” arriving at a “service station”. Again, the terms are generic. A customer might be a person and the service station a store, but also a customer might be a file request and the service station a web server.
- A device is placed in service and eventually fails. It is replaced by a device of the same type and the process is repeated. We do not count the replacement time in our analysis; equivalently we can assume that the replacement is immediate. The times of the replacements are the renewals
- The arrivals are times of some natural event, such as a lightening strike, a tornado or an earthquake, at a particular geographical point.
- The arrivals are emissions of elementary particles from a radioactive source.

Basic Processes

The basic model actually gives rise to several interrelated random processes: the sequence of interarrival times, the sequence of arrival times, and the counting process. The term *renewal process* can refer to any (or all) of these. There are also several natural “age” processes that arise. In this section we will define and study the basic properties of each of these processes in turn.

Interarrival Times

Let X_1 denote the time of the first arrival, and X_i the time between the $(i-1)$ st and i th arrivals for $i \in \{2, 3, \dots\}$. Our basic assumption is that the sequence of *interarrival times* $\mathbf{X} = (X_1, X_2, \dots)$ is an independent, identically distributed sequence of random variables. In statistical terms, \mathbf{X} corresponds to sampling from the distribution of a generic interarrival time X . We assume that X takes values in $[0, \infty)$ and $\mathbb{P}(X > 0) > 0$, so that the interarrival times are nonnegative, but not identically 0. Let $\mu = \mathbb{E}(X)$ denote the common mean of the interarrival times. We allow that possibility that $\mu = \infty$. On the other hand,

$\mu > 0$.

Proof

This is a basic fact from properties of expected value. For a simple proof, note that if $\mu = 0$ then $\mathbb{P}(X > x) = 0$ for every $x > 0$ by Markov's inequality. But then $\mathbb{P}(X = 0) = 1$.

If $\mu < \infty$, we will let $\sigma^2 = \text{var}(X)$ denote the common variance of the interarrival times. Let F denote the common distribution function of the interarrival times, so that

$$F(x) = \mathbb{P}(X \leq x), \quad x \in [0, \infty) \quad (15.1.1)$$

The distribution function F turns out to be of fundamental importance in the study of renewal processes. We will let f denote the probability density function of the interarrival times if the distribution is discrete or if the distribution is continuous and has a probability density function (that is, if the distribution is absolutely continuous with respect to Lebesgue measure on $[0, \infty)$). In the discrete case, the following definition turns out to be important:

If X takes values in the set $\{nd : n \in \mathbb{N}\}$ for some $d \in (0, \infty)$, then X (or its distribution) is said to be *arithmetic* (the terms *lattice* and *periodic* are also used). The largest such d is the *span* of X .

The reason the definition is important is because the limiting behavior of renewal processes turns out to be more complicated when the interarrival distribution is arithmetic.

The Arrival Times

Let

$$T_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (15.1.2)$$

We follow our usual convention that the sum over an empty index set is 0; thus $T_0 = 0$. On the other hand, T_n is the time of the n th arrival for $n \in \mathbb{N}_+$. The sequence $\mathbf{T} = (T_0, T_1, \dots)$ is called the *arrival time process*, although note that T_0 is not considered an arrival. A renewal process is so named because the process starts over, independently of the past, at each arrival time.




Figure 15.1.1: The interarrival times and arrival times

The sequence \mathbf{T} is the *partial sum process* associated with the independent, identically distributed sequence of interarrival times \mathbf{X} . Partial sum processes associated with independent, identically distributed sequences have been studied in several places in this project. In the remainder of this subsection, we will collect some of the more important facts about such processes. First, we can recover the interarrival times from the arrival times:

$$X_i = T_i - T_{i-1}, \quad i \in \mathbb{N}_+ \quad (15.1.3)$$

Next, let F_n denote the distribution function of T_n , so that

$$F_n(t) = \mathbb{P}(T_n \leq t), \quad t \in [0, \infty) \quad (15.1.4)$$

Recall that if X has probability density function f (in either the discrete or continuous case), then T_n has probability density function $f_n = f^{*n} = f * f * \dots * f$, the n -fold convolution power of f .

The sequence of arrival times \mathbf{T} has *stationary, independent increments*:

1. If $m \leq n$ then $T_n - T_m$ has the same distribution as T_{n-m} and thus has distribution function F_{n-m}
2. If $n_1 \leq n_2 \leq n_3 \leq \dots$ then $(T_{n_1}, T_{n_2} - T_{n_1}, T_{n_3} - T_{n_2}, \dots)$ is a sequence of independent random variables.

Proof

Recall that these are properties that hold generally for the partial sum sequence associated with a sequence of IID variables.

If $n, m \in \mathbb{N}$ then

1. $\mathbb{E}(T_n) = n\mu$
2. $\text{var}(T_n) = n\sigma^2$
3. $\text{cov}(T_m, T_n) = \min\{m, n\}\sigma^2$

Proof

Part (a) follows, of course, from the additive property of expected value, and part (b) from the additive property of variance for sums of *independent* variables. For part (c), assume that $m \leq n$. Then $T_n = T_m + (T_n - T_m)$. But T_m and $T_n - T_m$ are independent, so

$$\text{cov}(T_m, T_n) = \text{cov}[T_m, T_m + (T_n - T_m)] = \text{cov}(T_m, T_m) + \text{cov}(T_m, T_n - T_m) = \text{var}(T_m) = m\sigma^2 \quad (15.1.5)$$

Recall the law of large numbers: $T_n/n \rightarrow \mu$ as $n \rightarrow \infty$

1. With probability 1 (the strong law).
2. In probability (the weak law).

Note that $T_n \leq T_{n+1}$ for $n \in \mathbb{N}$ since the interarrival times are nonnegative. Also $\mathbb{P}(T_n = T_{n-1}) = \mathbb{P}(X_n = 0) = F(0)$. This can be positive, so with positive probability, more than one arrival can occur at the same time. On the other hand, the arrival times are unbounded:

$T_n \rightarrow \infty$ as $n \rightarrow \infty$ with probability 1.

Proof

Since $\mathbb{P}(X > 0) > 0$, there exists $t > 0$ such that $\mathbb{P}(X > t) > 0$. From the second Borel-Cantelli lemma it follows that with probability 1, $X_i > t$ for infinitely many $n \in \mathbb{N}_+$. Therefore $\sum_{i=1}^{\infty} X_i = \infty$ with probability 1.

The Counting Process

For $t \geq 0$, let N_t denote the number of arrivals in the interval $[0, t]$:

$$N_t = \sum_{n=1}^{\infty} \mathbf{1}(T_n \leq t), \quad t \in [0, \infty) \quad (15.1.6)$$

We will refer to the random process $\mathbf{N} = (N_t : t \geq 0)$ as the *counting process*. Recall again that $T_0 = 0$ is not considered an arrival, but it's possible to have $T_n = 0$ for $n \in \mathbb{N}_+$, so there may be one or more arrivals at time 0.

$$N_t = \max\{n \in \mathbb{N} : T_n \leq t\} \text{ for } t \geq 0.$$

If $s, t \in [0, \infty)$ and $s \leq t$ then $N_t - N_s$ is the number of arrivals in $(s, t]$.

Note that as a function of t , N_t is a (random) step function with jumps at the distinct values of (T_1, T_2, \dots) ; the size of the jump at an arrival time is the number of arrivals at that time. In particular, N is an increasing function of t .




Figure 15.1.2: The counting process

More generally, we can define the (random) counting measure corresponding to the sequence of random points (T_1, T_2, \dots) in $[0, \infty)$. Thus, if A is a (measurable) subset of $[0, \infty)$, we will let $N(A)$ denote the number of the random points in A :

$$N(A) = \sum_{n=1}^{\infty} \mathbf{1}(T_n \in A) \quad (15.1.7)$$

In particular, note that with our new notation, $N_t = N[0, t]$ for $t \geq 0$ and $N(s, t] = N_t - N_s$ for $s \leq t$. Thus, the random counting measure is completely determined by the counting process. The counting process is the “cumulative measure function” for the counting measure, analogous the cumulative distribution function of a probability measure.

For $t \geq 0$ and $n \in \mathbb{N}$,

1. $T_n \leq t$ if and only if $N_t \geq n$
2. $N_t = n$ if and only if $T_n \leq t < T_{n+1}$

Proof

Note that the event in part (a) means that there are at least n arrivals in $[0, t]$. The event in part (b) means that there are exactly n arrivals in $[0, t]$.

Of course, the complements of the events in (a) are also equivalent, so $T_n > t$ if and only if $N_t < n$. On the other hand, neither of the events $N_t \leq n$ and $T_n \geq t$ implies the other. For example, we could easily have $N_t = n$ and $T_n < t < T_{n+1}$. Taking complements, neither of the events $N_t > n$ and $T_n < t$ implies the other. The last result also shows that the arrival time process \mathbf{T} and the counting process \mathbf{N} are inverses of each other in a sense.

The following events have probability 1:

1. $N_t < \infty$ for all $t \in [0, \infty)$
2. $N_t \rightarrow \infty$ as $t \rightarrow \infty$

Proof

The event in part (a) occurs if and only if $T_n \rightarrow \infty$ as $n \rightarrow \infty$, which occurs with probability 1 by the result [above](#). The event in part (b) occurs if and only if $T_n < \infty$ for all $n \in \mathbb{N}$ which also occurs with probability 1.

All of the results so far in this subsection show that the arrival time process \mathbf{T} and the counting process \mathbf{N} are inverses of one another in a sense. The important [equivalences](#) above can be used to obtain the probability distribution of the counting variables in terms of the interarrival distribution function F .

For $t \geq 0$ and $n \in \mathbb{N}$,

1. $\mathbb{P}(N_t \geq n) = F_n(t)$
2. $\mathbb{P}(N_t = n) = F_n(t) - F_{n+1}(t)$

The next result is little more than a restatement of the result [above](#) relating the counting process and the arrival time process. However, you may need to review the section on filtrations and stopping times to understand the result

For $t \in [0, \infty)$, $N_t + 1$ is a stopping time for the sequence of interarrival times \mathbf{X}

Proof

Note that $N_t + 1$ takes values in \mathbb{N}_+ , so we need to show that the event $\{N_t + 1 = n\}$ is measurable with respect to $\mathcal{F}_n = \sigma\{X_1, X_2, \dots, X_n\}$ for $n \in \mathbb{N}_+$. But from the result [above](#), $N_t + 1 = n$ if and only if $N_t = n - 1$ if and only if $T_{n-1} \leq t < T_n$. The last event is clearly measurable with respect to \mathcal{F}_n .

The Renewal Function

The function M that gives the expected number of arrivals up to time t is known as the *renewal function*:

$$M(t) = \mathbb{E}(N_t), \quad t \in [0, \infty) \quad (15.1.8)$$

The renewal function turns out to be of fundamental importance in the study of renewal processes. Indeed, the renewal function essentially characterizes the renewal process. It will take awhile to fully understand this, but the following theorem is a first step:

The renewal function is given in terms of the interarrival distribution function by

$$M(t) = \sum_{n=1}^{\infty} F_n(t), \quad 0 \leq t < \infty \quad (15.1.9)$$

Proof

Recall that $N_t = \sum_{n=1}^{\infty} \mathbf{1}(T_n \leq t)$. Taking expected values gives the result. Note that the interchange of sum and expected value is valid because the terms are nonnegative.

Note that we have not yet shown that $M(t) < \infty$ for $t \geq 0$, and note also that this does not follow from the previous theorem. However, we will establish this finiteness condition in the subsection on [moment generating functions](#) below. If M is differentiable, the derivative $m = M'$ is known as the *renewal density*, so that $m(t)$ gives the expected rate of arrivals per unit time at $t \in [0, \infty)$.

More generally, if A is a (measurable) subset of $[0, \infty)$, let $M(A) = \mathbb{E}[N(A)]$, the expected number of arrivals in A .

M is a positive measure on $[0, \infty)$. This measure is known as the *renewal measure*.

Proof

N is a measure on $[0, \infty)$ (albeit a random one). So if (A_1, A_2, \dots) is a sequence of disjoint, measurable subsets of $[0, \infty)$ then

$$N\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} N(A_i) \quad (15.1.10)$$

Taking expected values gives

$$m\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} m(A_i) \quad (15.1.11)$$

Again, the interchange of sum and expected value is justified since the terms are nonnegative.

The renewal measure is also given by

$$M(A) = \sum_{n=1}^{\infty} \mathbb{P}(T_n \in A), \quad A \subset [0, \infty) \quad (15.1.12)$$

Proof

Recall that $N(A) = \sum_{n=1}^{\infty} \mathbf{1}(T_n \in A)$. Taking expected values gives the result. Again, the interchange of expected value and infinite series is justified since the terms are nonnegative.

If $s, t \in [0, \infty)$ with $s \leq t$ then $M(t) - M(s) = m(s, t]$, the expected number of arrivals in $(s, t]$.

The last theorem implies that the renewal function actually determines the entire renewal measure. The renewal function is the “cumulative measure function”, analogous to the cumulative distribution function of a probability measure. Thus, every renewal process naturally leads to two measures on $[0, \infty)$, the random counting measure corresponding to the arrival times, and the measure associated with the expected number of arrivals.

The Age Processes

For $t \in [0, \infty)$, $T_{N_t} \leq t < T_{N_t+1}$. That is, t is in the random renewal interval $[T_{N_t}, T_{N_t+1})$.

Consider the reliability setting in which whenever a device fails, it is immediately replaced by a new device of the same type. Then the sequence of interarrival times \mathbf{X} is the sequence of lifetimes, while T_n is the time that the n th device is placed in service. There are several other natural random processes that can be defined.

The random variable

$$C_t = t - T_{N_t}, \quad t \in [0, \infty) \quad (15.1.13)$$

is called the *current life* at time t . This variable takes values in the interval $[0, t]$ and is the age of the device that is in service at time t . The random process $\mathbf{C} = (C_t : t \geq 0)$ is the *current life process*.

The random variable

$$R_t = T_{N_t+1} - t, \quad t \in [0, \infty) \quad (15.1.14)$$

is called the *remaining life* at time t . This variable takes values in the interval $(0, \infty)$ and is the time remaining until the device that is in service at time t fails. The random process $\mathbf{R} = (R_t : t \geq 0)$ is the *remaining life process*.




Figure 15.1.3: The current and remaining life at time t

The random variable

$$L_t = C_t + R_t = T_{N_t+1} - T_{N_t} = X_{N_t+1}, \quad t \in [0, \infty) \quad (15.1.15)$$

is called the *total life* at time t . This variable takes values in $[0, \infty)$ and gives the total life of the device that is in service at time t . The random process $\mathbf{L} = (L_t : t \geq 0)$ is the *total life process*.

Tail events of the current and remaining life can be written in terms of each other and in terms of the counting variables.

Suppose that $t \in [0, \infty)$, $x \in [0, t]$, and $y \in [0, \infty)$. Then

1. $\{R_t > y\} = \{N_{t+y} - N_t = 0\}$
2. $\{C_t \geq x\} = \{R_{t-x} > x\} = \{N_t - N_{t-x} = 0\}$
3. $\{C_t \geq x, R_t > y\} = \{R_{t-x} > x+y\} = \{N_{t+y} - N_{t-x} = 0\}$

Proof




Figure 15.1.4: The events of interest for the current and remaining life

Of course, the various equivalent events in the last result must have the same probability. In particular, it follows that if we know the distribution of R_t for all t then we also know the distribution of C_t for all t , and in fact we know the joint distribution of (R_t, C_t) for all t and hence also the distribution of L_t for all t .

For fixed $t \in (0, \infty)$ the total life at t (the lifetime of the device in service at time t) is stochastically larger than a generic lifetime. This result, a bit surprising at first, is known as the *inspection paradox*. Let X denote fixed interarrival time.

$$\mathbb{P}(L_t > x) \geq \mathbb{P}(X > x) \text{ for } x \geq 0.$$

Proof

Recall that $L_t = X_{N_t+1}$. The proof is by conditioning on N_t . An important tool is the fact that if A and B are nested events in a probability space (one a subset of the other), then the events are positively correlated, so that $\mathbb{P}(A | B) \geq \mathbb{P}(A)$. Recall that F is the common CDF of the interarrival times. First

$$\mathbb{P}(X_{N_t+1} > x | N_t = 0) = \mathbb{P}(X_1 > x | X_1 > t) \geq \mathbb{P}(X_1 > x) = 1 - F(x) \quad (15.1.16)$$

Next, for $n \in \mathbb{N}_+$,

$$\mathbb{P}(X_{N_t+1} > x | N_t = n) = \mathbb{P}(X_{n+1} > x | T_n \leq t < T_{n+1}) = \mathbb{P}(X_{n+1} > x | T_n \leq t < T_n + X_{n+1}) \quad (15.1.17)$$

We condition this additionally on T_n , the time of the n th arrival. For $s \leq t$, and since X_{n+1} is independent of T_n , we have

$$\mathbb{P}(X_{n+1} > x | T_n = s, X_{n+1} > t-s) = \mathbb{P}(X_{n+1} > x | X_{n+1} > t-s) \geq \mathbb{P}(X_{n+1} > x) = 1 - F(x) \quad (15.1.18)$$

It follows that $\mathbb{P}(X_{N_t+1} > x | N_t = n) \geq 1 - F(x)$ for every $n \in \mathbb{N}$, and hence

$$\mathbb{P}(X_{N_t+1} > x) = \sum_{n=0}^{\infty} \mathbb{P}(X_{N_t+1} > x | N_t = n) \mathbb{P}(N_t = n) \geq \sum_{n=0}^{\infty} [1 - F(x)] \mathbb{P}(N_t = n) = 1 - F(x) \quad (15.1.19)$$

Basic Comparison

The basic comparison in the following result is often useful, particularly for obtaining various bounds. The idea is very simple: if the interarrival times are shortened, the arrivals occur more frequently.

Suppose now that we have two interarrival sequences, $\mathbf{X} = (X_1, X_2, \dots)$ and $\mathbf{Y} = (Y_1, Y_2, \dots)$ defined on the same probability space, with $Y_i \leq X_i$ (with probability 1) for each i . Then for $n \in \mathbb{N}$ and $t \in [0, \infty)$,

1. $T_{Y,n} \leq T_{X,n}$
2. $N_{Y,t} \geq N_{X,t}$
3. $m_Y(t) \geq m_X(t)$

Examples and Special Cases

Bernoulli Trials

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter $p \in (0, 1)$. Recall that \mathbf{X} is a sequence of independent, identically distributed indicator variables with $\mathbb{P}(X = 1) = p$.

Recall the random processes derived from \mathbf{X} :

1. $\mathbf{Y} = (Y_0, Y_1, \dots)$ where Y_n the number of success in the first n trials. The sequence \mathbf{Y} is the partial sum process associated with \mathbf{X} . The variable Y_n has the binomial distribution with parameters n and p .

1. $\mathbf{U} = (U_1, U_2, \dots)$ where U_n the number of trials needed to go from success number $n - 1$ to success number n . These are independent variables, each having the geometric distribution on \mathbb{N}_+ with parameter p .
2. $\mathbf{V} = (V_0, V_1, \dots)$ where V_n is the trial number of success n . The sequence \mathbf{V} is the partial sum process associated with \mathbf{U} . The variable V_n has the negative binomial distribution with parameters n and p .

It is natural to view the successes as arrivals in a discrete-time renewal process.

Consider the renewal process with interarrival sequence \mathbf{U} . Then

1. The basic assumptions are satisfied and that the mean interarrival time is $\mu = 1/p$.
2. \mathbf{V} is the sequence of arrival times.
3. \mathbf{Y} is the counting process (restricted to \mathbb{N}).
4. The renewal function is $m(n) = np$ for $n \in \mathbb{N}$.

It follows that the renewal measure is proportional to counting measure on \mathbb{N}_+ .

Run the binomial timeline experiment 1000 times for various values of the parameters n and p . Compare the empirical distribution of the counting variable to the true distribution.

Run the negative binomial experiment 1000 times for various values of the parameters k and p . Compare the empirical distribution of the arrival time to the true distribution.

Consider again the renewal process with interarrival sequence \mathbf{U} . For $n \in \mathbb{N}$,

1. The current life and remaining life at time n are independent.
2. The remaining life at time n has the same distribution as an interarrival time U , namely the geometric distribution on \mathbb{N}_+ with parameter p .
3. The current life at time n has a truncated geometric distribution with parameters n and p :

$$\mathbb{P}(C_n = k) = \begin{cases} p(1-p)^k, & k \in \{0, 1, \dots, n-1\} \\ (1-p)^n, & k = n \end{cases} \quad (15.1.20)$$

Proof

These results follow from [age process events](#) above.

This renewal process starts over, independently of the past, not only at the arrival times, but at fixed times $n \in \mathbb{N}$ as well. The Bernoulli trials process (with the successes as arrivals) is the only discrete-time renewal process with this property, which is a consequence of the memoryless property of the geometric interarrival distribution.

We can also use the indicator variables as the interarrival times. This may seem strange at first, but actually turns out to be useful.

Consider the renewal process with interarrival sequence \mathbf{X} .

1. The basic assumptions are satisfied and that the mean interarrival time is $\mu = p$.
2. \mathbf{Y} is the sequence of arrival times.
3. The number of arrivals at time 0 is $U_1 - 1$ and the number of arrivals at time $i \in \mathbb{N}_+$ is U_{i+1} .
4. The number of arrivals in the interval $[0, n]$ is $V_{n+1} - 1$ for $n \in \mathbb{N}$. This gives the counting process.
5. The renewal function is $m(n) = \frac{n+1}{p} - 1$ for $n \in \mathbb{N}$.

The age processes are not very interesting for this renewal process.

For $n \in \mathbb{N}$ (with probability 1),

1. $C_n = 0$
2. $R_n = 1$

The Moment Generating Function of the Counting Variables

As an application of the last renewal process, we can show that the moment generating function of the counting variable N_t in an arbitrary renewal process is finite in an interval about 0 for every $t \in [0, \infty)$. This implies that N_t has finite moments of all orders and in particular that $m(t) < \infty$ for every $t \in [0, \infty)$.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is the interarrival sequence for a renewal process. By the basic assumptions, there exists $a > 0$ such that $p = \mathbb{P}(X \geq a) > 0$. We now consider the renewal process with interarrival sequence $\mathbf{X}_a = (X_{a,1}, X_{a,2}, \dots)$, where $X_{a,i} = a \mathbf{1}(X_i \geq a)$ for $i \in \mathbb{N}_+$. The renewal process with interarrival sequence \mathbf{X}_a is just like the renewal process with [Bernoulli interarrivals](#), except that the arrival times occur at the points in the sequence $(0, a, 2a, \dots)$, instead of $(0, 1, 2, \dots)$.

For each $t \in [0, \infty)$, N_t has finite moment generating function in an interval about 0, and hence N_t has moments of all orders at 0.

Proof

Note first that $X_{a,i} \leq X_i$ for each $i \in \mathbb{N}_+$. Recall the moment generating function Γ of the geometric distribution with parameter p is

$$\Gamma(s) = \frac{e^s p}{1 - (1-p)e^s}, \quad s < -\ln(1-p) \quad (15.1.21)$$

But as with the process with [Bernoulli interarrival times](#), $N_{a,t}$ can be written as $a(V_{n+1} - 1)$ where $n = \lfloor a/t \rfloor$ and where V_{n+1} is a sum of $n+1$ IID geometric variables, each with parameter p . We don't really care about the explicit form of the MGF of $N_{a,t}$, but it is clearly finite in an interval of the form $(-\infty, \epsilon)$ where $\epsilon > 0$. But $N_t \leq N_{t,a}$, so its MGF is also finite on this interval.

The Poisson Process

The *Poisson process*, named after Simeon Poisson, is the most important of all renewal processes. The Poisson process is so important that it is treated in a separate chapter in this project. Please review the essential properties of this process:

Properties of the Poisson process with rate $r \in (0, \infty)$.

1. The interarrival times have an exponential distribution with rate parameter r . Thus, the basic assumptions above are satisfied and the mean interarrival time is $\mu = 1/r$.
2. The exponential distribution is the only distribution with the memoryless property on $[0, \infty)$.
3. The time of the n th arrival T_n has the gamma distribution with shape parameter n and rate parameter r .
4. The counting process $\mathbf{N} = (N_t : t \geq 0)$ has stationary, independent increments and N_t has the Poisson distribution with parameter rt for $t \in [0, \infty)$.
5. In particular, the renewal function is $m(t) = rt$ for $t \in [0, \infty)$. Hence, the renewal measure is a multiple of the standard length measure (Lebesgue measure) on $[0, \infty)$.

Consider again the Poisson process with rate parameter r . For $t \in [0, \infty)$,

1. The current life and remaining life at time t are independent.
2. The remaining life at time t has the same distribution as an interarrival time X , namely the exponential distribution with rate parameter r .
3. The current life at time t has a truncated exponential distribution with parameters t and r :

$$\mathbb{P}(C_t \geq s) = \begin{cases} e^{-rs}, & 0 \leq s \leq t \\ 0, & s > t \end{cases} \quad (15.1.22)$$

Proof

These results follow from [age process events](#) given above.

The Poisson process starts over, independently of the past, not only at the arrival times, but at fixed times $t \in [0, \infty)$ as well. The Poisson process is the only renewal process with this property, which is a consequence of the memoryless property of the exponential interarrival distribution.

Run the Poisson experiment 1000 times for various values of the parameters t and r . Compare the empirical distribution of the counting variable to the true distribution.

Run the gamma experiment 1000 times for various values of the parameters n and r . Compare the empirical distribution of the arrival time to the true distribution.

Simulation Exercises

Open the renewal experiment and set $t = 10$. For each of the following interarrival distributions, run the simulation 1000 times and note the shape and location of the empirical distribution of the counting variable. Note also the mean of the interarrival distribution in each case.

1. The continuous uniform distribution on the interval $[0, 1]$ (the standard uniform distribution).
2. the discrete uniform distribution starting at $a = 0$, with step size $h = 0.1$, and with $n = 10$ points.
3. The gamma distribution with shape parameter $k = 2$ and scale parameter $b = 1$.
4. The beta distribution with left shape parameter $a = 3$ and right shape parameter $b = 2$.
5. The exponential-logarithmic distribution with shape parameter $p = 0.1$ and scale parameter $b = 1$.
6. The Gompertz distribution with shape parameter $a = 1$ and scale parameter $b = 1$.
7. The Wald distribution with mean $\mu = 1$ and shape parameter $\lambda = 1$.
8. The Weibull distribution with shape parameter $k = 2$ and scale parameter $b = 1$.

This page titled [15.1: Introduction](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

15.2: Renewal Equations

Many quantities of interest in the study of renewal processes can be described by a special type of integral equation known as a *renewal equation*. Renewal equations almost always arise by conditioning on the time of the first arrival and by using the defining property of a renewal process—the fact that the process restarts at each arrival time, independently of the past. However, before we can study renewal equations, we need to develop some additional concepts and tools involving measures, convolutions, and transforms. Some of the results in the advanced sections on measure theory, general distribution functions, the integral with respect to a measure, properties of the integral, and density functions are needed for this section. You may need to review some of these topics as necessary. As usual, we assume that all functions and sets that are mentioned are measurable with respect to the appropriate σ -algebras. In particular, $[0, \infty)$ which is our basic temporal space, is given the usual Borel σ -algebra generated by the intervals.

Measures, Integrals, and Transforms

Distribution Functions and Positive Measures

Recall that a distribution function on $[0, \infty)$ is a function $G : [0, \infty) \rightarrow [0, \infty)$ that is increasing and continuous from the right. The distribution function G defines a positive measure on $[0, \infty)$, which we will also denote by G , by means of the formula $G[0, t] = G(t)$ for $t \in [0, \infty)$.



Figure 15.2.1: $G(t)$ is the cumulative measure at t

Hopefully, our notation will not cause confusion and it will be clear from context whether G refers to the positive measure (a set function) or the distribution function (a point function). More generally, if $a, b \in [0, \infty)$ and $a \leq b$ then $G(a, b] = G(b) - G(a)$. Note that the positive measure associated with a distribution function is *locally finite* in the sense that $G(A) < \infty$ if $A \subseteq [0, \infty)$ is bounded. Of course, if A is unbounded, $G(A)$ may well be infinite. The basic structure of a distribution function and its associated positive measure occurred several times in our preliminary discussion of renewal processes:

Distributions associated with a renewal process.

1. The distribution function F of the interarrival times defines a probability measure on $[0, \infty)$
2. The counting process N defines a (random) counting measure on $[0, \infty)$
3. the renewal function M defines a (deterministic) positive measure on $[0, \infty)$

Suppose again that G is a distribution function on $[0, \infty)$. Recall that the integral associated with the positive measure G is also called the *Lebesgue-Stieltjes integral* associated with the distribution function G (named for Henri Lebesgue and Thomas Stieltjes). If $f : [0, \infty) \rightarrow \mathbb{R}$ and $A \subseteq [0, \infty)$ (measurable of course), the integral of f over A (if it exists) is denoted

$$\int_A f(t) dG(t) \quad (15.2.1)$$

We use the more conventional $\int_0^t f(x) dG(x)$ for the integral over $[0, t]$ and $\int_0^\infty f(x) dG(x)$ for the integral over $[0, \infty)$. On the other hand, $\int_s^t f(x) dG(x)$ means the integral over $(s, t]$ for $s < t$, and $\int_s^\infty f(x) dG(x)$ means the integral over (s, ∞) . Thus, the additivity of the integral over disjoint domains holds, as it must. For example, for $t \in [0, \infty)$,

$$\int_0^\infty f(x) dG(x) = \int_0^t f(x) dG(x) + \int_t^\infty f(x) dG(x) \quad (15.2.2)$$

This notation would be ambiguous without the clarification, but is consistent with how the measure works: $G[0, t] = G(t)$ for $t \geq 0$, $G(s, t] = G(t) - G(s)$ for $0 \leq s < t$, etc. Of course, if G is continuous as a *function*, so that G is also continuous as a *measure*, then none of this matters—the integral over an interval is the same whether or not endpoints are included. The following definition is a natural complement to the locally finite property of the positive measures that we are considering.

A function $f : [0, \infty) \rightarrow \mathbb{R}$ is *locally bounded* if it is measurable and is bounded on $[0, t]$ for each $t \in [0, \infty)$.

The locally bounded functions form a natural class for which our integrals of interest exist.

Suppose that G is a distribution function on $[0, \infty)$ and $f : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded. Then $g : [0, \infty) \rightarrow \mathbb{R}$ defined by $g(t) = \int_0^t f(s) dG(s)$ is also locally bounded.

Proof

Suppose that $|f(s)| \leq C_t$ for $s \in [0, t]$ and $t \in [0, \infty)$. Then

$$\int_0^s |f(x)| dG(x) \leq C_t G(s) \leq C_t G(t), \quad t \in [0, \infty) \quad (15.2.3)$$

Hence f is integrable on $[0, s]$ and the integral is bounded by $C_t G(t)$ for $s \in [0, t]$.

Note that if f and g are locally bounded, then so are $f + g$ and fg . If f is increasing on $[0, \infty)$ then f is locally bounded, so in particular, a distribution function on $[0, \infty)$ is locally bounded. If f is continuous on $[0, \infty)$ then f is locally bounded. Similarly, if G and H are distribution functions on $[0, \infty)$ and if $c \in (0, \infty)$, then $G + H$ and cG are also distribution functions on $[0, \infty)$. Convolution, which we consider next, is another way to construct new distributions on $[0, \infty)$ from ones that we already have.

Convolution

The term *convolution* means different things in different settings. Let's start with the definition we know, the convolution of probability density functions, on our space of interest $[0, \infty)$.

Suppose that X and Y are independent random variables with values in $[0, \infty)$ and with probability density functions f and g , respectively. Then $X + Y$ has probability density function $f * g$ given as follows, in the [discrete](#) and [continuous](#) cases, respectively

$$(f * g)(t) = \sum_{s \in [0, t]} f(t-s)g(s) \quad (15.2.4)$$

$$(f * g)(t) = \int_0^t f(t-s)g(s) ds \quad (15.2.5)$$

In the discrete case, it's understood that t is a possible value of $X + Y$, and the sum is over the countable collection of $s \in [0, t]$ with s a value of X and $t - s$ a value of Y . Often in this case, the random variables take values in \mathbb{N} , in which case the sum is simply over the set $\{0, 1, \dots, t\}$ for $t \in \mathbb{N}$. The discrete and continuous cases could be unified by defining convolution with respect to a general positive measure on $[0, \infty)$. Moreover, the definition clearly makes sense for functions that are not necessarily probability density functions.

Suppose that $f, g : [0, \infty) \rightarrow \mathbb{R}$ are locally bounded and that H is a distribution function on $[0, \infty)$. The *convolution* of f and g with respect to H is the function on $[0, \infty)$ defined by

$$t \mapsto \int_0^t f(t-s)g(s) dH(s) \quad (15.2.6)$$

If f and g are probability density functions for discrete distributions on a countable set $C \subseteq [0, \infty)$ and if H is counting measure on C , we get discrete convolution, as above. If f and g are probability density functions for continuous distributions on $[0, \infty)$ and if H is Lebesgue measure, we get continuous convolution, as above. Note however, that if g is nonnegative then $G(t) = \int_0^t g(s) dH(s)$ for $t \in [0, \infty)$ defines another distribution function on $[0, \infty)$, and the convolution integral above is simply $\int_0^t f(t-s) dG(s)$. This motivates our next version of convolution, the one that we will use in the remainder of this section.

Suppose that $f : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded and that G is a distribution function on $[0, \infty)$. The convolution of the *function* f with the *distribution* G is the function $f * G$ defined by

$$(f * G)(t) = \int_0^t f(t-s) dG(s), \quad t \in [0, \infty) \quad (15.2.7)$$

Note that if F and G are distribution functions on $[0, \infty)$, the convolution $F * G$ makes sense, with F simply as a *function* and G as a *distribution function*. The result is another distribution function. Moreover in this case, the operation is *commutative*.

If F and G are distribution functions on $[0, \infty)$ then $F * G$ is also a distribution function on $[0, \infty)$, and $F * G = G * F$

Proof

Let $F \otimes G$ and $G \otimes F$ denote the usual product measures on $[0, \infty)^2 = [0, \infty) \times [0, \infty)$. For $t \in [0, \infty)$, let $T_t = \{(r, s) \in [0, \infty)^2 : r + s \leq t\}$, the triangular region with vertices $(0, 0)$, $(t, 0)$, and $(0, t)$. Then

$$(F * G)(t) = \int_0^t F(t-s) dG(s) = \int_0^t \int_0^{t-s} dF(r) dG(s) = (F \otimes G)(T_t) \quad (15.2.8)$$

This clearly defines a distribution function. Specifically, if $0 \leq s \leq t < \infty$ then $T_s \subseteq T_t$ so $(F * G)(s) = (F \otimes G)(T_s) \leq (F \otimes G)(T_t) = (F * G)(t)$. Hence $F * G$ is decreasing. If $t \in [0, \infty)$ and $t_n \in [0, \infty)$ for $n \in \mathbb{N}_+$ with $t_n \downarrow t$ as $n \rightarrow \infty$ then $T_{t_n} \downarrow T_t$ (in the subset sense) as $n \rightarrow \infty$ so by the continuity property of $F \otimes G$ we have $(F * G)(t_n) = (F \otimes G)(T_{t_n}) \downarrow (F \otimes G)(T_t) = (F * G)(t)$ as $n \rightarrow \infty$. Hence $F * G$ is continuous from the right.

For the commutative property, we have $(F * G)(t) = (F \otimes G)(T_t)$ and $(G * F)(t) = (G \otimes F)(T_t)$. By the symmetry of the triangle T_t with respect to the diagonal $\{(s, s) : s \in [0, \infty)\}$, these are the same.

If F and G are probability distribution functions corresponding to independent random variables X and Y with values in $[0, \infty)$, then $F * G$ is the probability distribution function of $X + Y$. Suppose now that $f : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded and that G and H are distribution functions on $[0, \infty)$. From the previous result, both $(f * G) * H$ and $f * (G * H)$ make sense. Fortunately, they are the same so that convolution is *associative*.

Suppose that $f : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded and that G and H are distribution functions on $[0, \infty)$. Then

$$(f * G) * H = f * (G * H) \quad (15.2.9)$$

Proof

For $t \in [0, \infty)$,

$$[(f * G) * H](t) = \int_0^t (f * G)(t - s) dH(s) = \int_0^t \int_0^{t-s} f(t - s - r) dG(r) dH(s) = [f * (G * H)](t) \quad (15.2.10)$$

Finally, convolution is a *linear operation*. That is, convolution preserves sums and scalar multiples, whenever these make sense.

Suppose that $f, g : [0, \infty) \rightarrow \mathbb{R}$ are locally bounded, H is a distribution function on $[0, \infty)$, and $c \in \mathbb{R}$. Then

1. $(f + g) * H = (f * H) + (g * H)$
2. $(cf) * H = c(f * H)$

Proof

These properties follow easily from linearity properties of the integral.

1. $[(f + g) * H](t) = \int_0^t (f + g)(t - s) dH(s) = \int_0^t f(t - s) dH(s) + \int_0^t g(t - s) dH(s) = (f * H)(t) + (g * H)(t)$
2. $[(cf) * H](t) = \int_0^t cf(t - s) dH(s) = c \int_0^t f(t - s) dH(s) = c(f * H)(t)$

Suppose that $f : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded, G and H are distribution functions on $[0, \infty)$, and that $c \in (0, \infty)$. Then

1. $f * (G + H) = (f * G) + (f * H)$
2. $f * (cG) = c(f * G)$

Proof

These properties also follow from linearity properties of the integral.

1. $[f * (G + H)](t) = \int_0^t f(t - s) d(G + H)(s) = \int_0^t f(t - s) dG(s) + \int_0^t f(t - s) dH(s) = (f * G)(t) + (f * H)(t)$
2. $[f * (cG)](t) = \int_0^t f(t - s) d(cG)(s) = c \int_0^t f(t - s) dG(s) = c(f * G)(t)$

Laplace Transforms

Like convolution, the term *Laplace transform* (named for Pierre Simon Laplace of course) can mean slightly different things in different settings. We start with the usual definition that you may have seen in your study of differential equations or other subjects:

The *Laplace transform* of a function $f : [0, \infty) \rightarrow \mathbb{R}$ is the function ϕ defined as follows, for all $s \in (0, \infty)$ for which the integral exists in \mathbb{R} :

$$\phi(s) = \int_0^\infty e^{-st} f(t) dt \quad (15.2.11)$$

Suppose that f is nonnegative, so that the integral defining the transform exists in $[0, \infty]$ for every $s \in (0, \infty)$. If $\phi(s_0) < \infty$ for some $s_0 \in (0, \infty)$ then $\phi(s) < \infty$ for $s \geq s_0$. The transform of a general function f exists (in \mathbb{R}) if and only if the transform of $|f|$ is finite at s . It follows that if f has a Laplace transform, then the transform ϕ is defined on an interval of the form (a, ∞) for some $a \in (0, \infty)$. The

actual domain is of very little importance; the main point is that the Laplace transform, if it exists, will be defined for all sufficiently large s . Basically, a nonnegative function will fail to have a Laplace transform if it grows at a “hyper-exponential rate” as $t \rightarrow \infty$.

We could generalize the Laplace transform by replacing the Riemann or Lebesgue integral with the integral over a positive measure on $[0, \infty)$.

Suppose that G is a distribution on $[0, \infty)$. The Laplace transform of $f : [0, \infty) \rightarrow \mathbb{R}$ with respect to G is the function given below, defined for all $s \in (0, \infty)$ for which the integral exists in \mathbb{R} :

$$s \mapsto \int_0^\infty e^{-st} f(t) dG(t) \quad (15.2.12)$$

However, as before, if f is nonnegative, then $H(t) = \int_0^t f(x) dG(x)$ for $t \in [0, \infty)$ defines another distribution function, and the previous integral is simply $\int_0^\infty e^{-st} dH(t)$. This motivates the definition for the Laplace transform of a distribution.

The Laplace transform of a distribution F on $[0, \infty)$ is the function Φ defined as follows, for all $s \in (0, \infty)$ for which the integral is finite:

$$\Phi(s) = \int_0^\infty e^{-st} dF(t) \quad (15.2.13)$$

Once again if F has a Laplace transform, then the transform will be defined for all sufficiently large $s \in (0, \infty)$. We will try to be explicit in explaining which of the Laplace transform definitions is being used. For a generic function, the first definition applies, and we will use a lower case Greek letter. If the function is a *distribution function*, either definition makes sense, but it is usually the latter that is appropriate, in which case we use an upper case Greek letter. Fortunately, there is a simple relationship between the two.

Suppose that F is a distribution function on $[0, \infty)$. Let Φ denote the Laplace transform of the *distribution* F and ϕ the Laplace transform of the *function* F . Then $\Phi(s) = s\phi(s)$.

Proof

The main tool is Fubini's theorem (named for Guido Fubini), which allow us to interchange the order of integration for a nonnegative function.

$$\phi(s) = \int_0^\infty e^{-st} F(t) dt = \int_0^\infty e^{-st} \left(\int_0^t dF(x) \right) dt \quad (15.2.14)$$

$$= \int_0^\infty \left(\int_x^\infty e^{-st} dt \right) dF(x) = \int_0^\infty \frac{1}{s} e^{-sx} dF(x) = \frac{1}{s} \Phi(s) \quad (15.2.15)$$

For a probability distribution, there is also a simple relationship between the Laplace transform and the moment generating function.

Suppose that X is a random variable with values in $[0, \infty)$ and with probability distribution function F . The Laplace transform Φ and the moment generating function Γ of the distribution F are given as follows, and so $\Phi(s) = \Gamma(-s)$ for all $s \in (0, \infty)$.

$$\Phi(s) = \mathbb{E}(e^{-sX}) = \int_0^\infty e^{-st} dF(t) \quad (15.2.16)$$

$$\Gamma(s) = \mathbb{E}(e^{sX}) = \int_0^\infty e^{st} dF(t) \quad (15.2.17)$$

In particular, a probability distribution F on $[0, \infty)$ always has a Laplace transform Φ , defined on $(0, \infty)$. Note also that if $F(0) < 1$ (so that X is not deterministically 0), then $\Phi(s) < 1$ for $s \in (0, \infty)$.

Laplace transforms are important for general distributions on $[0, \infty)$ for the same reasons that moment generating functions are important for probability distributions: the transform of a distribution uniquely determines the distribution, and the transform of a convolution is the product of the corresponding transforms (and products are much nicer mathematically than convolutions). The following theorems give the essential properties of Laplace transforms. We assume that the transforms exist, of course, and it should be understood that equations involving transforms hold for sufficiently large $s \in (0, \infty)$.

Suppose that F and G are distributions on $[0, \infty)$ with Laplace transforms Φ and Γ , respectively. If $\Phi(s) = \Gamma(s)$ for s sufficiently large, then $G = F$.

In the case of general functions on $[0, \infty)$, the conclusion is that $f = g$ except perhaps on a subset of $[0, \infty)$ of measure 0. The Laplace transform is a linear operation.

Suppose that $f, g : [0, \infty) \rightarrow \mathbb{R}$ have Laplace transforms ϕ and γ , respectively, and $c \in \mathbb{R}$ then

1. $f + g$ has Laplace transform $\phi + \gamma$
2. cf has Laplace transform $c\phi$

Proof

These properties follow from the linearity of the integral. For s sufficiently large,

1. $\int_0^\infty e^{-st} [f(t) + g(t)] dt = \int_0^\infty e^{-st} f(t) dt + \int_0^\infty e^{-st} g(t) dt = \phi(s) + \gamma(s)$
2. $\int_0^\infty e^{-st} cf(t) dt = c \int_0^\infty e^{-st} f(t) dt = c\phi(s)$

The same properties holds for distributions on $[0, \infty)$ with $c \in (0, \infty)$. Integral transforms have a smoothing effect. Laplace transforms are differentiable, and we can interchange the derivative and integral operators.

Suppose that $f : [0, \infty) \rightarrow \mathbb{R}$ has Laplace transform ϕ . Then ϕ has derivatives of all orders and

$$\phi^{(n)}(s) = \int_0^\infty (-1)^n t^n e^{-st} f(t) dt \quad (15.2.18)$$

Restated, $(-1)^n \phi^{(n)}$ is the Laplace transform of the function $t \mapsto t^n f(t)$. Again, one of the most important properties is that the Laplace transform turns convolution into products.

Suppose that $f : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded with Laplace transform ϕ , and that G is a distribution function on $[0, \infty)$ with Laplace transform Γ . Then $f * G$ has Laplace transform $\phi \cdot \Gamma$.

Proof

By definition, the Laplace transform of $f * G$ is

$$\int_0^\infty e^{-st} (f * G)(t) dt = \int_0^\infty e^{-st} \left(\int_0^t f(t-x) dG(x) \right) dt \quad (15.2.19)$$

Writing $e^{-st} = e^{-s(t-x)} e^{-sx}$ and reversing the order of integration, the last iterated integral can be written as

$$\int_0^\infty e^{-sx} \left(\int_x^\infty e^{-s(t-x)} f(t-x) dt \right) dG(x) \quad (15.2.20)$$

The interchange is justified, once again, by Fubini's theorem, since our functions are integrable (for sufficiently large $s \in (0, \infty)$).

Finally with the substitution $y = t - x$ the last iterated integral can be written as a product

$$\left(\int_0^\infty e^{-sy} f(y) dy \right) \left(\int_0^\infty e^{-sx} dG(x) \right) = \phi(s) \Gamma(s) \quad (15.2.21)$$

If F and G are distributions on $[0, \infty)$, then so is $F * G$. The result above applies, of course, with F and $F * G$ thought of as *functions* and G as a *distribution*, but multiplying through by s and using the theorem [above](#), it's clear that the result is also true with all three as distributions.

Renewal Equations and Their Solutions

Armed with our new analytic machinery, we can return to the study of renewal processes. Thus, suppose that we have a renewal process with interarrival sequence $\mathbf{X} = (X_1, X_2, \dots)$, arrival time sequence $\mathbf{T} = (T_0, T_1, \dots)$, and counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$. As usual, let F denote the common distribution function of the interarrival times, and let M denote the renewal function, so that $M(t) = \mathbb{E}(N_t)$ for $t \in [0, \infty)$. Of course, the probability distribution function F defines a probability measure on $[0, \infty)$, but as noted earlier, M is also a distribution function and so defines a positive measure on $[0, \infty)$. Recall that $F^c = 1 - F$ is the right distribution function (or reliability function) of an interarrival time.

The distributions of the arrival times are the convolution powers of F . That is, $F_n = F^{*n} = F * F * \dots * F$.

Proof

This follows from the definitions: F_n is the distribution function of T_n , and $T_n = \sum_{i=1}^n X_i$. Since \mathbf{X} is an independent, identically distributed sequence, $F_n = F^{*n}$

The next definition is the central one for this section.

Suppose that $a : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded. An integral equation of the form

$$u = a + u * F \quad (15.2.22)$$

for an unknown function $u : [0, \infty) \rightarrow \mathbb{R}$ is called a *renewal equation* for u .

Often $u(t) = \mathbb{E}(U_t)$ where $\{U_t : t \in [0, \infty)\}$ is a random process of interest associated with the renewal process. The renewal equation comes from conditioning on the first arrival time $T_1 = X_1$, and then using the defining property of the renewal process—the fact that the process starts over, interdependently of the past, at the arrival time. Our next important result illustrates this.

Renewal equations for M and F :

1. $M = F + M * F$
2. $F = M - F * M$

Proof

1. We condition on the time of the first arrival X_1 and break the domain of integration $[0, \infty)$ into the two parts $[0, t]$ and (t, ∞) :

$$M(t) = \mathbb{E}(N_t) = \int_0^\infty \mathbb{E}(N_t | X_1 = s) dF(s) = \int_0^t \mathbb{E}(N_t | X_1 = s) dF(s) + \int_t^\infty \mathbb{E}(N_t | X_1 = s) dF(s) \quad (15.2.23)$$

If $s > t$ then $\mathbb{E}(N_t | X_1 = s) = 0$. If $0 \leq s \leq t$, then by the renewal property, $\mathbb{E}(N_t | X_1 = s) = 1 + M(t - s)$. Hence we have

$$M(t) = \int_0^t [1 + M(t - s)] dF(s) = F(t) + (M * F)(t) \quad (15.2.24)$$

2. From (a) and the commutative property of convolution given [above](#) (recall that M is also a distribution function), we have $F = M - M * F = M - F * M$

Thus, the renewal function itself satisfies a renewal equation. Of course, we already have a “formula” for M , namely $M = \sum_{n=1}^\infty F_n$. However, sometimes M can be computed more easily from the renewal equation directly. The next result is the transform version of the previous result:

The distributions F and M have Laplace transforms Φ and Γ , respectively, that are related as follows:

$$\Gamma = \frac{\Phi}{1 - \Phi}, \quad \Phi = \frac{\Gamma}{\Gamma + 1} \quad (15.2.25)$$

Proof from the renewal equation

Taking Laplace transforms through the renewal equation $M = F + M * F$ (and treating all terms as distributions), we have $\Gamma = \Phi + \Gamma\Phi$. Solving for Γ gives the result. Recall that since F is a probability distribution on $[0, \infty)$ with $F(0) < 1$, we know that $0 < \Phi(s) < 1$ for $s \in (0, \infty)$. The second equation follows from the first by simple algebra.

Proof from convolution

Recall that $M = \sum_{n=1}^\infty F^{*n}$. Taking Laplace transforms (again treating all terms as distributions), and using geometric series we have

$$\Gamma = \sum_{n=1}^\infty \Phi^n = \frac{\Phi}{1 - \Phi} \quad (15.2.26)$$

Recall again that $0 < \Phi(s) < 1$ for $s \in (0, \infty)$. Once again, the second equation follows from the first by simple algebra.

In particular, the renewal distribution M always has a Laplace transform. The following theorem gives the fundamental results on the solution of the renewal equation.

Suppose that $a : [0, \infty) \rightarrow \mathbb{R}$ is locally bounded. Then the unique locally bounded solution to the renewal equation $u = a + u * F$ is $u = a + a * M$.

Direct proof

Suppose that $u = a + a * M$. Then $u * F = a * F + a * M * F$. But from the renewal equation for M [above](#), $M * F = M - F$. Hence we have $u * F = a * F + a * (M - F) = a * [F + (M - F)] = a * M$. But $a * M = u - a$ by definition of u , so $u = a + u * F$ and hence u is a solution to the renewal equation. Next since a is locally bounded, so is $u = a + a * M$. Suppose now that v is another locally bounded solution of the integral equation, and let $w = u - v$. Then w is locally bounded and

$w * F = (u * F) - (v * F) = [(u - a) - (v - a) = u - v = w$. Hence $w = w * F_n$ for $n \in \mathbb{N}_+$. Suppose that $|w(s)| \leq D_t$ for $0 \leq s \leq t$. Then $|w(t)| \leq D_t F_n(t)$ for $n \in \mathbb{N}_+$. Since $M(t) = \sum_{n=1}^{\infty} F_n(t) < \infty$ it follows that $F_n(t) \rightarrow 0$ as $n \rightarrow \infty$. Hence $w(t) = 0$ for $t \in [0, \infty)$ and so $u = v$.

Proof from Laplace transforms

Let α and θ denote the Laplace transforms of the functions a and u , respectively, and Φ the Laplace transform of the distribution F . Taking Laplace transforms through the renewal equations gives the simple algebraic equation $\theta = \alpha + \theta\Phi$. Solving give

$$\theta = \frac{\alpha}{1 - \Phi} = \alpha \left(1 + \frac{\Phi}{1 - \Phi} \right) = \alpha + \alpha\Gamma \quad (15.2.27)$$

where $\Gamma = \frac{\Phi}{1 - \Phi}$ is the Laplace transform of the distribution M . Thus θ is the transform of $a + a * M$.

Returning to the renewal equations for M and F [above](#), we now see that the renewal function M completely determines the renewal process: from M we can obtain F , and everything is ultimately constructed from the interarrival times. Of course, this is also clear from the Laplace transform result [above](#) which gives simple algebraic equations for each transform in terms of the other.

The Distribution of the Age Variables

Let's recall the definition of the age variables. A deterministic time $t \in [0, \infty)$ falls in the *random* renewal interval $[T_{N_t}, T_{N_t+1})$. The *current life* (or age) at time t is $C_t = t - T_{N_t}$, the *remaining life* at time t is $R_t = T_{N_t+1} - t$, and the *total life* at time t is $L_t = T_{N_t+1} - T_{N_t}$. In the usual reliability setting, C_t is the age of the device that is in service at time t , while R_t is the time until that device fails, and L_t is the total lifetime of the device.

For $t, y \in [0, \infty)$, let

$$r_y(t) = \mathbb{P}(R_t > y) = \mathbb{P}(N(t, t+y] = 0) \quad (15.2.28)$$

and let $F_y^c(t) = F^c(t+y)$. Note that $y \mapsto r_y(t)$ is the right distribution function of R_t . We will derive and then solve a renewal equation for r_y by conditioning on the time of the first arrival. We can then find integral equations that describe the distribution of the current age and the joint distribution of the current and remaining ages.

For $y \in [0, \infty)$, r_y satisfies the renewal equation $r_y = F_y^c + r_y * F$ and hence for $t \in [0, \infty)$,

$$\mathbb{P}(R_t > y) = F^c(t+y) + \int_0^t F^c(t+y-s) dM(s), \quad y \geq 0 \quad (15.2.29)$$


Proof

As usual, we condition on the time of the first renewal:

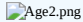
$$\mathbb{P}(R_t > y) = \int_0^\infty \mathbb{P}(R_t > y \mid X_1 = s) dF(s) \quad (15.2.30)$$

We are naturally led to break the domain $[0, \infty)$ of the integral into three parts $[0, t]$, $(t, t+y]$, and $(t+y, \infty)$, which we take one at a time.

Note first that $\mathbb{P}(R_t > y \mid X_1 = s) = \mathbb{P}(R_{t-s} > y)$ for $s \in [0, t]$

The event $R_t > y$ given $X_1 = s$ when $0 \leq s \leq t$


Next note that $\mathbb{P}(R_t > y \mid X_1 = s) = 0$ for $s \in (t, t+y]$

The event $R_t > y$ given $X_1 = s$ when $t < s \leq t+y$


Finally note that $\mathbb{P}(R_t > y \mid X_1 = s) = 1$ for $s \in (t+y, \infty)$

The event $R_t > y$ given $X_1 = s$ when $s > t+y$


Putting the pieces together we have

$$\mathbb{P}(R_t > y) = \int_0^t \mathbb{P}(R_{t-s} > y) dF(s) + \int_t^{t+y} 0 dF(s) + \int_{t+y}^\infty 1 dF(s) \quad (15.2.31)$$

In terms of our function notation, the first integral is $(r_y * F)(t)$, the second integral is 0 of course, and the third integral is $1 - F(t+y) = F_y^c(t)$. Thus the renewal equation is satisfied and the formula for $\mathbb{P}(R_t > y)$ follows the [fundamental theorem](#) on

renewal equations.

We can now describe the distribution of the current age.

For $t \in [0, \infty)$,

$$\mathbb{P}(C_t \geq x) = F^c(t) + \int_0^{t-x} F^c(t-s) dM(s), \quad x \in [0, t] \quad (15.2.32)$$

Proof

This follows from the [previous theorem](#) and the fact that $\mathbb{P}(C_t \geq x) = \mathbb{P}(R_{t-x} > x)$ for $x \in [0, t]$.

Finally we get the joint distribution of the current and remaining ages.

For $t \in [0, \infty)$,

$$\mathbb{P}(C_t \geq x, R_t > y) = F^c(t+y) + \int_0^{t-x} F^c(t+y-s) dM(s), \quad x \in [0, t], y \in [0, \infty) \quad (15.2.33)$$

Proof

Recall that $\mathbb{P}(C_t \geq x, R_t > y) = \mathbb{P}(R_{t-x} > x+y)$. The result now follows from the [result above](#) for the remaining life.

Examples and Special Cases

Uniformly Distributed Interarrivals

Consider the renewal process with interarrival times uniformly distributed on $[0, 1]$. Thus the distribution function of an interarrival time is $F(x) = x$ for $0 \leq x \leq 1$. The renewal function M can be computed from the general [renewal equation for \$M\$](#) by successively solving differential equations. The following exercise give the first two cases.

On the interval $[0, 2]$, show that M is given as follows:

1. $M(t) = e^t - 1$ for $0 \leq t \leq 1$
2. $M(t) = (e^t - 1) - (t-1)e^{t-1}$ for $1 \leq t \leq 2$

Solution




Figure 15.2.2: The graph of M on the interval $[0, 2]$

Show that the Laplace transform Φ of the interarrival distribution F and the Laplace transform Γ of the renewal distribution M are given by

$$\Phi(s) = \frac{1-e^{-s}}{s}, \quad \Gamma(s) = \frac{1-e^{-s}}{s-1+e^{-s}}; \quad s \in (0, \infty) \quad (15.2.34)$$

Solution

First note that

$$\Phi(s) = \int_0^\infty e^{-st} dF(t) = \int_0^1 e^{-st} dt = \frac{1-e^{-s}}{s}, \quad s \in (0, \infty) \quad (15.2.35)$$

The formula for Γ follows from $\Gamma = \Phi/(1-\Phi)$.

Open the renewal experiment and select the uniform interarrival distribution on the interval $[0, 1]$. For each of the following values of the time parameter, run the experiment 1000 times and note the shape and location of the empirical distribution of the counting variable.

1. $t = 5$
2. $t = 10$
3. $t = 15$
4. $t = 20$
5. $t = 25$
6. $t = 30$

The Poisson Process

Recall that the Poisson process has interarrival times that are exponentially distributed with rate parameter $r > 0$. Thus, the interarrival distribution function F is given by $F(x) = 1 - e^{-rx}$ for $x \in [0, \infty)$. The following exercises give alternate proofs of fundamental results obtained in the Introduction.

Show that the renewal function M is given by $M(t) = rt$ for $t \in [0, \infty)$

1. Using the renewal equation
2. Using Laplace transforms

Solution

1. The renewal equation gives

$$M(t) = 1 - e^{-rt} + \int_0^t M(t-s)re^{-rs} ds \quad (15.2.36)$$

Substituting $x = t - s$ in the integral gives

$$M(t) = 1 - e^{-rt} + re^{-rt} \int_0^t M(x)e^{rx} dx \quad (15.2.37)$$

Multiplying through by e^{rt} , differentiating with respect to t , and simplifying gives $M'(t) = r$ for $t \geq 0$. Since $M(0) = 0$, the result follows.

2. The Laplace transform Φ of the distribution F is given by

$$\Phi(s) = \int_0^\infty e^{-st} re^{-rt} dt = \int_0^\infty rne^{-(s+r)t} dt = \frac{r}{r+s}, \quad s \in (0, \infty) \quad (15.2.38)$$

So the Laplace transform Γ of the distribution M is given by

$$\Gamma(s) = \frac{\Phi(s)}{1 - \Phi(s)} = \frac{r}{s}, \quad s \in (0, \infty) \quad (15.2.39)$$

But this is the Laplace transform of the distribution $t \mapsto rt$.

Show that the current and remaining life at time $t \geq 0$ satisfy the following properties:

1. C_t and R_t are independent.
2. R_t has the same distribution as an interarrival time, namely the exponential distribution with rate parameter r .
3. C_t has a truncated exponential distribution with parameters t and r :

$$\mathbb{P}(C_t \geq x) = \begin{cases} e^{-rx}, & 0 \leq x \leq t \\ 0, & x > t \end{cases} \quad (15.2.40)$$

Solution

Recall again that $M(t) = rt$ for $t \in [0, \infty)$. Using the result above on the [joint distribution of the current and remaining life](#), and some standard calculus, we have

$$\mathbb{P}(C_t \geq x, R_t \geq y) = e^{-r(t+y)} + \int_0^{t-x} e^{-r(t+y-s)} r ds = r^{-rx} e^{-ry}, \quad x \in [0, t], y \in [0, \infty) \quad (15.2.41)$$

Letting $y = 0$ gives $\mathbb{P}(C_t \geq x) = e^{-rx}$ for $x \in [0, t]$. Letting $x = 0$ gives $\mathbb{P}(R_t \geq y) = e^{-ry}$ for $y \in [0, \infty)$. But then also $\mathbb{P}(C_t \geq x, R_t \geq y) = \mathbb{P}(C_t \geq x)\mathbb{P}(R_t \geq y)$ for $x \in [0, t]$ and $y \in [0, \infty)$ so the variables are independent.

Bernoulli Trials

Consider the renewal process for which the interarrival times have the geometric distribution with parameter p . Recall that the probability density function is

$$f(n) = (1-p)^{n-1}p, \quad n \in \mathbb{N}_+ \quad (15.2.42)$$

The arrivals are the successes in a sequence of Bernoulli trials. The number of successes Y_n in the first n trials is the counting variable for $n \in \mathbb{N}$. The renewal equations in this section can be used to give alternate proofs of some of the fundamental results in the Introduction.

Show that the renewal function is $M(n) = np$ for $n \in \mathbb{N}$

1. Using the renewal equation
2. Using Laplace transforms

Proof

1. The [renewal equation for \$M\$](#) is

$$M(n) = F(n) + (M * F)(n) = 1 - (1-p)^n + \sum_{k=1}^n M(n-k)p(1-p)^{k-1}, \quad n \in \mathbb{N} \quad (15.2.43)$$

So substituting values of n successively we have

$$M(0) = 1 - (1-p)^0 = 0 \quad (15.2.44)$$

$$M(1) = 1 - (1-p) + M(0)p = p \quad (15.2.45)$$

$$M(2) = 1 - (1-p)^2 + M(1)p + M(0)p(1-p) = 2p \quad (15.2.46)$$

and so forth.

2. The Laplace transform Φ of the distribution F is

$$\Phi(s) = \sum_{n=1}^{\infty} e^{-sn} p(1-p)^{n-1} = \frac{pe^{-s}}{1 - (1-p)e^{-s}}, \quad s \in (0, \infty) \quad (15.2.47)$$

Hence the Laplace transform of the distribution M is

$$\Gamma(s) = \frac{\Phi(s)}{1 - \Phi(s)} = p \frac{e^{-s}}{1 - e^{-s}}, \quad s \in (0, \infty) \quad (15.2.48)$$

But $s \mapsto e^{-s} / (1 - e^{-s})$ is the transform of the distribution $n \mapsto n$ on \mathbb{N} . That is,

$$\sum_{n=1}^{\infty} e^{-sn} \cdot 1 = \frac{e^{-s}}{1 - e^{-s}}, \quad s \in (0, \infty) \quad (15.2.49)$$

Show that the current and remaining life at time $n \in \mathbb{N}$ satisfy the following properties:

1. C_n and R_n are independent.
2. R_n has the same distribution as an interarrival time, namely the geometric distribution with parameter p .
3. C_n has a truncated geometric distribution with parameters n and p :

$$\mathbb{P}(C_n = j) = \begin{cases} p(1-p)^j, & j \in \{0, 1, \dots, n-1\} \\ (1-p)^n, & j = n \end{cases} \quad (15.2.50)$$

Solution

Recall again that $M(n) = np$ for $n \in \mathbb{N}$. Using the result above on the [joint distribution of the current and remaining life](#) and geometric series, we have

$$\mathbb{P}(C_n \geq j, R_n > k) = (1-p)^{n+k} + \sum_{i=1}^{n-j} p(1-p)^{n+k-i} = (1-p)^{j+k}, \quad j \in \{0, 1, \dots, n\}, k \in \mathbb{N} \quad (15.2.51)$$

Letting $k = 0$ gives $\mathbb{P}(C_n \geq j) = (1-p)^j$ for $j \in \{0, 1, \dots, n\}$. Letting $j = 0$ gives $\mathbb{P}(R_n > k) = (1-p)^k$ for $k \in \mathbb{N}$. But then also $\mathbb{P}(C_n \geq j, R_n > k) = \mathbb{P}(C_n \geq j)\mathbb{P}(R_n > k)$ for $j \in \{0, 1, \dots, n\}$ and $k \in \mathbb{N}$ so the variables are independent.

A Gamma Interarrival Distribution

Consider the renewal process whose interarrival distribution F is gamma with shape parameter 2 and rate parameter $r \in (0, \infty)$. Thus

$$F(t) = 1 - (1 + rt)e^{-rt}, \quad t \in [0, \infty) \quad (15.2.52)$$

Recall also that F is the distribution of the sum of two independent random variables, each having the exponential distribution with rate parameter r .

Show that the renewal distribution function M is given by

$$M(t) = -\frac{1}{4} + \frac{1}{2}rt + \frac{1}{4}e^{-2rt}, \quad t \in [0, \infty) \quad (15.2.53)$$

Solution

The exponential distribution with rate parameter r has Laplace transform $s \mapsto r/(r+s)$ and hence the Laplace transform Φ of the interarrival distribution F is given by

$$\Phi(s) = \left(\frac{r}{r+s} \right)^2 \quad (15.2.54)$$

So the Laplace transform Γ of the distribution M is

$$\Gamma(s) = \frac{\Phi(s)}{1 - \Phi(s)} = \frac{r^2}{s(s+2r)} \quad (15.2.55)$$

Using a partial fraction decomposition,

$$\Gamma(s) = \frac{r}{2s} - \frac{r}{2(s+2r)} = \frac{1}{2} \frac{r}{s} - \frac{1}{4} \frac{2r}{s+2r} \quad (15.2.56)$$

But the r/s is the Laplace transform of the distribution rt and $2r/(s+2r)$ is the Laplace transform of the distribution $1 - e^{-2rt}$ (the exponential distribution with parameter $2r$).

Note that $M(t) \approx -\frac{1}{4} + \frac{1}{2}rt$ as $t \rightarrow \infty$.




Figure 15.2.3: The graph of M on the interval $[0, 5]$ when $r = 1$

Open the renewal experiment and select the gamma interarrival distribution with shape parameter $k = 2$ and scale parameter $b = 1$ (so the rate parameter $r = \frac{1}{b}$ is also 1). For each of the following values of the time parameter, run the experiment 1000 times and note the shape and location of the empirical distribution of the counting variable.

1. $t = 5$
2. $t = 10$
3. $t = 15$
4. $t = 20$
5. $t = 25$
6. $t = 30$

15.3: Renewal Limit Theorems

We start with a renewal process as constructed in the introduction. Thus, $\mathbf{X} = (X_1, X_2, \dots)$ is the sequence of interarrival times. These are independent, identically distributed, nonnegative variables with common distribution function F (satisfying $F(0) < 1$) and common mean μ . When $\mu = \infty$, we let $1/\mu = 0$. When $\mu < \infty$, we let σ denote the common standard deviation. Recall also that $F^c = 1 - F$ is the right distribution function (or reliability function). Then, $\mathbf{T} = (T_0, T_1, \dots)$ is the arrival time sequence, where $T_0 = 0$ and

$$T_n = \sum_{i=1}^n X_i \quad (15.3.1)$$

is the time of the n th arrival for $n \in \mathbb{N}_+$. Finally, $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is the counting process, where for $t \in [0, \infty)$,

$$N_t = \sum_{n=1}^{\infty} \mathbf{1}(T_n \leq t) \quad (15.3.2)$$

is the number of arrivals in $[0, t]$. The renewal function M is defined by $M(t) = \mathbb{E}(N_t)$ for $t \in [0, \infty)$.

We noted earlier that the arrival time process and the counting process are inverses, in a sense. The arrival time process is the partial sum process for a sequence of independent, identically distributed variables. Thus, it seems reasonable that the fundamental limit theorems for partial sum processes (the law of large numbers and the central limit theorem), should have analogs for the counting process. That is indeed the case, and the purpose of this section is to explore the limiting behavior of renewal processes. The main results that we will study, known appropriately enough as *renewal theorems*, are important for other stochastic processes, particularly Markov chains.

Basic Theory

The Law of Large Numbers

Our first result is a strong law of large numbers for the renewal counting process, which comes as you might guess, from the law of large numbers for the sequence of arrival times.

If $\mu < \infty$ then $N_t/t \rightarrow 1/\mu$ as $t \rightarrow \infty$ with probability 1.

Proof

Recall that $T_{N_t} \leq t < T_{N_t+1}$ for $t > 0$. Hence, if $N_t > 0$,

$$\frac{T_{N_t}}{N_t} \leq \frac{t}{N_t} < \frac{T_{N_t+1}}{N_t} \quad (15.3.3)$$

Recall that $N_t \rightarrow \infty$ as $t \rightarrow \infty$ with probability 1. Recall also that by the strong law of large numbers that $T_n/n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1. It follows that $T_{N_t}/N_t \rightarrow \mu$ as $t \rightarrow \infty$ with probability 1. Also, $(N_t + 1)/N_t \rightarrow 1$ as $t \rightarrow \infty$ with probability 1. Therefore

$$\frac{T_{N_t+1}}{N_t} = \frac{T_{N_t+1}}{N_t+1} \frac{N_t+1}{N_t} \rightarrow \mu \quad (15.3.4)$$

as $t \rightarrow \infty$ with probability 1. Hence by the squeeze theorem for limits, $t/N_t \rightarrow \mu$ as $t \rightarrow \infty$ with probability 1.

Thus, $1/\mu$ is the limiting average rate of arrivals per unit time.

Open the renewal experiment and set $t = 50$. For a variety of interarrival distributions, run the simulation 1000 times and note how the empirical distribution is concentrated near t/μ .

The Central Limit Theorem

Our next goal is to show that the counting variable N_t is asymptotically normal.

Suppose that μ and σ are finite, and let

$$Z_t = \frac{N_t - t/\mu}{\sigma\sqrt{t/\mu^3}}, \quad t > 0 \quad (15.3.5)$$

The distribution of Z_t converges to the standard normal distribution as $t \rightarrow \infty$.

Proof

For $n \in \mathbb{N}_+$, let

$$W_n = \frac{T_n - n\mu}{\sigma\sqrt{n}} \quad (15.3.6)$$

The distribution of W_n converges to the standard normal distribution as $n \rightarrow \infty$, by the ordinary central limit theorem. Next, for $z \in \mathbb{R}$, $\mathbb{P}(Z_t \leq z) = \mathbb{P}(T_{n(z,t)} > t)$ where $n(z, t) = \lfloor t/\mu + z\sigma\sqrt{t/\mu^3} \rfloor$. Also, $\mathbb{P}(Z_t \leq z) = \mathbb{P}[W_{n(z,t)} > w(z, t)]$ where

$$w(z, t) = -\frac{z}{\sqrt{1 + z\sigma/\sqrt{t/\mu}}} \quad (15.3.7)$$

But $n(z, t) \rightarrow \infty$ as $t \rightarrow \infty$ and $w(z, t) \rightarrow -z$ as $t \rightarrow \infty$. Recall that $1 - \Phi(-z) = \Phi(z)$, where as usual, Φ is the standard normal distribution function. Thus, we conclude that $\mathbb{P}(Z_t \leq z) \rightarrow \Phi(z)$ as $t \rightarrow \infty$.

Open the renewal experiment and set $t = 50$. For a variety of interarrival distributions, run the simulation 1000 times and note the “normal” shape of the empirical distribution. Compare the empirical mean and standard deviation to t/μ and $\sigma\sqrt{t/\mu^3}$, respectively

The Elementary Renewal Theorem

The *elementary renewal theorem* states that the basic limit in the law of large numbers [above](#) holds in mean, as well as with probability 1. That is, the limiting mean average rate of arrivals is $1/\mu$. The elementary renewal theorem is of fundamental importance in the study of the limiting behavior of Markov chains, but the proof is not as easy as one might hope. In particular, recall that convergence with probability 1 does not imply convergence in mean, so the elementary renewal theorem does not follow from the law of large numbers.

$M(t)/t \rightarrow 1/\mu$ as $t \rightarrow \infty$.

Proof

We first show that $\liminf_{t \rightarrow \infty} M(t)/t \geq 1/\mu$. Note first that this result is trivial if $\mu = \infty$, so assume that $\mu < \infty$. Next, recall that $N_t + 1$ is a stopping time for the sequence of interarrival times \mathbf{X} . Recall also that $T_{N_t+1} > t$ for $t > 0$. From Wald's equation it follows that

$$\mathbb{E}(T_{N_t+1}) = \mathbb{E}(N_t + 1)\mu = [M(t) + 1]\mu > t \quad (15.3.8)$$

Therefore $M(t)/t > 1/\mu - 1/t$ for $t > 0$. Hence $\liminf_{t \rightarrow \infty} M(t)/t \geq 1/\mu$.

Next we show that $\limsup_{t \rightarrow \infty} M(t)/t \leq 1/\mu$. For this part of the proof, we need to *truncate* the arrival times, and use the basic comparison method. For $a > 0$, let

$$X_{a,i} = \begin{cases} X_i, & X_i \leq a \\ a, & X_i > a \end{cases} \quad (15.3.9)$$

and consider the renewal process with the sequence of interarrival times $\mathbf{X}_a = (X_{a,1}, X_{a,2}, \dots)$. We will use the standard notation developed in the introductory section. First note that $T_{a, N_{a,t}+1} \leq t + a$ for $t > 0$ and $a > 0$. From Wald's equation again, it follows that $[M_a(t) + 1]\mu_a \leq t + a$. Therefore

$$\frac{M_a(t)}{t} \leq \left(\frac{1}{\mu_a} + \frac{a}{t\mu_a} \right) - \frac{1}{t}, \quad a, t > 0 \quad (15.3.10)$$

But $M(t) \leq M_a(t)$ for $t > 0$ and $a > 0$ and therefore

$$\frac{M(t)}{t} \leq \left(\frac{1}{\mu_a} + \frac{a}{t\mu_a} \right) - \frac{1}{t}, \quad a, t > 0 \quad (15.3.11)$$

Hence $\limsup_{t \rightarrow \infty} M(t)/t \leq 1/\mu_a$ for $a > 0$. Finally, $\mu_a \rightarrow \mu$ as $a \rightarrow \infty$ by the monotone convergence theorem, so it follows that $\limsup_{t \rightarrow \infty} M(t)/t \leq 1/\mu$

Open the renewal experiment and set $t = 50$. For a variety of interarrival distributions, run the experiment 1000 times and once again compare the empirical mean and standard deviation to t/μ and $\sigma\sqrt{t/\mu^3}$, respectively.

The Renewal Theorem

The *renewal theorem* states that the expected number of renewals in an interval is asymptotically proportional to the length of the interval; the proportionality constant is $1/\mu$. The precise statement is different, depending on whether the renewal process is arithmetic or not. Recall that for an arithmetic renewal process, the interarrival times take values in a set of the form $\{nd : n \in \mathbb{N}\}$ for some $d \in (0, \infty)$, and the largest such d is the *span* of the distribution.

For $h > 0$, $M(t, t+h] \rightarrow \frac{h}{\mu}$ as $t \rightarrow \infty$ in each of the following cases:

1. The renewal process is non-arithmetic
2. The renewal process is arithmetic with span d , and h is a multiple of d

The renewal theorem is also known as *Blackwell's theorem* in honor of David Blackwell. The final limit theorem we will study is the most useful, but before we can state the theorem, we need to define and study the class of functions to which it applies.

Direct Riemann Integration

Recall that in the ordinary theory of Riemann integration, the integral of a function on the interval $[0, t]$ exists if the upper and lower Riemann sums converge to a common number as the partition is refined. Then, the integral of the function on $[0, \infty)$ is defined to be the limit of the integral on $[0, t]$, as $t \rightarrow \infty$. For our new definition, a function is said to be *directly* Riemann integrable if the lower and upper Riemann sums on the entire unbounded interval $[0, \infty)$ converge to a common number as the partition is refined, a more restrictive definition than the usual one.

Suppose that $g : [0, \infty) \rightarrow [0, \infty)$. For $h \in [0, \infty)$ and $k \in \mathbb{N}$, let $m_k(g, h) = \inf\{g(t) : t \in [kh, (k+1)h)\}$ and $M_k(g, h) = \sup\{g(t) : t \in [kh, (k+1)h)\}$. The *lower* and *upper* Riemann sums of g on $[0, \infty)$ corresponding to h are

$$L_g(h) = h \sum_{k=0}^{\infty} m_k(g, h), \quad U_g(h) = h \sum_{k=0}^{\infty} M_k(g, h) \quad (15.3.12)$$

The sums exist in $[0, \infty]$ and satisfy the following properties:

1. $L_g(h) \leq U_g(h)$ for $h > 0$
2. $L_g(h)$ increases as h decreases
3. $U_g(h)$ decreases as h decreases

It follows that $\lim_{h \downarrow 0} L_g(h)$ and $\lim_{h \downarrow 0} U_g(h)$ exist in $[0, \infty]$ and $\lim_{h \downarrow 0} L_g(h) \leq \lim_{h \downarrow 0} U_g(h)$. Naturally, the case where the limits are finite and agree is what we're after.

A function $g : [0, \infty) \rightarrow [0, \infty)$ is *directly Riemann integrable* if $U_g(h) < \infty$ for every $h > 0$ and

$$\lim_{h \downarrow 0} L_g(h) = \lim_{h \downarrow 0} U_g(h) \quad (15.3.13)$$

The common value is $\int_0^{\infty} g(t) dt$.

Ordinary Riemann integrability on $[0, \infty)$ allows functions that are unbounded and oscillate wildly as $t \rightarrow \infty$, and these are the types of functions that we want to exclude for the renewal theorems. The following result connects ordinary Riemann integrability with direct Riemann integrability.

If $g: [0, \infty) \rightarrow [0, \infty)$ is integrable (in the ordinary Riemann sense) on $[0, t]$ for every $t \in [0, \infty)$ and if $U_g(h) < \infty$ for some $h \in (0, \infty)$ then g is directly Riemann integrable.

Here is a simple and useful class of functions that are directly Riemann integrable.

Suppose that $g: [0, \infty) \rightarrow [0, \infty)$ is decreasing with $\int_0^\infty g(t) dt < \infty$. Then g is directly Riemann integrable.

The Key Renewal Theorems

The *key renewal theorem* is an integral version of the renewal theorem, and is the most useful of the various limit theorems.

Suppose that the renewal process is non-arithmetic and that $g: [0, \infty) \rightarrow [0, \infty)$ is directly Riemann integrable. Then

$$(g * M)(t) = \int_0^t g(t-s) dM(s) \rightarrow \frac{1}{\mu} \int_0^\infty g(x) dx \text{ as } t \rightarrow \infty \quad (15.3.14)$$

Connections

Our next goal is to see how the various renewal theorems relate.

The renewal theorem implies the elementary renewal theorem:

Proof

Let $a_n = M(n, n+1]$ for $n \in \mathbb{N}$. From the renewal theorem, $a_n \rightarrow 1/\mu$ as $n \rightarrow \infty$. Therefore $\frac{1}{n} \sum_{k=0}^{n-1} a_k \rightarrow \frac{1}{\mu}$ as $n \rightarrow \infty$. It follows that $M(n)/n \rightarrow 1/\mu$ as $n \rightarrow \infty$. But the renewal function is increasing so for $t > 0$,

$$\frac{\lfloor t \rfloor}{t} \frac{M(\lfloor t \rfloor)}{\lfloor t \rfloor} \leq \frac{M(t)}{t} \leq \frac{\lceil t \rceil}{t} \frac{M(\lceil t \rceil)}{\lceil t \rceil} \quad (15.3.15)$$

From the squeeze theorem for limits it follows that $M(t)/t \rightarrow 1/\mu$ as $t \rightarrow \infty$.

Conversely, the elementary renewal theorem *almost* implies the renewal theorem.

Proof

Assume that $g(x) = \lim_{t \rightarrow \infty} [M(t+x) - M(t)]$ exists for each $x > 0$. (This assumption is the reason that the proof is incomplete.) Note that

$$M(t+x+y) - M(t) = [M(t+x+y) - M(t+x)] + [M(t+x) - M(t)] \quad (15.3.16)$$

Let $t \rightarrow \infty$ to conclude that $g(x+u) = g(x) + g(y)$ for all $x \geq 0$ and $y \geq 0$. It follows that g is increasing and $g(x) = cx$ for $x \geq 0$ where c is a constant. Exactly as in the proof of the previous theorem, it follows that $M(n)/n \rightarrow c$ as $n \rightarrow \infty$. From the elementary renewal theorem, we can conclude that $c = 1/\mu$.

The key renewal theorem implies the renewal theorem

Proof

This result follows by applying the key renewal theorem to the function $g_h(x) = \mathbf{1}(0 \leq x \leq h)$ where $h > 0$.

Conversely, the renewal theorem implies the key renewal theorem.

The Age Processes

The key renewal theorem can be used to find the limiting distributions of the current and remaining age. Recall that for $t \in [0, \infty)$ the current life at time t is $C_t = t - T_{N_t}$ and the remaining life at time t is $R_t = T_{N_t+1} - t$.

If the renewal process is non-arithmetic, then

$$\mathbb{P}(R_t > x) \rightarrow \frac{1}{\mu} \int_x^\infty F^c(y) dy \text{ as } t \rightarrow \infty, \quad x \in [0, \infty) \quad (15.3.17)$$

Proof

Recall that

$$\mathbb{P}(R_t > x) = F^c(t+x) + \int_0^t F^c(t+x-s) dM(s), \quad x \in [0, \infty) \quad (15.3.18)$$

But $F^c(t+x) \rightarrow 0$ as $t \rightarrow \infty$, and by the key renewal theorem, the integral converges to $\frac{1}{\mu} \int_0^\infty F^c(x+y) dy$. Finally a change of variables in the limiting integral gives the result.

If the renewal process is aperiodic, then

$$\mathbb{P}(R_t > x) \rightarrow \frac{1}{\mu} \int_x^\infty F^c(y) dy \text{ as } t \rightarrow \infty, \quad x \in [0, \infty) \quad (15.3.19)$$

Proof

Recall that, since the renewal process is aperiodic,

$$\mathbb{P}(C_t > x) = F^c(t) + \int_0^{t-x} F^c(t-s) dM(s), \quad x \in [0, t] \quad (15.3.20)$$

Again, $F^c(t) \rightarrow 0$ as $t \rightarrow \infty$. The change of variables $u = t - x$ changes the integral into $\int_0^u F^c(u+x-s) dM(s)$. By the key renewal theorem, this integral converges $\frac{1}{\mu} \int_0^\infty F^c(y+x) dy = \int_x^\infty F^c(y+x) dy$.

The current and remaining life have the same limiting distribution. In particular,

$$\lim_{t \rightarrow \infty} \mathbb{P}(C_t \leq x) = \lim_{t \rightarrow \infty} \mathbb{P}(R_t \leq x) = \frac{1}{\mu} \int_0^x F^c(y) dy, \quad x \in [0, \infty) \quad (15.3.21)$$

Proof

By the previous two theorems, the limiting right distribution functions of R_t and C_t are the same. The ordinary (left) limiting distribution function is

$$1 - \frac{1}{\mu} \int_x^\infty F^c(y) dy = \frac{1}{\mu} \left(\mu - \int_x^\infty F^c(y) dy \right) \quad (15.3.22)$$

But recall that $\mu = \int_0^\infty F^c(y) dy$ so the result follows since $\int_0^\infty F^c(y) dy - \int_x^\infty F^c(y) dy = \int_0^x F^c(y) dy$

The fact that the current and remaining age processes have the same limiting distribution may seem surprising at first, but there is a simple intuitive explanation. After a long period of time, the renewal process looks just about the same backward in time as forward in time. But reversing the direction of time reverses the roles of current and remaining age.

Examples and Special Cases

The Poisson Process

Recall that the Poisson process, the most important of all renewal processes, has interarrival times that are exponentially distributed with rate parameter $r > 0$. Thus, the interarrival distribution function is $F(x) = 1 - e^{-rx}$ for $x \geq 0$ and the mean interarrival time is $\mu = 1/r$.

Verify each of the following directly:

1. The law of large numbers for the counting process.
2. The central limit theorem for the counting process.
3. The elementary renewal theorem.
4. The renewal theorem.

Bernoulli Trials

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of Bernoulli trials with success parameter $p \in (0, 1)$. Recall that \mathbf{X} is a sequence of independent, identically distributed indicator variables with $p = \mathbb{P}(X = 1)$. We have studied a number of random processes

derived from \mathbf{X} :

Random processes associated with Bernoulli trials.

1. $\mathbf{Y} = (Y_0, Y_1, \dots)$ where Y_n the number of successes in the first n trials. The sequence \mathbf{Y} is the partial sum process associated with \mathbf{X} . The variable Y_n has the binomial distribution with parameters n and p .
2. $\mathbf{U} = (U_1, U_2, \dots)$ where U_n the number of trials needed to go from success number $n - 1$ to success number n . These are independent variables, each having the geometric distribution on \mathbb{N}_+ with parameter p .
3. $\mathbf{V} = (V_0, V_1, \dots)$ where V_n is the trial number of success n . The sequence \mathbf{V} is the partial sum process associated with \mathbf{U} . The variable V_n has the negative binomial distribution with parameters n and p .

Consider the renewal process with interarrival sequence \mathbf{U} . Thus, $\mu = 1/p$ is the mean interarrival time, and \mathbf{Y} is the counting process. Verify each of the following directly:

1. The law of large numbers for the counting process.
2. The central limit theorem for the counting process.
3. The elementary renewal theorem.

Consider the renewal process with interarrival sequence \mathbf{X} . Thus, the mean interarrival time is $\mu = p$ and the number of arrivals in the interval $[0, n]$ is $V_{n+1} - 1$ for $n \in \mathbb{N}$. Verify each of the following directly:

1. The law of large numbers for the counting process.
2. The central limit theorem for the counting process.
3. The elementary renewal theorem.

This page titled [15.3: Renewal Limit Theorems](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

15.4: Delayed Renewal Processes

Basic Theory

Preliminaries

A *delayed renewal process* is just like an ordinary renewal process, except that the first arrival time is allowed to have a different distribution than the other interarrival times. Delayed renewal processes arise naturally in applications and are also found embedded in other random processes. For example, in a Markov chain (which we study in the next chapter), visits to a fixed state, *starting in that state* form the random times of an ordinary renewal process. But visits to a fixed state, *starting in another state* form a delayed renewal process.

Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent variables taking values in $[0, \infty)$, with (X_2, X_3, \dots) identically distributed. Suppose also that $\mathbb{P}(X_i > 0) > 0$ for $i \in \mathbb{N}_+$. The stochastic process with \mathbf{X} as the sequence of interarrival times is a *delayed renewal process*.

As before, the actual arrival times are the partial sums of \mathbf{X} . Thus let

$$T_n = \sum_{i=1}^n X_i \quad (15.4.1)$$

so that $T_0 = 0$ and T_n is the time of the n th arrival for $n \in \{1, 2, \dots\}$. Also as before, N_t is the number of arrivals in $[0, t]$ (not counting T_0):

$$N_t = \sum_{n=1}^{\infty} \mathbf{1}(T_n \leq t) = \max\{n \in \mathbb{N} : T_n \leq t\} \quad (15.4.2)$$

If we restart the clock at time $T_1 = X_1$, we have an ordinary renewal process with interarrival sequence (X_2, X_3, \dots) . We use some of the standard notation developed in the Introduction for this renewal process. In particular, F denotes the common distribution function and μ the common mean of X_i for $i \in \{2, 3, \dots\}$. Similarly $F_n = F^{*n}$ denotes the distribution function of the sum of n independent variables with distribution function F , and M denotes the renewal function:

$$M(t) = \sum_{n=1}^{\infty} F_n(t), \quad t \in [0, \infty) \quad (15.4.3)$$

On the other hand, we will let G denote the distribution function of X_1 (the special interarrival time, different from the rest), and we will let G_n denote the distribution function of T_n for $n \in \mathbb{N}_+$. As usual, $F^c = 1 - F$ and $G^c = 1 - G$ are the corresponding right-tail distribution functions.

$$G_n = G * F_{n-1} = F_{n-1} * G \quad \text{for } n \in \mathbb{N}_+.$$

Proof

The follows from the fact that T_n is the sum of n independent random variables; the first has distribution function G and the remaining $n - 1$ have distribution function F .

Finally, we will let U denote the renewal function for the delayed renewal process. Thus, $U(t) = \mathbb{E}(N_t)$ is the expected number of arrivals in $[0, t]$ for $t \in [0, \infty)$.

The delayed renewal function satisfies

$$U(t) = \sum_{n=1}^{\infty} G_n(t), \quad t \in [0, \infty) \quad (15.4.4)$$

Proof

The proof is just as before.

$$U(t) = \mathbb{E}(N_t) = \mathbb{E} \left(\sum_{n=1}^{\infty} \mathbf{1}(T_n \leq t) \right) = \sum_{n=1}^{\infty} \mathbb{P}(T_n \leq t) = \sum_{n=1}^{\infty} G_n(t) \quad (15.4.5)$$

The delayed renewal function U satisfies the equation $U = G + M * G$; that is,

$$U(t) = G(t) + \int_0^t M(t-s) dG(s), \quad t \in [0, \infty) \quad (15.4.6)$$

Proof

The proof follows from conditioning on the time of the first arrival $T_1 = X_1$. Note first that $\mathbb{E}(N_t | X_1 = s) = 0$ if $s > t$ and $\mathbb{E}(N_t | X_1 = s) = 1 + M(t-s)$ if $0 \leq s \leq t$. Hence

$$U(t) = \int_0^{\infty} \mathbb{E}(N_t | X_1 = s) dG(s) = \int_0^t [1 + M(t-s)] dG(s) = G(t) + \int_0^t M(t-s) dG(s) \quad (15.4.7)$$

The delayed renewal function U satisfies the renewal equation $U = G + U * F$; that is,

$$U(t) = G(t) + \int_0^t U(t-s) dF(s), \quad t \in [0, \infty) \quad (15.4.8)$$

Proof

Note that

$$U = \sum_{n=1}^{\infty} G_n = G + \sum_{n=2}^{\infty} G_n = G + \sum_{n=2}^{\infty} (G_{n-1} * F) = G + \left(\sum_{n=1}^{\infty} G_n \right) * F = G + U * F \quad (15.4.9)$$

Asymptotic Behavior

In a delayed renewal process only the first arrival time is changed. Thus, it's not surprising that the asymptotic behavior of a delayed renewal process is the same as the asymptotic behavior of the corresponding regular renewal process. Our first result is the *strong law of large numbers* for the delayed renewal process.

$N_t/t \rightarrow 1/\mu$ as $t \rightarrow \infty$ with probability 1.

Proof

We will show that $T_n/n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1. Then, the proof is exactly like the proof of the law of large numbers for a regular renewal process. For $n \in \{2, 3, \dots\}$,

$$\frac{T_n}{n} = \frac{X_1}{n} + \frac{n-1}{n} \frac{1}{n-1} \sum_{i=2}^n X_i \quad (15.4.10)$$

But $\frac{X_1}{n} \rightarrow 0$ as $n \rightarrow \infty$ with probability 1; of course $\frac{n-1}{n} \rightarrow 1$ as $n \rightarrow \infty$; and $\frac{1}{n-1} \sum_{i=2}^n X_i \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1 by the ordinary strong law of large numbers.

Our next result is the *elementary renewal theorem* for the delayed renewal process.

$U(t)/t \rightarrow 1/\mu$ as $t \rightarrow \infty$.

Next we have the *renewal theorem* for the delayed renewal process, also known as *Blackwell's theorem*, named for David Blackwell.

For $h > 0$, $U(t, t+h] = U(t+h) - U(t) \rightarrow h/\mu$ as $t \rightarrow \infty$ in each of the following cases:

1. F is non-arithmetic
2. F is arithmetic with span $d \in (0, \infty)$, and h is a multiple of d .

Finally we have the *key renewal theorem* for the delayed renewal process.

Suppose that the renewal process is non-arithmetic and that $g: [0, \infty) \rightarrow [0, \infty)$ is directly Riemann integrable. Then

$$(g * U)(t) = \int_0^t g(t-s) dU(s) \rightarrow \frac{1}{\mu} \int_0^\infty g(x) dx \text{ as } t \rightarrow \infty \quad (15.4.11)$$

Stationary Point Processes

Recall that a *point process* is a stochastic process that models a discrete set of random points in a measure space $(S, \mathcal{S}, \lambda)$. Often, of course, $S \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}_+$ and λ is the corresponding n -dimensional Lebesgue measure. The special cases $S = \mathbb{N}$ with counting measure and $S = [0, \infty)$ with length measure are of particular interest, in part because renewal and delayed renewal processes give rise to point processes in these spaces.

For a general point process on S , we use our standard notation and denote the number of random points $A \in \mathcal{S}$ by $N(A)$. There are a couple of natural properties that a point process may have. In particular, the process is said to be *stationary* if $\lambda(A) = \lambda(B)$ implies that $N(A)$ and $N(B)$ have the same distribution for $A, B \in \mathcal{S}$. In $[0, \infty)$ the term *stationary increments* is often used, because the stationarity property means that for $s, t \in [0, \infty)$, the distribution of $N(s, s+t] = N_{s+t} - N_s$ depends only on t .

Consider now a regular renewal process. We showed earlier that the asymptotic distributions of the current life and remaining life are the same. Intuitively, after a very long period of time, the renewal process looks pretty much the same forward in time or backward in time. This suggests that if we make the renewal process into a delayed renewal process by giving the first arrival time this asymptotic distribution, then the resulting point process will be stationary. This is indeed the case. Consider the setting and notation of the preliminary subsection above.

For the delayed renewal process, the point process \mathbf{N} is stationary if and only if the initial arrival time has distribution function

$$G(t) = \frac{1}{\mu} \int_0^t F^c(s) ds, \quad t \in [0, \infty) \quad (15.4.12)$$

in which case the renewal function is $U(t) = t/\mu$ for $t \in [0, \infty)$.

Proof

Suppose first that \mathbf{N} has stationary increments. In particular, this means that the arrival times have continuous distributions. For $s, t \in [0, \infty)$,

$$U(s+t) = \mathbb{E}(N_{s+t}) = \mathbb{E}[(N_{s+t} - N_t) + N_t] = \mathbb{E}(N_{s+t} - N_t) + \mathbb{E}(N_t) = U(s) + U(t) \quad (15.4.13)$$

A theorem from analysis states that the only increasing solutions to such a functional equation are linear functions, and hence $U(t) = ct$ for some positive constant c . Substituting m_d into the [renewal equation](#) above gives

$$ct = G(t) + \int_0^t c(t-s) dF(s) = G(t) + ctF(t) - \int_0^t cs dF(s) \quad (15.4.14)$$

Integrating by parts in the last integral and simplifying gives

$$G(t) = c \int_0^t F^c(s) ds \quad (15.4.15)$$

Finally, if we let $t \rightarrow \infty$, the left side converges to 1 and the right side to $c\mu$, so $c = 1/\mu$. Thus G has the form given in the statement of the theorem and $U(t) = t/\mu$ for $t \in [0, \infty)$.

Conversely, suppose that G has the form given in the theorem. Note that this is a continuous distribution with density function $t \mapsto F^c(t)/\mu$. Substituting into the [renewal equation](#) above, it follows that the renewal density U' satisfies

$$U' = \frac{1}{\mu} F^c + \frac{1}{\mu} F^c * \sum_{n=1}^{\infty} F_n = \frac{1}{\mu} \quad (15.4.16)$$

Hence $U(t) = t/\mu$ for $t \geq 0$. Next, the process \mathbf{N} has stationary increments if and only if the remaining life R_t at time t has distribution function G for each t . Arguing just as in Section 2, we have

$$\mathbb{P}(R_t > y) = G^c(t+y) + \int_0^t F^c(t+y-s) dm_d(s), \quad y \geq 0 \quad (15.4.17)$$

But $G^c(t+y) = \frac{1}{\mu} \int_{t+y}^{\infty} F^c(u) du$ and $dU(s) = \frac{1}{\mu} ds$, so substituting into the last displayed equation and using a simple substitution in the integral gives

$$\mathbb{P}(R_t > y) = \frac{1}{\mu} \int_{t+y}^{\infty} F^c(u) du + \frac{1}{\mu} \int_y^{t+y} F^c(u) du = \frac{1}{\mu} \int_y^{\infty} F^c(u) du \quad (15.4.18)$$

Examples and Applications

Patterns in Multinomial Trials

Suppose that $\mathbf{L} = (L_1, L_2, \dots)$ is a sequence of independent, identically distributed random variables taking values in a finite set S , so that \mathbf{L} is a sequence of multinomial trials. Let f denote the common probability density function so that for a generic trial variable L , we have $f(a) = \mathbb{P}(L = a)$ for $a \in S$. We assume that all outcomes in S are actually possible, so $f(a) > 0$ for $a \in S$.

In this section, we interpret S as an *alphabet*, and we write the sequence of variables in concatenation form, $\mathbf{L} = L_1 L_2 \dots$ rather than standard sequence form. Thus the sequence is an infinite string of letters from our alphabet S . We are interested in the repeated occurrence of a particular finite substring of letters (that is, a “word” or “pattern”) in the infinite sequence.

So, fix a word \mathbf{a} (again, a finite string of elements of S), and consider the successive random trial numbers (T_1, T_2, \dots) where the word \mathbf{a} is completed in \mathbf{L} . Since the sequence \mathbf{L} is independent and identically distributed, it seems reasonable that these variables are the arrival times of a renewal process. However there is a slight complication. An example may help.

Suppose that \mathbf{L} is a sequence of Bernoulli trials (so $S = \{0, 1\}$). Suppose that the outcome of \mathbf{L} is

$$101100101010001101000110 \dots \quad (15.4.19)$$

1. For the word $\mathbf{a} = 001$ note that $T_1 = 7, T_2 = 15, T_3 = 22$
2. For the word $\mathbf{b} = 010$, note that $T_1 = 8, T_2 = 10, T_3 = 12, T_4 = 19$

In this example, you probably noted an important difference between the two words. For \mathbf{b} , a *suffix* of the word (a proper substring at the end) is also a *prefix* of the word (a proper substring at the beginning). Word \mathbf{a} does not have this property. So, once we “arrive” at \mathbf{b} , there are ways to get to \mathbf{b} again (taking advantage of the suffix-prefix) that do not exist starting from the beginning of the trials. On the other hand, once we arrive at \mathbf{a} , arriving at \mathbf{a} again is just like with a new sequence of trials. Thus we are lead to the following definition.

Suppose that \mathbf{a} is a finite word from the alphabet S . If no proper suffix of \mathbf{a} is also a prefix, then \mathbf{a} is *simple*. Otherwise, \mathbf{a} is *compound*.

Returning to the general setting, let $T_0 = 0$ and then let $X_n = T_n - T_{n-1}$ for $n \in \mathbb{N}_+$. For $k \in \mathbb{N}$, let $N_k = \sum_{n=1}^{\infty} \mathbf{1}(T_n \leq k)$. For occurrences of the word \mathbf{a} , $\mathbf{X} = (X_1, X_2, \dots)$ is the sequence of interarrival times, $\mathbf{T} = (T_0, T_1, \dots)$ is the sequence of arrival times, and $\mathbf{N} = \{N_k : k \in \mathbb{N}\}$ is the counting process. If \mathbf{a} is simple, these form an ordinary renewal process. If \mathbf{a} is compound, they form a delayed renewal process, since X_1 will have a different distribution than (X_2, X_3, \dots) . Since the structure of a delayed renewal process subsumes that of an ordinary renewal process, we will work with the notation above for the delayed process. In particular, let U denote the renewal function. Everything in this paragraph depends on the word \mathbf{a} of course, but we have suppressed this in the notation.

Suppose $\mathbf{a} = a_1 a_2 \dots a_k$, where $a_i \in S$ for each $i \in \{1, 2, \dots, k\}$, so that \mathbf{a} is a word of length k . Note that X_1 takes values in $\{k, k+1, \dots\}$. If \mathbf{a} is simple, this applies to the other interarrival times as well. If \mathbf{a} is compound, the situation is more complicated X_2, X_3, \dots will have some minimum value $j < k$, but the possible values are positive integers, of course, and include $\{k+1, k+2, \dots\}$. In any case, the renewal process is arithmetic with span 1. Expanding the definition of the probability density function f , let

$$f(\mathbf{a}) = \prod_{i=1}^k f(a_i) \quad (15.4.20)$$

so that $f(\mathbf{a})$ is the probability of forming \mathbf{a} with k consecutive trials. Let $\mu(\mathbf{a})$ denote the common mean of X_n for $n \in \{2, 3, \dots\}$, so $\mu(\mathbf{a})$ is the mean number of trials between occurrences of \mathbf{a} . Let $\nu(\mathbf{a}) = \mathbb{E}(X_1)$, so that $\nu(\mathbf{a})$ is the mean time

number of trials until \mathbf{a} occurs for the first time. Our first result is an elegant connection between $\mu(\mathbf{a})$ and $f(\mathbf{a})$, which has a wonderfully simple proof from renewal theory.

If \mathbf{a} is a word in S then

$$\mu(\mathbf{a}) = \frac{1}{f(\mathbf{a})} \quad (15.4.21)$$

Proof

Suppose that \mathbf{a} has length k , and consider the discrete interval $(n, n+k] = \{n+1, n+2, \dots, n+k\}$. By the [renewal theorem](#), $U(n, n+k] \rightarrow 1/\mu(\mathbf{a})$ as $n \rightarrow \infty$. But $N(n, n+k]$, the number of times that \mathbf{a} occurs in the interval, is either 1 or 0. Hence $U(n, n+k] = f(\mathbf{a})$ for any n .

Our next goal is to compute $\nu(\mathbf{a})$ in the case that \mathbf{a} is a compound word.

Suppose that \mathbf{a} is a compound word, and that \mathbf{b} is the largest word that is a proper suffix and prefix of \mathbf{a} . Then

$$\nu(\mathbf{a}) = \nu(\mathbf{b}) + \mu(\mathbf{a}) = \nu(\mathbf{b}) + \frac{1}{f(\mathbf{a})} \quad (15.4.22)$$

Proof

Since \mathbf{b} is the largest prefix-suffix, the expected number of trials to go from \mathbf{b} to \mathbf{a} is the same as the expected number of trials to go from \mathbf{a} to \mathbf{a} , namely $\mu(\mathbf{a})$. (Note that the paths from \mathbf{b} to \mathbf{a} are the same as the paths from \mathbf{a} to \mathbf{a} .) But to form the word \mathbf{a} initially, the word \mathbf{b} must be formed first, so this result follows from the additivity of expected value and the previous result.

By repeated use of the last result, we can compute the expected number of trials needed to form any compound word.

Consider Bernoulli trials with success probability $p \in (0, 1)$, and let $q = 1 - p$. For each of the following strings, find the expected number of trials between occurrences and the expected number of trials to the first occurrence.

1. $\mathbf{a} = 001$
2. $\mathbf{b} = 010$
3. $\mathbf{c} = 1011011$
4. $\mathbf{d} = 11 \cdots 1$ (k times)

Answer

1. $\mu(\mathbf{a}) = \nu(\mathbf{a}) = \frac{1}{pq^2}$
2. $\mu(\mathbf{b}) = \frac{1}{pq^2}$, $\nu(\mathbf{b}) = \frac{1}{q} + \frac{1}{pq^2}$
3. $\mu(\mathbf{c}) = \frac{1}{p^5q^2}$, $\nu(\mathbf{c}) = \frac{1}{p} + \frac{1}{p^3q} + \frac{1}{p^5q^2}$
4. $\mu(\mathbf{d}) = \frac{1}{p^k}$, $\nu(\mathbf{d}) = \sum_{i=1}^k \frac{1}{p^i}$

Recall that an *ace-six flat die* is a six-sided die for which faces 1 and 6 have probability $\frac{1}{4}$ each while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each. Ace-six flat dice are sometimes used by gamblers to cheat.

Suppose that an ace-six flat die is thrown repeatedly. Find the expected number of throws until the pattern 6165616 first occurs.

Solution

From our main theorem,

$$\begin{aligned} \nu(6165616) &= \frac{1}{f(6165616)} + \nu(616) = \frac{1}{f(6165616)} + \frac{1}{f(616)} + \nu(6) \\ &= \frac{1}{f(6165616)} + \frac{1}{f(616)} + \frac{1}{f(6)} = \frac{1}{(1/4)^6(1/8)} + \frac{1}{(1/4)^3} + \frac{1}{1/4} = 32836 \end{aligned}$$

Suppose that a monkey types randomly on a keyboard that has the 26 lower-case letter keys and the space key (so 27 keys). Find the expected number of keystrokes until the monkey produces each of the following phrases:

1. *it was the best of times*
2. *to be or not to be*

Proof

1. $27^{24} \approx 2.258 \times 10^{34}$
2. $27^5 + 27^{18} \approx 5.815 \times 10^{25}$

This page titled [15.4: Delayed Renewal Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

15.5: Alternating Renewal Processes

Basic Theory

Preliminaries

An *alternating renewal process* models a system that, over time, alternates between two states, which we denote by 1 and 0 (so the system starts in state 1). Generically, we can imagine a device that, over time, alternates between *on* and *off* states. Specializing further, suppose that a device operates until it fails, and then is replaced with an identical device, which in turn operates until failure and is replaced, and so forth. In this setting, the times that the device is functioning correspond to the *on* state, while the replacement times correspond to the *off* state. (The device might actually be *repaired* rather than *replaced*, as long as the repair returns the device to pristine, new condition.) The basic assumption is that the pairs of random times successively spent in the two states form an independent, identically distributed sequence. Clearly the model of a system alternating between two states is basic and important, but moreover, such alternating processes are often found embedded in other stochastic processes.

Let's set up the mathematical notation. Let $\mathbf{U} = (U_1, U_2, \dots)$ denote the successive lengths of time that the system is in state 1, and let $\mathbf{V} = (V_1, V_2, \dots)$ the successive lengths of time that the system is in state 0. So to be clear, the system starts in state 1 and remains in that state for a period of time U_1 , then goes to state 0 and stays in this state for a period of time V_1 , then back to state 1 for a period of time U_2 , and so forth. Our basic assumption is that $\mathbf{W} = ((U_1, V_1), (U_2, V_2), \dots)$ is an independent, identically distributed sequence. It follows that \mathbf{U} and \mathbf{V} each are independent, identically distributed sequences, but \mathbf{U} and \mathbf{V} might well be dependent. In fact, V_n might be a function of U_n for $n \in \mathbb{N}_+$. Let $\mu = \mathbb{E}(U)$ denote the mean of a generic time period U in state 1 and let $\nu = \mathbb{E}(V)$ denote the mean of a generic time period V in state 0. Let G denote the distribution function of a time period U in state 1, and as usual, let $G^c = 1 - G$ denote the right distribution function (or reliability function) of U .

Clearly it's natural to consider returns to state 1 as the arrivals in a renewal process. Thus, let $X_n = U_n + V_n$ for $n \in \mathbb{N}_+$ and consider the renewal process with interarrival times $\mathbf{X} = (X_1, X_2, \dots)$. Clearly this makes sense, since \mathbf{X} is an independent, identically distributed sequence of nonnegative variables. For the most part, we will use our usual notation for a renewal process, so the common distribution function of $X_n = U_n + V_n$ is denoted by F , the arrival time process is $\mathbf{T} = (T_0, T_1, \dots)$, the counting process is $\{N_t : t \in [0, \infty)\}$, and the renewal function is M . But note that the mean interarrival time is now $\mu + \nu$.

The renewal process associated with $\mathbf{W} = ((U_1, V_1), (U_2, V_2), \dots)$ as constructed above is known as an *alternating renewal process*.

The State Process

Our interest is the state I_t of the system at time $t \in [0, \infty)$, so $\mathbf{I} = \{I_t : t \in [0, \infty)\}$ is a stochastic process with state space $\{0, 1\}$. Clearly the stochastic processes \mathbf{W} and \mathbf{I} are equivalent in the sense that we can recover one from the other. Let $p(t) = \mathbb{P}(I_t = 1)$, the probability that the device is on at time $t \in [0, \infty)$. Our first main result is a renewal equation for the function p .

The function p satisfies the renewal equation $p = G^c + p * F$ and hence $p = G^c + G^c * M$.

Proof

By now, the approach should be clear—we're going to condition on the first arrival X_1 :

$$\mathbb{P}(I_t = 1) = \mathbb{P}(I_t = 1, X_1 > t) + \mathbb{P}(I_t = 1, X_1 \leq t) = \mathbb{P}(I_t = 1, X_1 > t) + \int_0^t \mathbb{P}(I_t = 1 \mid X_1 = s) dF(s) \quad (15.5.1)$$

But $\{I_t = 1, X_1 > t\} = \{U_1 > t\}$ so $\mathbb{P}(I_t = 1, X_1 > t) = \mathbb{P}(U_1 > t) = G^c(t)$. By the fundamental renewal property (the process restarts, independently of the past, at each arrival) $\mathbb{P}(I_t = 1 \mid X_1 = s) = p(t - s)$ for $s \leq t$. Hence we have

$$p(t) = G^c(t) + \int_0^t p(t - s) dF(s), \quad t \in [0, \infty) \quad (15.5.2)$$

or equivalently, $p = G^c + p * F$. By the fundamental theorem on renewal equations, the solution is $p = G^c + G^c * M$, so $p(t) = G^c(t) + \int_0^t G^c(t - s) dM(s)$, $t \in [0, \infty)$

We can now apply the key renewal theorem to get the asymptotic behavior of p .

If the renewal process is non-arithmetic, then

$$p(t) \rightarrow \frac{\mu}{\mu + \nu} \text{ as } t \rightarrow \infty \quad (15.5.3)$$

Proof

From the result [above](#), $p = G^c + G^c * m$. First, $G^c(t) \rightarrow 0$ as $t \rightarrow \infty$ as a basic property of the right distribution function. Next, by the key renewal theorem,

$$(G^c * M)(t) \rightarrow \frac{1}{\mu + \nu} \int_0^\infty G^c(s) ds \text{ as } t \rightarrow \infty \quad (15.5.4)$$

But by another basic property of the right distribution function, $\int_0^\infty G^c(s) ds = \mu$.

Thus, the limiting probability that the system is on is simply the ratio of the mean of an *on* period to the mean of an *on-off* period. It follows, of course, that

$$\mathbb{P}(I_t = 0) = 1 - p(t) \rightarrow \frac{\nu}{\mu + \nu} \text{ as } t \rightarrow \infty \quad (15.5.5)$$

so in particular, the fact that the system starts in the on state makes no difference in the limit. We will return to the asymptotic behavior of the alternating renewal process in the next section on renewal reward processes.

Applications and Special Cases

With a clever definition of *on* and *off*, many stochastic processes can be turned into alternating renewal processes, leading in turn to interesting limits, via the basic limit theorem above.

Age Processes

The last remark applies in particular to the age processes of a standard renewal process. So, suppose that we have a renewal process with interarrival sequence \mathbf{X} , arrival sequence \mathbf{T} , and counting process \mathbf{N} . As usual, let μ denote the mean and F the probability distribution function of an interarrival time, and let $F^c = 1 - F$ denote the right distribution function (or reliability function).

For $t \in [0, \infty)$, recall that the *current life*, *remaining life* and *total life* at time t are

$$C_t = t - T_{N_t}, \quad R_t = T_{N_t+1} - t, \quad L_t = C_t + R_t = T_{N_t+1} - T_{N_t} = X_{N_t+1} \quad (15.5.6)$$

respectively. In the usual terminology of reliability, C_t is the age of the device in service at time t , R_t is the time remaining until this device fails, and L_t is total life of the device. We will use limit theorem [above](#) to derive the limiting distributions these age processes. The limiting distributions were obtained earlier, in the section on renewal limit theorems, by a direct application of the key renewal theorem. So the results are not new, but the method of proof is interesting.

If the renewal process is non-arithmetic then

$$\lim_{t \rightarrow \infty} \mathbb{P}(C_t \leq x) = \lim_{t \rightarrow \infty} \mathbb{P}(R_t \leq x) = \frac{1}{\mu} \int_0^x F^c(y) dy, \quad x \in [0, \infty) \quad (15.5.7)$$

Proof

Fix $x \in [0, \infty)$. For the current life limit, define the *on* period corresponding to the interarrival time X_n to be $U_n = \min\{X_n, x\}$ for $n \in \mathbb{N}_+$, so that the *off* period is $V_n = X_n - \min\{X_n, x\}$. Note that the system is on at time $t \in [0, \infty)$ if and only if $C_t \leq x$, and hence $p(t) = \mathbb{P}(C_t \leq x)$. It follows from the limit theorem [above](#) that

$$\mathbb{P}(C_t \leq x) \rightarrow \frac{\mathbb{E}[\min\{X, x\}]}{\mu} \text{ as } t \rightarrow \infty \quad (15.5.8)$$

where X is a generic interarrival time. But

$$G^c(y) = \mathbb{P}[\min\{X, x\} > y] = \mathbb{P}(X > y) \mathbf{1}(x > y) = F^c(y) \mathbf{1}(x > y), \quad y \in [0, \infty) \quad (15.5.9)$$

Hence

$$\mathbb{E}[\min(X, x)] = \int_0^\infty G^c(y) dy = \int_0^x F^c(y) dy \quad (15.5.10)$$

For the remaining life limit we reverse the on-off periods. Thus, define the *on* period corresponding to the interarrival time X_n to be $U_n = X_n - \min\{X_n, x\}$ for $n \in \mathbb{N}_+$, so that the *off* period is $V_n = \min\{X_n, x\}$. Note that the system is off at time t if and only if $R_t \leq x$, and hence $1 - p(t) = \mathbb{P}(R_t \leq x)$. From the limit theorem above,

$$\mathbb{P}(R_t \leq x) \rightarrow \frac{\mathbb{E}[\min\{X, x\}]}{\mu} \text{ as } t \rightarrow \infty \quad (15.5.11)$$

As we have noted before, the fact that the limiting distributions are the same is not surprising after a little thought. After a long time, the renewal process looks the same forward and backward in time, and reversing the arrow of time reverses the roles of current and remaining time.

If the renewal process is non-arithmetic then

$$\lim_{t \rightarrow \infty} \mathbb{P}(L_t \leq x) = \frac{1}{\mu} \int_0^x y dF(y), \quad x \in [0, \infty) \quad (15.5.12)$$

Proof

Fix $x \in [0, \infty)$. For $n \in \mathbb{N}_+$, define the *on* period associated with interarrival time X_n by $U_n = X_n \mathbf{1}(X_n > x)$. Of course, the *off* period corresponding to X_n is $V_n = X_n - U_n$. Thus, each renewal period is either totally on or totally off, depending on whether or not the interarrival time is greater than x . Note that the system is on at time $t \in [0, \infty)$ if and only if $L_t > x$, so from the basic limit theorem [above](#),

$$\mathbb{P}(L_t > x) \rightarrow \frac{1}{\mu} \mathbb{E}[X \mathbf{1}(X > x)] \quad (15.5.13)$$

where X is a generic interarrival time. But

$$\mathbb{E}[X \mathbf{1}(X > x)] = \int_x^\infty y dF(y) \quad (15.5.14)$$

This page titled [15.5: Alternating Renewal Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

15.6: Renewal Reward Processes

Basic Theory

Preliminaries

In a *renewal reward process*, each interarrival time is associated with a random variable that is generically thought of as the *reward* associated with that interarrival time. Our interest is in the process that gives the total reward up to time t . So let's set up the usual notation. Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ are the interarrival times of a renewal process, so that \mathbf{X} is a sequence of independent, identically distributed, nonnegative variables with common distribution function F and mean μ . As usual, we assume that $F(0) < 1$ so that the interarrival times are not deterministically 0, and in this section we also assume that $\mu < \infty$. Let

$$T_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (15.6.1)$$

so that T_n is the time of the n th arrival for $n \in \mathbb{N}_+$ and $\mathbf{T} = (T_0, T_1, \dots)$ is the arrival time sequence. Finally, Let

$$N_t = \sum_{n=1}^{\infty} \mathbf{1}(T_n \leq t), \quad t \in [0, \infty) \quad (15.6.2)$$

so that N_t is the number of arrivals in $[0, t]$ and $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is the counting process. As usual, let $M(t) = \mathbb{E}(N_t)$ for $t \in [0, \infty)$ so that M is the renewal function.

Suppose now that $\mathbf{Y} = (Y_1, Y_2, \dots)$ is a sequence of real-valued random variables, where Y_n is thought of as the reward associated with the interarrival time X_n . However, the term *reward* should be interpreted generically since Y_n might actually be a *cost* or some other value associated with the interarrival time, and in any event, may take negative as well as positive values. Our basic assumption is that the interarrival time and reward pairs $\mathbf{Z} = ((X_1, Y_1), (X_2, Y_2), \dots)$ form an independent and identically distributed sequence. Recall that this implies that \mathbf{X} is an IID sequence, as required by the definition of the renewal process, and that \mathbf{Y} is also an IID sequence. But \mathbf{X} and \mathbf{Y} might well be dependent, and in fact Y_n might be a function of X_n for $n \in \mathbb{N}_+$. Let $\nu = \mathbb{E}(Y)$ denote the mean of a generic reward Y , which we assume exists in \mathbb{R} .

The stochastic process $\mathbf{R} = \{R_t : t \in [0, \infty)\}$ defined by

$$R_t = \sum_{i=1}^{N_t} Y_i, \quad t \in [0, \infty) \quad (15.6.3)$$

is the *reward renewal process* associated with \mathbf{Z} . The function r given by $r(t) = \mathbb{E}(R_t)$ for $t \in [0, \infty)$ is the *reward function*.

As promised, R_t is the total reward up to time $t \in [0, \infty)$. Here are some typical examples:

- The arrivals are customers at a store. Each customer spends a random amount of money.
- The arrivals are visits to a website. Each visitor spends a random amount of time at the site.
- The arrivals are failure times of a complex system. Each failure requires a random repair time.
- The arrivals are earthquakes at a particular location. Each earthquake has a random severity, a measure of the energy released.

So R_t is a random sum of random variables for each $t \in [0, \infty)$. In the special case that \mathbf{Y} and \mathbf{X} independent, the distribution of R_t is known as a *compound distribution*, based on the distribution of N_t and the distribution of a generic reward Y . Specializing further, if the renewal process is Poisson and is independent of \mathbf{Y} , the process \mathbf{R} is a compound Poisson process.

Note that a renewal reward process generalizes an ordinary renewal process. Specifically, if $Y_n = 1$ for each $n \in \mathbb{N}_+$, then $R_t = N_t$ for $t \in [0, \infty)$, so that the reward process simply reduces to the counting process, and then r reduces to the renewal function M .

The Renewal Reward Theorem

For $t \in (0, \infty)$, the average reward on the interval $[0, t]$ is R_t/t , and the expected average reward on that interval is $r(t)/t$. The fundamental theorem on renewal reward processes gives the asymptotic behavior of these averages.

The renewal reward theorem

1. $R_t/t \rightarrow \nu/\mu$ as $t \rightarrow \infty$ with probability 1.
2. $r(t)/t \rightarrow \nu/\mu$ as $t \rightarrow \infty$

Proof

1. Note that

$$\frac{R_t}{t} = \frac{R_t}{N_t} \frac{N_t}{t} \quad (15.6.4)$$

But by the ordinary strong law of large numbers for the IID sequence \mathbf{Y} ,

$$\frac{1}{n} \sum_{i=1}^n Y_i \rightarrow \nu \quad (15.6.5)$$

as $n \rightarrow \infty$ with probability 1. Recall also that $N_t \rightarrow \infty$ as $t \rightarrow \infty$ with probability 1. Hence it follows that

$$\frac{R_t}{N_t} = \frac{1}{N_t} \sum_{i=1}^{N_t} Y_i \rightarrow \nu \quad (15.6.6)$$

as $t \rightarrow \infty$ with probability 1. From the law of large numbers for the renewal process, we know that $N_t/t \rightarrow 1/\mu$ as $t \rightarrow \infty$ with probability 1.

2. Note first that

$$R_t = \sum_{i=1}^{N_t} Y_i = \sum_{i=1}^{N_t+1} Y_i - Y_{N(t)+1} \quad (15.6.7)$$

Next Recall that $N_t + 1$ is a stopping time for the sequence of interarrival times \mathbf{X} for $t \in (0, \infty)$, and hence is also a stopping time for the sequence of interarrival time, reward pairs \mathbf{Z} . (If a random time is a stopping time for a filtration, then it's a stopping time for any larger filtration.) By Wald's equation,

$$\mathbb{E} \left(\sum_{i=1}^{N_t+1} Y_i \right) = \nu \mathbb{E}(N_t + 1) = \nu[M(t) + 1] = \nu M(t) + \nu \quad (15.6.8)$$

By the elementary renewal theorem,

$$\frac{\nu M(t) + \nu}{t} = \nu \frac{M(t)}{t} + \frac{\nu}{t} \rightarrow \frac{\nu}{\mu} \text{ as } t \rightarrow \infty \quad (15.6.9)$$

Thus returning to the first displayed equation above, it remains to show that

$$\frac{\mathbb{E}[Y_{N_t+1}]}{t} \rightarrow 0 \text{ as } t \rightarrow \infty \quad (15.6.10)$$

Let $u(t) = \mathbb{E}[Y_{N_t+1}]$ for $t \in [0, \infty)$. Taking cases for the first arrival time X_1 we have

$$u(t) = \mathbb{E}[Y_{N_t+1} \mathbf{1}(X_1 > t)] + \mathbb{E}[Y_{N_t+1} \mathbf{1}(X_1 \leq t)] \quad (15.6.11)$$

But $X_1 > t$ if and only if $N_t = 0$ so the first term is $\mathbb{E}[Y_1 \mathbf{1}(X_1 > t)]$, which we will note by $a(t)$. We have assumed that the expected reward ν exists in \mathbb{R} . Hence $|a(t)| \leq \mathbb{E}(|Y_1|) < \infty$ so that a is bounded, and $a(t) \rightarrow 0$ as $t \rightarrow \infty$. For the second term, if the first arrival occurs at time $s \in [0, t]$, then the renewal process restarts, independently of the past, so

$$\mathbb{E}[Y_{N_t+1} \mathbf{1}(X_1 \leq t)] = \int_0^t u(t-s) dF(s), \quad t \in [0, \infty) \quad (15.6.12)$$

It follows that u satisfies the renewal equation $u = a + u * F$. By the fundamental theorem on renewal equations, the solution is $u = a + a * M$. Now, fix $\epsilon > 0$. There exists $T \in (0, \infty)$ such that $|a(t)| < \epsilon$ for $t > T$. So for $t > T$,

$$\begin{aligned} \left| \frac{u(t)}{t} \right| &\leq \frac{1}{t} \left[|a(t)| + \int_0^{t-T} |a(t-s)| dM(s) + \int_{t-M}^t |a(t-s)| dM(s) \right] \\ &\leq \frac{1}{t} [\epsilon + \epsilon M(t-T) + \mathbb{E}|Y_1| [M(t) - M(t-T)]] \end{aligned}$$

Using the elementary renewal theorem again, the last expression converges to ϵ/μ as $t \rightarrow \infty$. Since $\epsilon > 0$ is arbitrary, it follows that $u(t)/t \rightarrow 0$ as $t \rightarrow \infty$.

Part (a) generalizes the law of large numbers and part (b) generalizes elementary renewal theorem, for a basic renewal process. Once again, if $Y_n = 1$ for each n , then (a) becomes $N_t/t \rightarrow 1/\mu$ as $t \rightarrow \infty$ and (b) becomes $M(t)/t \rightarrow 1/\mu$ as $t \rightarrow \infty$. It's not surprising then that these two theorems play a fundamental role in the proof of the renewal reward theorem.

General Reward Processes

The renewal reward process $\mathbf{R} = \{R_t : t \in [0, \infty)\}$ [above](#) is constant, taking the value $\sum_{i=1}^n Y_i$, on the renewal interval $[T_n, T_{n+1})$ for each $n \in \mathbb{N}$. Effectively, the rewards are received discretely: Y_1 at time T_1 , an additional Y_2 at time T_2 , and so forth. It's possible to modify the construction so the rewards accrue continuously in time or in a mixed discrete/continuous manner. Here is a simple set of conditions for a general reward process.

Suppose again that $\mathbf{Z} = ((X_1, Y_1), (X_2, Y_2), \dots)$ is the sequence of interarrival times and rewards. A stochastic process $\mathbf{V} = \{V_t : t \in [0, \infty)\}$ (on our underlying probability space) is a *reward process* associated with \mathbf{Z} if the following conditions hold:

1. $V_{T_n} = \sum_{i=1}^n Y_i$ for $n \in \mathbb{N}$
2. V_t is between V_{T_n} and $V_{T_{n+1}}$ for $t \in (T_n, T_{n+1})$ and $n \in \mathbb{N}$

In the continuous case, with nonnegative rewards (the most important case), the reward process will typically have the following form:

Suppose that the rewards are nonnegative and that $\mathbf{U} = \{U_t : t \in [0, \infty)\}$ is a nonnegative stochastic process (on our underlying probability space) with

1. $t \mapsto U_t$ piecewise continuous
2. $\int_{T_n}^{T_{n+1}} U_t dt = Y_{n+1}$ for $n \in \mathbb{N}$

Let $V_t = \int_0^t U_s ds$ for $t \in [0, \infty)$. Then $\mathbf{V} = \{V_t : t \in [0, \infty)\}$ is a reward process associated with \mathbf{Z} .

Proof

By the additivity of the integral and (b), $V_{T_n} = \sum_{i=1}^n Y_i$ for $n \in \mathbb{N}$. Since \mathbf{U} is nonnegative, \mathbf{V} is increasing, so $V_{T_n} \leq V_t \leq V_{T_{n+1}}$ for $t \in (T_n, T_{n+1})$

Thus in this special case, the rewards are being accrued continuously and U_t is the *rate* at which the reward is being accrued at time t . So \mathbf{U} plays the role of a reward density process. For a general reward process, the basic renewal reward theorem still holds.

Suppose that $\mathbf{V} = \{V_t : t \in [0, \infty)\}$ is a reward process associated with $\mathbf{Z} = ((X_1, Y_1), (X_2, Y_2), \dots)$, and let $v(t) = \mathbb{E}(V_t)$ for $t \in [0, \infty)$ be the corresponding reward function.

1. $V_t/t \rightarrow \nu/\mu$ as $t \rightarrow \infty$ with probability 1.
2. $v(t)/t \rightarrow \nu/\mu$ as $t \rightarrow \infty$.

Proof

Suppose first that the reward variables Y are nonnegative. Then

$$\frac{R_t}{t} \leq \frac{V_t}{t} \leq \frac{R_t}{t} + \frac{Y_{N_t+1}}{t} \quad (15.6.13)$$

From the proof of the renewal reward theorem [above](#), $R_t/t \rightarrow \nu/\mu$ as $t \rightarrow \infty$ with probability 1, and $Y_{N_t+1}/t \rightarrow 0$ as $t \rightarrow \infty$ with probability 1. Hence (a) holds. Taking expected values,

$$\frac{r(t)}{t} \leq \frac{v(t)}{t} \leq \frac{r(t)}{t} + \frac{\mathbb{E}(Y_{N_t+1})}{t} \quad (15.6.14)$$

But again from the renewal reward theorem above, $r(t)/t \rightarrow \nu/\mu$ as $t \rightarrow \infty$ and $E(Y_{N_t+1})/t \rightarrow 0$ as $t \rightarrow \infty$. Hence (b) holds. A similar argument works if the reward variables are negative. If the reward variables take positive and negative values, we split the variables into positive and negative parts in the usual way.

Here is the corollary for a continuous reward process.

Suppose that the rewards are positive, and consider the continuous reward process with density process $U = \{U_t : t \in [0, \infty)\}$ as [above](#). Let $u(t) = \mathbb{E}(U_t)$ for $t \in [0, \infty)$. Then

1. $\frac{1}{t} \int_0^t U_s ds \rightarrow \frac{\nu}{\mu}$ as $t \rightarrow \infty$ with probability 1
2. $\frac{1}{t} \int_0^t u(s) ds \rightarrow \frac{\nu}{\mu}$ as $t \rightarrow \infty$

Special Cases and Applications

With a clever choice of the “rewards”, many interesting renewal processes can be turned into renewal reward processes, leading in turn to interesting limits via the renewal reward theorem.

Alternating Renewal Processes

Recall that in an alternating renewal process, a system alternates between *on* and *off* states (starting in the *on* state). If we let $U = (U_1, U_2, \dots)$ be the lengths of the successive time periods in which the system is on, and $V = (V_1, V_2, \dots)$ the lengths of the successive time periods in which the system is off, then the basic assumptions are that $((U_1, V_1), (U_2, V_2), \dots)$ is an independent, identically distributed sequence, and that the variables $X_n = U_n + V_n$ for $n \in \mathbb{N}_+$ form the interarrival times of a standard renewal process. Let $\mu = \mathbb{E}(U)$ denote the mean of a time period that the device is on, and $\nu = \mathbb{E}(V)$ the mean of a time period that the device is off. Recall that I_t denotes the state (1 or 0) of the system at time $t \in [0, \infty)$, so that $I = \{I_t : t \in [0, \infty)\}$ is the state process. The state probability function p is given by $p(t) = \mathbb{P}(I_t = 1)$ for $t \in [0, \infty)$.

Limits for the alternating renewal process.

1. $\frac{1}{t} \int_0^t I_s ds \rightarrow \frac{\mu}{\mu+\nu}$ as $t \rightarrow \infty$ with probability 1
2. $\frac{1}{t} \int_0^t p(s) ds \rightarrow \frac{\mu}{\mu+\nu}$ as $t \rightarrow \infty$

Proof

Consider the renewal reward process where the reward associated with the interarrival time X_n is U_n , the *on* period for that renewal period. The rewards U_n are nonnegative and clearly $\int_{T_n}^{T_{n+1}} I_s ds = U_{n+1}$. So $t \mapsto \int_0^t I_s ds$ for $t \in [0, \infty)$ defines a continuous reward process of the form given [above](#). Parts (a) and (b) follow directly from the reward renewal theorem [above](#).

Thus, the asymptotic average time that the device is on, and the asymptotic mean average time that the device is on, are both simply the ratio of the mean of an on period to the mean of an on-off period. In our previous study of alternating renewal processes, the fundamental result was that in the non-arithmetic case, $p(t) \rightarrow \mu/(\mu + \nu)$ as $t \rightarrow \infty$. This result implies part (b) in the theorem above.

Age Processes

Renewal reward processes can be used to derive some asymptotic results for the age processes of a standard renewal process. So, suppose that we have a renewal process with interarrival sequence \mathbf{X} , arrival sequence \mathbf{T} , and counting process \mathbf{N} . As usual, let $\mu = \mathbb{E}(X)$ denote the mean of an interarrival time, but now we will also need $\nu = \mathbb{E}(X^2)$, the second moment. We assume that both moments are finite.

For $t \in [0, \infty)$, recall that the *current life*, *remaining life* and *total life* at time t are

$$A_t = t - T_{N_t}, \quad B_t = T_{N_t+1} - t, \quad L_t = A_t + B_t = T_{N_t+1} - T_{N_t} = X_{N_t+1} \quad (15.6.15)$$

respectively. In the usual terminology of reliability, A_t is the age of the device in service at time t , B_t is the time remaining until this device fails, and L_t is total life of the device. (To avoid notational clashes, we are using different notation than in past

sections.) Let $a(t) = \mathbb{E}(A_t)$, $b(t) = \mathbb{E}(B_t)$, and $l(t) = \mathbb{E}(L_t)$ for $t \in [0, \infty)$, the corresponding mean functions. To derive our asymptotic results, we simply use the current life and the remaining life as reward densities (or rates) in a renewal reward process.

Limits for the current life process.

1. $\frac{1}{t} \int_0^t A_s ds \rightarrow \frac{\nu}{2\mu}$ as $t \rightarrow \infty$ with probability 1
2. $\frac{1}{t} \int_0^t a(s) ds \rightarrow \frac{\nu}{2\mu}$ as $t \rightarrow \infty$

Proof

Consider the renewal reward process where the reward associated with the interarrival time X_n is $\frac{1}{2}X_n^2$ for $n \in \mathbb{N}$. The process $t \mapsto \int_0^t A_s ds$ for $t \in [0, \infty)$ is a continuous reward process for this sequence of rewards, as defined [above](#). To see this, note that for $t \in [T_n, T_{n+1})$, we have $A_t = t - T_n$, so with a change of variables and noting that $T_{n+1} = T_n + X_{n+1}$ we have

$$\int_{T_n}^{T_{n+1}} A_t dt = \int_0^{X_{n+1}} s ds = \frac{1}{2} X_{n+1}^2 \quad (15.6.16)$$

The results now follow from the renewal reward theorem [above](#).

Limits for the remaining life process.

1. $\frac{1}{t} \int_0^t B_s ds \rightarrow \frac{\nu}{2\mu}$ as $t \rightarrow \infty$ with probability 1
2. $\frac{1}{t} \int_0^t b(s) ds \rightarrow \frac{\nu}{2\mu}$ as $t \rightarrow \infty$

Proof

Consider again the renewal reward process where the reward associated with the interarrival time X_n is $\frac{1}{2}X_n^2$ for $n \in \mathbb{N}$. The process $t \mapsto \int_0^t B_s ds$ for $t \in [0, \infty)$ is a continuous reward process for this sequence of rewards, as defined [above](#). To see this, note that for $t \in [T_n, T_{n+1})$, we have $B_t = T_{n+1} - t$, so once again with a change of variables and noting that $T_{n+1} = T_n + X_{n+1}$ we have

$$\int_{T_n}^{T_{n+1}} B_t dt = \int_0^{X_{n+1}} s ds = \frac{1}{2} X_{n+1}^2 \quad (15.6.17)$$

The results now follow from the renewal reward theorem [above](#).

With a little thought, it's not surprising that the limits for the current life and remaining life processes are the same. After a long period of time, a renewal process looks stochastically the same forward or backward in time. Changing the “arrow of time” reverses the role of the current and remaining life. Asymptotic results for the total life process now follow trivially from the results for the current and remaining life processes.

Limits for the total life process

1. $\frac{1}{t} \int_0^t L_s ds \rightarrow \frac{\nu}{\mu}$ as $t \rightarrow \infty$ with probability 1
2. $\frac{1}{t} \int_0^t l(s) ds = \frac{\nu}{\mu}$ as $t \rightarrow \infty$

Replacement Models

Consider again a standard renewal process as defined in the Introduction, with interarrival sequence $\mathbf{X} = (X_1, X_2, \dots)$, arrival sequence $\mathbf{T} = (T_0, T_1, \dots)$, and counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$. One of the most basic applications is to reliability, where a device operates for a random lifetime, fails, and then is replaced by a new device, and the process continues. In this model, X_n is the lifetime and T_n the failure time of the n th device in service, for $n \in \mathbb{N}_+$, while N_t is the number of failures in $[0, t]$ for $t \in [0, \infty)$. As usual, F denotes the distribution function of a generic lifetime X , and $F^c = 1 - F$ the corresponding right distribution function (reliability function). Sometimes, the device is actually a *system* with a number of critical components—the failure of any of the critical components causes the system to fail.

Replacement models are variations on the basic model in which the device is replaced (or the critical components replaced) at times other than failure. Often the cost a of a *planned replacement* is less than the cost b of an *emergency replacement* (at failure), so

replacement models can make economic sense. We will consider the the most common model.

In the *age replacement model*, the device is replaced either when it fails or when it reaches a specified age $s \in (0, \infty)$. This model gives rise to a new renewal process with interarrival sequence $\mathbf{U} = (U_1, U_2, \dots)$ where $U_n = \min\{X_n, s\}$ for $n \in \mathbb{N}_+$. If $a, b \in (0, \infty)$ are the costs of planned and unplanned replacements, respectively, then the cost associated with the renewal period U_n is

$$Y_n = a\mathbf{1}(U_n = s) + b\mathbf{1}(U_n < s) = a\mathbf{1}(X_n \geq s) + b\mathbf{1}(X_n < s) \quad (15.6.18)$$

Clearly $((U_1, Y_1), (U_2, Y_2), \dots)$ satisfies the assumptions of a renewal reward process given [above](#). The model makes mathematical sense for any $a, b \in (0, \infty)$ but if $a \geq b$, so that the planned cost of replacement is at least as large as the unplanned cost of replacement, then $Y_n \geq b$ for $n \in \mathbb{N}_+$, so the model makes no financial sense. Thus we assume that $a < b$.

In the age replacement model, with planned replacement at age $s \in (0, \infty)$,

1. The expected cost of a renewal period is $\mathbb{E}(Y) = aF^c(s) + bF(s)$.
2. The expected length of a renewal period is $\mathbb{E}(U) = \int_0^s F^c(x) dx$

The limiting expected cost per unit time is

$$C(s) = \frac{aF^c(s) + bF(s)}{\int_0^s F^c(x) dx} \quad (15.6.19)$$

Proof

Parts (a) and (b) follow from the definition of the reward Y and the renewal period U , and then the formula for $C(s)$ follows from the reward renewal theorem [above](#)

So naturally, given the costs a and b , and the lifetime distribution function F , the goal is be to find the value of s that minimizes $C(s)$; this value of s is the *optimal replacement time*. Of course, the optimal time may not exist.

Properties of C

1. $C(s) \rightarrow \infty$ as $s \downarrow 0$
2. $C(s) \rightarrow \mu/b$ as $s \uparrow \infty$

Proof

1. Recall that $F^c(0) > 0$ and $\int_0^s F^c(x) dx \rightarrow 0$ as $s \downarrow 0$
2. As $s \rightarrow \infty$ note that $F^c(s) \rightarrow 0$, $F(s) \rightarrow 1$ and $\int_0^s F^c(x) dx \rightarrow \int_0^\infty F^c(x) dx = \mu$

As $s \rightarrow \infty$, the age replacement model becomes the standard (unplanned) model with limiting expected average cost b/μ .

Suppose that the lifetime of the device (in appropriate units) has the standard exponential distribution. Find $C(s)$ and solve the optimal age replacement problem.

Answer

The exponential reliability function is $F^c(t) = e^{-t}$ for $t \in [0, \infty)$. After some algebra, the long term expected average cost per unit time is

$$C(s) = \frac{a}{e^s - 1} + b, \quad s \in [0, \infty) \quad (15.6.20)$$

But clearly $C(s)$ is strictly decreasing in s , with limit b , so there is no minimum value.

The last result is hardly surprising. A device with an exponentially distributed lifetime does not age—if it has not failed, it's just as good as new. More generally, age replacement does not make sense for any device with decreasing failure rate. Such devices *improve* with age.

Suppose that the lifetime of the device (in appropriate units) has the gamma distribution with shape parameter 2 and scale parameter 1. Suppose that the costs (in appropriate units) are $a = 1$ and $b = 5$.

1. Find $C(s)$.

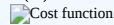
2. Sketch the graph of $C(s)$.
3. Solve numerically the optimal age replacement problem.

Answer

The gamma reliability function is $F^c(t) = e^{-t}(1+t)$ for $t \in [0, \infty)$

1.
$$C(s) = \frac{4+4s-5e^s}{2+s-2e^s}, \quad s \in (0, \infty) \quad (15.6.21)$$

2. The graph of $C(s)$ on the interval $(0, 5]$

 Cost function

3. C is minimized for replacement time $s \approx 1.3052$. The optimal cost is about 2.26476.

Suppose again that the lifetime of the device (in appropriate units) has the gamma distribution with shape parameter 2 and scale parameter 1. But suppose now that the costs (in appropriate units) are $a = 1$ and $b = 2$.

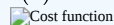
1. Find $C(s)$.
2. Sketch the graph of $C(s)$.
3. Solve the optimal age replacement problem.

Answer

The gamma reliability function is $F^c(t) = e^{-t}(1+t)$ for $t \in [0, \infty)$

1.
$$C(s) = \frac{2e^s - (1+s)}{2e^s - (2+s)}, \quad s \in (0, \infty) \quad (15.6.22)$$

2. The graph of $C(s)$ on the interval $(0, 4]$

 Cost function

3. C is strictly decreasing on $[0, \infty)$ with limit 1, so there is no minimum value.

In the last case, the difference between the cost of an emergency replacement and a planned replacement is not great enough for age replacement to make sense.

Suppose that the lifetime of the device (in appropriately scaled units) is uniformly distributed on the interval $[0, 1]$. Find $C(s)$ and solve the optimal replacement problem. Give the results explicitly for the following costs:

1. $a = 4, b = 6$
2. $a = 2, b = 5$
3. $a = 1, b = 10$

Proof

The reliability function is $F^c(t) = 1 - t$ for $t \in [0, 1]$. After standard computations,

$$C(s) = 2 \frac{a(1-s) + bs}{s(2-s)}, \quad s \in (0, 1] \quad (15.6.23)$$

After more standard calculus, the optimal replacement time is

$$s = \frac{\sqrt{2ab - a^2} - a}{b - a} \quad (15.6.24)$$

1. $s = 2(\sqrt{2} - 1) \approx 0.828, C \approx 11.657$
2. $s = \frac{2}{3}, C = 9$
3. $s = \frac{\sqrt{19}-1}{9} \approx 0.373, C \approx 14.359$

Thinning

We start with a standard renewal process with interarrival sequence $\mathbf{X} = (X_1, X_2, \dots)$, arrival sequence $\mathbf{T} = (T_0, T_1, \dots)$ and counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$. As usual, let $\mu = \mathbb{E}(X)$ denote the mean of an interarrival time. For $n \in \mathbb{N}_+$, suppose

now that arrival n is either *accepted* or *rejected*, and define random variable Y_n to be 1 in the first case and 0 in the second. Let $Z_n = (X_n, Y_n)$ denote the interarrival time and rejection variable pair for $n \in \mathbb{N}_+$, and assume that $\mathbf{Z} = (Z_1, Z_2, \dots)$ is an independent, identically distributed sequence.

Note that we have the structure of a renewal reward process, and so in particular, $\mathbf{Y} = (Y_1, Y_2, \dots)$ is a sequence of Bernoulli trials. Let p denote the parameter of this sequence, so that p is the probability of accepting an arrival. The procedure of accepting or rejecting points in a point process is known as *thinning* the point process. We studied thinning of the Poisson process. In the notation of this section, note that the reward process $\mathbf{R} = \{R_t : t \in [0, \infty)\}$ is the thinned counting process. That is,

$$R_t = \sum_{i=1}^{N_t} Y_i \quad (15.6.25)$$

is the number of accepted points in $[0, t]$ for $t \in [0, \infty)$. So then $r(t) = \mathbb{E}(R_t)$ is the expected number of accepted points in $[0, t]$. The renewal reward theorem gives the asymptotic behavior.

Limits for the thinned process.

1. $R_t/t \rightarrow p/\mu$ as $t \rightarrow \infty$
2. $r(t)/t \rightarrow p/\mu$ as $t \rightarrow \infty$

Proof

This follows immediately from the renewal reward theorem [above](#), since $\nu = p$.

This page titled [15.6: Renewal Reward Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

16: Markov Processes

A Markov process is a random process in which the future is independent of the past, given the present. Thus, Markov processes are the natural stochastic analogs of the deterministic processes described by differential and difference equations. They form one of the most important classes of random processes.

- [16.1: Introduction to Markov Processes](#)
- [16.2: Potentials and Generators for General Markov Processes](#)
- [16.3: Introduction to Discrete-Time Chains](#)
- [16.4: Transience and Recurrence for Discrete-Time Chains](#)
- [16.5: Periodicity of Discrete-Time Chains](#)
- [16.6: Stationary and Limiting Distributions of Discrete-Time Chains](#)
- [16.7: Time Reversal in Discrete-Time Chains](#)
- [16.8: The Ehrenfest Chains](#)
- [16.9: The Bernoulli-Laplace Chain](#)
- [16.10: Discrete-Time Reliability Chains](#)
- [16.11: Discrete-Time Branching Chain](#)
- [16.12: Discrete-Time Queuing Chains](#)
- [16.13: Discrete-Time Birth-Death Chains](#)
- [16.14: Random Walks on Graphs](#)
- [16.15: Introduction to Continuous-Time Markov Chains](#)
- [16.16: Transition Matrices and Generators of Continuous-Time Chains](#)
- [16.17: Potential Matrices](#)
- [16.18: Stationary and Limiting Distributions of Continuous-Time Chains](#)
- [16.19: Time Reversal in Continuous-Time Chains](#)
- [16.20: Chains Subordinate to the Poisson Process](#)
- [16.21: Continuous-Time Birth-Death Chains](#)
- [16.22: Continuous-Time Queuing Chains](#)
- [16.23: Continuous-Time Branching Chains](#)

This page titled [16: Markov Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.1: Introduction to Markov Processes

A *Markov process* is a random process indexed by time, and with the property that the future is independent of the past, given the present. Markov processes, named for Andrei Markov, are among the most important of all random processes. In a sense, they are the stochastic analogs of [differential equations and recurrence relations](#), which are of course, among the most important deterministic processes.

The complexity of the theory of Markov processes depends greatly on whether the time space T is \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time) and whether the state space is *discrete* (countable, with all subsets measurable) or a more general topological space. When $T = [0, \infty)$ or when the state space is a general space, continuity assumptions usually need to be imposed in order to rule out various types of weird behavior that would otherwise complicate the theory.

When the state space is discrete, Markov processes are known as *Markov chains*. The general theory of Markov chains is mathematically rich and relatively simple.

- When $T = \mathbb{N}$ and the state space is discrete, Markov processes are known as *discrete-time Markov chains*. The theory of such processes is mathematically elegant and complete, and is understandable with minimal reliance on measure theory. Indeed, the main tools are basic probability and linear algebra. Discrete-time Markov chains are studied in this chapter, along with a number of special models.
- When $T = [0, \infty)$ and the state space is discrete, Markov processes are known as *continuous-time Markov chains*. If we avoid a few technical difficulties (created, as always, by the continuous time space), the theory of these processes is also reasonably simple and mathematically very nice. The Markov property implies that the process, sampled at the random times when the state changes, forms an embedded *discrete-time* Markov chain, so we can apply the theory that we will have already learned. The Markov property also implies that the holding time in a state has the memoryless property and thus must have an exponential distribution, a distribution that we know well. In terms of what you may have already studied, the Poisson process is a simple example of a continuous-time Markov chain.

For a general state space, the theory is more complicated and technical, as noted above. However, we can distinguish a couple of classes of Markov processes, depending again on whether the time space is discrete or continuous.

- When $T = \mathbb{N}$ and $S = \mathbb{R}$, a simple example of a Markov process is the partial sum process associated with a sequence of independent, identically distributed real-valued random variables. Such sequences are studied in the chapter on random samples (but not as Markov processes), and revisited [below](#).
- In the case that $T = [0, \infty)$ and $S = \mathbb{R}$ or more generally $S = \mathbb{R}^k$, the most important Markov processes are the *diffusion processes*. Generally, such processes can be constructed via stochastic differential equations from Brownian motion, which thus serves as the quintessential example of a Markov process in continuous time and space.

The goal of this section is to give a broad sketch of the general theory of Markov processes. Some of the statements are not completely rigorous and some of the proofs are omitted or are sketches, because we want to emphasize the main ideas without getting bogged down in technicalities. If you are a new student of probability you may want to just browse this section, to get the basic ideas and notation, but skipping over the proofs and technical details. Then jump ahead to the study of discrete-time Markov chains. On the other hand, to understand this section in more depth, you will need to review topics in the chapter on foundations and in the chapter on stochastic processes.

Basic Theory

Preliminaries

As usual, our starting point is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so that Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on (Ω, \mathcal{F}) . The time set T is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). In the first case, T is given the discrete topology and in the second case T is given the usual Euclidean topology. In both cases, T is given the Borel σ -algebra \mathcal{T} , the σ -algebra generated by the open sets. In the discrete case when $T = \mathbb{N}$, this is simply the power set of T so that every subset of T is measurable; every function from T to another measurable space is measurable; and every function from T to another topological space is continuous. The time space (T, \mathcal{T}) has a natural measure; counting measure $\#$ in the discrete case, and Lebesgue in the continuous case.

The set of states S also has a σ -algebra \mathcal{S} of admissible subsets, so that (S, \mathcal{S}) is the *state space*. Usually S has a topology and \mathcal{S} is the Borel σ -algebra generated by the open sets. A typical set of assumptions is that the topology on S is *LCCB*: locally compact, Hausdorff, and with a countable base. These particular assumptions are general enough to capture all of the most important processes that occur in applications and yet are restrictive enough for a nice mathematical theory. Usually, there is a natural positive measure λ on the state space (S, \mathcal{S}) . When S has an LCCB topology and \mathcal{S} is the Borel σ -algebra, the measure λ will usually be a *Borel measure* satisfying $\lambda(C) < \infty$ if $C \subseteq S$ is compact. The term *discrete state space* means that S is countable with $\mathcal{S} = \mathcal{P}(S)$, the collection of all subsets of S . Thus every subset of S is measurable, as is every function from S to another measurable space. This is the Borel σ -algebra for the discrete topology on S , so that every function from S to another topological space is continuous. The compact sets are simply the finite sets, and the reference measure is $\#$, counting measure. If $S = \mathbb{R}^k$ for some $k \in \mathbb{N}$ (another common case), then we usually give S the Euclidean topology (which is LCCB) so that \mathcal{S} is the usual Borel σ -algebra. The compact sets are the closed, bounded sets, and the reference measure λ is k -dimensional Lebesgue measure.

Clearly, the topological and measure structures on T are not really necessary when $T = \mathbb{N}$, and similarly these structures on S are not necessary when S is countable. But the main point is that the assumptions unify the discrete and the common continuous cases. Also, it should be noted that much more general state spaces (and more general time spaces) are possible, but most of the important Markov processes that occur in applications fit the setting we have described here.

Various spaces of real-valued functions on S play an important role. Let \mathcal{B} denote the collection of bounded, measurable functions $f : S \rightarrow \mathbb{R}$. With the usual (pointwise) addition and scalar multiplication, \mathcal{B} is a vector space. We give \mathcal{B} the supremum norm, defined by $\|f\| = \sup\{|f(x)| : x \in S\}$.

Suppose now that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on $(\Omega, \mathcal{F}, \mathbb{P})$ with state space S and time space T . Thus, X_t is a random variable taking values in S for each $t \in T$, and we think of $X_t \in S$ as the *state* of a system at *time* $t \in T$. We also assume that we have a collection $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ of σ -algebras with the properties that X_t is measurable with respect to \mathcal{F}_t for $t \in T$, and the $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ for $s, t \in T$ with $s \leq t$. Intuitively, \mathcal{F}_t is the collection of event up to time $t \in T$. Technically, the assumptions mean that \mathfrak{F} is a filtration and that the process \mathbf{X} is *adapted* to \mathfrak{F} . The most basic (and coarsest) filtration is the *natural filtration* $\mathfrak{F}^0 = \{\mathcal{F}_t^0 : t \in T\}$ where $\mathcal{F}_t^0 = \sigma\{X_s : s \in T, s \leq t\}$, the σ -algebra generated by the process up to time $t \in T$. In continuous time, however, it is often necessary to use slightly finer σ -algebras in order to have a nice mathematical theory. In particular, we often need to assume that the filtration \mathfrak{F} is *right continuous* in the sense that $\mathcal{F}_{t+} = \mathcal{F}_t$ for $t \in T$ where $\mathcal{F}_{t+} = \bigcap\{\mathcal{F}_s : s \in T, s > t\}$. We can accomplish this by taking $\mathfrak{F} = \mathfrak{F}_+^0$ so that $\mathcal{F}_t = \mathcal{F}_{t+}^0$ for $t \in T$, and in this case, \mathfrak{F} is referred to as the *right continuous refinement* of the natural filtration. We also sometimes need to assume that \mathfrak{F} is *complete* with respect to \mathbb{P} in the sense that if $A \in \mathcal{S}$ with $\mathbb{P}(A) = 0$ and $B \subseteq A$ then $B \in \mathcal{F}_0$. That is, \mathcal{F}_0 contains all of the *null events* (and hence also all of the *almost certain* events), and therefore so does \mathcal{F}_t for all $t \in T$.

Definitions

The random process \mathbf{X} is a *Markov process* if

$$\mathbb{P}(X_{s+t} \in A \mid \mathcal{F}_s) = \mathbb{P}(X_{s+t} \in A \mid X_s) \quad (16.1.1)$$

for all $s, t \in T$ and $A \in \mathcal{S}$.

The defining condition, known appropriately enough as the *Markov property*, states that the conditional distribution of X_{s+t} given \mathcal{F}_s is the same as the conditional distribution of X_{s+t} just given X_s . Think of s as the present time, so that $s+t$ is a time in the future. If we know the present state X_s , then any additional knowledge of events in the past is irrelevant in terms of predicting the future state X_{s+t} . Technically, the conditional probabilities in the definition are random variables, and the equality must be interpreted as holding with probability 1. As you may recall, conditional expected value is a more general and useful concept than conditional probability, so the following theorem may come as no surprise.

The random process \mathbf{X} is a Markov process if and only if

$$\mathbb{E}[f(X_{s+t}) \mid \mathcal{F}_s] = \mathbb{E}[f(X_{s+t}) \mid X_s] \quad (16.1.2)$$

for every $s, t \in T$ and every $f \in \mathcal{B}$.

Proof sketch

The condition in this theorem clearly implies the Markov property, by letting $f = \mathbf{1}_A$, the indicator function of $A \in \mathcal{S}$. The converse is a classical bootstrapping argument: the Markov property implies the expected value condition

1. First when $f = \mathbf{1}_A$ for $A \in \mathcal{S}$ (by definition).
2. Next when $f \in \mathcal{B}$ is a simple function, by linearity.
3. Next when $f \in \mathcal{B}$ is nonnegative, by the monotone convergence theorem.
4. Finally for general $f \in \mathcal{B}$ by considering positive and negative parts.

Technically, we should say that \mathbf{X} is a Markov process *relative to the filtration* \mathfrak{F} . If \mathbf{X} satisfies the Markov property relative to a filtration, then it satisfies the Markov property relative to any coarser filtration.

Suppose that the stochastic process $\mathbf{X} = \{X_t : t \in T\}$ is adapted to the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ and that $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ is a filtration that is finer than \mathfrak{F} . If \mathbf{X} is a Markov process relative to \mathfrak{G} then \mathbf{X} is a Markov process relative to \mathfrak{F} .

Proof

First recall that \mathbf{X} is adapted to \mathfrak{G} since \mathbf{X} is adapted to \mathfrak{F} . If $s, t \in T$ and $f \in \mathcal{B}$ then

$$\mathbb{E}[f(X_{s+t}) \mid \mathcal{F}_s] = \mathbb{E}(\mathbb{E}[f(X_{s+t}) \mid \mathcal{G}_s] \mid \mathcal{F}_s) = \mathbb{E}(\mathbb{E}[f(X_{s+t}) \mid X_s] \mid \mathcal{F}_s) = \mathbb{E}[f(X_{s+t}) \mid X_s] \quad (16.1.3)$$

The first equality is a basic property of conditional expected value. The second uses the fact that \mathbf{X} is Markov relative to \mathfrak{G} , and the third follows since X_s is measurable with respect to \mathcal{F}_s .

In particular, if \mathbf{X} is a Markov process, then \mathbf{X} satisfies the Markov property relative to the natural filtration \mathfrak{F}^0 . The theory of Markov processes is simplified considerably if we add an additional assumption.

A Markov process \mathbf{X} is *time homogeneous* if

$$\mathbb{P}(X_{s+t} \in A \mid X_s = x) = \mathbb{P}(X_t \in A \mid X_0 = x) \quad (16.1.4)$$

for every $s, t \in T$, $x \in S$ and $A \in \mathcal{S}$.

So if \mathbf{X} is homogeneous (we usually don't bother with the *time* adjective), then the process $\{X_{s+t} : t \in T\}$ given $X_s = x$ is equivalent (in distribution) to the process $\{X_t : t \in T\}$ given $X_0 = x$. For this reason, the initial distribution is often unspecified in the study of Markov processes—if the process

is in state $x \in S$ at a particular time $s \in T$, then it doesn't really matter how the process got to state x ; the process essentially “starts over”, independently of the past. The term *stationary* is sometimes used instead of homogeneous.

From now on, we will usually assume that our Markov processes are homogeneous. This is not as big of a loss of generality as you might think. A non-homogeneous process can be turned into a homogeneous process by enlarging the state space, as shown [below](#). For a homogeneous Markov process, if $s, t \in T$, $x \in S$, and $f \in \mathcal{B}$, then

$$\mathbb{E}[f(X_{s+t}) \mid X_s = x] = \mathbb{E}[f(X_t) \mid X_0 = x] \quad (16.1.5)$$

Feller Processes

In continuous time, or with general state spaces, Markov processes can be very strange without additional continuity assumptions. Suppose (as is usually the case) that S has an LCCB topology and that \mathcal{S} is the Borel σ -algebra. Let \mathcal{C} denote the collection of bounded, continuous functions $f : S \rightarrow \mathbb{R}$. Let \mathcal{C}_0 denote the collection of continuous functions $f : S \rightarrow \mathbb{R}$ that *vanish at ∞* . The last phrase means that for every $\epsilon > 0$, there exists a compact set $C \subseteq S$ such that $|f(x)| < \epsilon$ if $x \notin C$. With the usual (pointwise) operations of addition and scalar multiplication, \mathcal{C}_0 is a vector subspace of \mathcal{C} , which in turn is a vector subspace of \mathcal{B} . Just as with \mathcal{B} , the supremum norm is used for \mathcal{C} and \mathcal{C}_0 .

A Markov process $\mathbf{X} = \{X_t : t \in T\}$ is a *Feller process* if the following conditions are satisfied.

1. *Continuity in space*: For $t \in T$ and $y \in S$, the distribution of X_t given $X_0 = x$ converges to the distribution of X_t given $X_0 = y$ as $x \rightarrow y$.
2. *Continuity in time*: Given $X_0 = x$ for $x \in S$, X_t converges in probability to x as $t \downarrow 0$.

Additional details

1. This means that $\mathbb{E}[f(X_t) \mid X_0 = x] \rightarrow \mathbb{E}[f(X_t) \mid X_0 = y]$ as $x \rightarrow y$ for every $f \in \mathcal{C}$.
2. This means that $\mathbb{P}[X_t \in U \mid X_0 = x] \rightarrow 1$ as $t \downarrow 0$ for every neighborhood U of x .

Feller processes are named for William Feller. Note that if S is discrete, (a) is automatically satisfied and if T is discrete, (b) is automatically satisfied. In particular, every discrete-time Markov chain is a Feller Markov process. There are certainly more general Markov processes, but most of the important processes that occur in applications are Feller processes, and a number of nice properties flow from the assumptions. Here is the first:

If $\mathbf{X} = \{X_t : t \in T\}$ is a Feller process, then there is a version of \mathbf{X} such that $t \mapsto X_t(\omega)$ is continuous from the right and has left limits for every $\omega \in \Omega$.

Again, this result is only interesting in continuous time $T = [0, \infty)$. Recall that for $\omega \in \Omega$, the function $t \mapsto X_t(\omega)$ is a *sample path* of the process. So we will often assume that a Feller Markov process has sample paths that are right continuous have left limits, since we know there is a version with these properties.

Stopping Times and the Strong Markov Property

For our next discussion, you may need to review again the section on filtrations and stopping times. To give a quick review, suppose again that we start with our probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ (so that we have a *filtered probability space*).

Since time (past, present, future) plays such a fundamental role in Markov processes, it should come as no surprise that random times are important. We often need to allow random times to take the value ∞ , so we need to enlarge the set of times to $T_\infty = T \cup \{\infty\}$. The topology on T is extended to T_∞ by the rule that for $s \in T$, the set $\{t \in T_\infty : t > s\}$ is an open neighborhood of ∞ . This is the one-point compactification of T and is used so that the notion of *time converging to infinity* is preserved. The Borel σ -algebra \mathcal{T}_∞ is used on T_∞ , which again is just the power set in the discrete case.

If $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process on the sample space (Ω, \mathcal{F}) , and if τ is a random time, then naturally we want to consider the state X_τ at the random time. There are two problems. First if τ takes the value ∞ , X_τ is not defined. The usual solution is to add a new “death state” δ to the set of states S , and then to give $S_\delta = S \cup \{\delta\}$ the σ algebra $\mathcal{S}_\delta = \mathcal{S} \cup \{A \cup \{\delta\} : A \in \mathcal{S}\}$. A function $f \in \mathcal{B}$ is extended to S_δ by the rule $f(\delta) = 0$. The second problem is that X_τ may not be a valid random variable (that is, measurable) unless we assume that the stochastic process \mathbf{X} is measurable. Recall that this means that $\mathbf{X} : \Omega \times T \rightarrow S$ is measurable relative to $\mathcal{F} \otimes \mathcal{T}$ and \mathcal{S} . (This is always true in discrete time.)

Recall next that a random time τ is a stopping time (also called a *Markov time* or an *optional time*) relative to \mathfrak{F} if $\{\tau \leq t\} \in \mathcal{F}_t$ for each $t \in T$. Intuitively, we can tell whether or not $\tau \leq t$ from the information available to us at time t . In a sense, a stopping time is a random time that does not require that we see into the future. Of course, the concept depends critically on the filtration. Recall that if a random time τ is a stopping time for a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ then it is also a stopping time for a finer filtration $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$, so that $\mathcal{F}_t \subseteq \mathcal{G}_t$ for $t \in T$. Thus, the finer the filtration, the larger the collection of stopping times. In fact if the filtration is the trivial one where $\mathcal{F}_t = \mathcal{F}$ for all $t \in T$ (so that all information is available to us from the beginning of time), then *any* random time is a stopping time. But of course, this trivial filtration is usually not sensible.

Next, recall that if τ is a stopping time for the filtration \mathfrak{F} , then the σ -algebra \mathcal{F}_τ associated with τ is given by

$$\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \in T\} \quad (16.1.6)$$

Intuitively, \mathcal{F}_τ is the collection of events up to the random time τ , analogous to the \mathcal{F}_t which is the collection of events up to the deterministic time $t \in T$. If $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process adapted to \mathfrak{F} and if τ is a stopping time relative to \mathfrak{F} , then we would hope that X_τ is measurable with respect to \mathcal{F}_τ just as X_t is measurable with respect to \mathcal{F}_t for deterministic $t \in T$. However, this will generally not be the case unless \mathbf{X} is *progressively measurable* relative to \mathfrak{F} , which means that $\mathbf{X} : \Omega \times T_t \rightarrow S$ is measurable with respect to $\mathcal{F}_t \otimes \mathcal{T}_t$ and \mathcal{S} where $T_t = \{s \in T : s \leq t\}$

and \mathcal{F}_t the corresponding Borel σ -algebra. This is always true in discrete time, of course, and more generally if S has an LCCB topology with \mathcal{S} the Borel σ -algebra, and \mathbf{X} is right continuous. If \mathbf{X} is progressively measurable with respect to \mathfrak{F} then \mathbf{X} is measurable and \mathbf{X} is adapted to \mathfrak{F} .

The *strong Markov property* for our stochastic process $\mathbf{X} = \{X_t : t \in T\}$ states that the future is independent of the past, given the present, when the present time is a stopping time.

The random process \mathbf{X} is a *strong Markov process* if

$$\mathbb{E}[f(X_{\tau+t}) \mid \mathcal{F}_\tau] = \mathbb{E}[f(X_{\tau+t}) \mid X_\tau] \quad (16.1.7)$$

for every $t \in T$, stopping time τ , and $f \in \mathcal{B}$.

As with the regular Markov property, the strong Markov property depends on the underlying filtration \mathfrak{F} . If the property holds with respect to a given filtration, then it holds with respect to a coarser filtration.

Suppose that the stochastic process $\mathbf{X} = \{X_t : t \in T\}$ is progressively measurable relative to the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ and that the filtration $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ is finer than \mathfrak{F} . If \mathbf{X} is a strong Markov process relative to \mathfrak{G} then \mathbf{X} is a strong Markov process relative to \mathfrak{F} .

Proof

Recall again that since \mathbf{X} is adapted to \mathfrak{F} , it is also adapted to \mathfrak{G} . Suppose that τ is a finite stopping time for \mathfrak{F} and that $t \in T$ and $f \in \mathcal{B}$. Then τ is also a stopping time for \mathfrak{G} , and $\mathcal{F}_\tau \subseteq \mathcal{G}_\tau$. Hence

$$\mathbb{E}[f(X_{\tau+t}) \mid \mathcal{F}_\tau] = \mathbb{E}(\mathbb{E}[f(X_{\tau+t}) \mid \mathcal{G}_\tau] \mid \mathcal{F}_\tau) = \mathbb{E}(\mathbb{E}[f(X_{\tau+t}) \mid X_\tau] \mid \mathcal{F}_\tau) = \mathbb{E}[f(X_{\tau+t}) \mid X_\tau] \quad (16.1.8)$$

The first equality is a basic property of conditional expected value. The second uses the fact that \mathbf{X} has the strong Markov property relative to \mathfrak{G} , and the third follows since X_τ measurable with respect to \mathcal{F}_τ . In continuous time, it's last step that requires progressive measurability.

So if \mathbf{X} is a strong Markov process, then \mathbf{X} satisfies the strong Markov property relative to its natural filtration. Again there is a tradeoff: finer filtrations allow more stopping times (generally a good thing), but make the strong Markov property harder to satisfy and may not be reasonable (not so good). So we usually don't want filtrations that are too much finer than the natural one.

With the strong Markov and homogeneous properties, the process $\{X_{\tau+t} : t \in T\}$ given $X_\tau = x$ is equivalent in distribution to the process $\{X_t : t \in T\}$ given $X_0 = x$. Clearly, the strong Markov property implies the ordinary Markov property, since a fixed time $t \in T$ is trivially also a stopping time. The converse is true in discrete time.

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a (homogeneous) Markov process in discrete time. Then \mathbf{X} is a strong Markov process.

As always in continuous time, the situation is more complicated and depends on the continuity of the process \mathbf{X} and the filtration \mathfrak{F} . Here is the standard result for Feller processes.

If $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Feller Markov process, then \mathbf{X} is a strong Markov process relative to filtration \mathfrak{F}_+^0 , the right-continuous refinement of the natural filtration.

Transition Kernels of Markov Processes

For our next discussion, you may need to review the section on kernels and operators in the chapter on expected value. Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a (homogeneous) Markov process with state space S and time space T , as described above. The kernels in the following definition are of fundamental importance in the study of \mathbf{X}

For $t \in T$, let

$$P_t(x, A) = \mathbb{P}(X_t \in A \mid X_0 = x), \quad x \in S, A \in \mathcal{S} \quad (16.1.9)$$

Then P_t is a probability kernel on (S, \mathcal{S}) , known as the *transition kernel* of \mathbf{X} for time t .

Proof

Fix $t \in T$. The measurability of $x \mapsto \mathbb{P}(X_t \in A \mid X_0 = x)$ for $A \in \mathcal{S}$ is built into the definition of conditional probability. Also, of course, $A \mapsto \mathbb{P}(X_t \in A \mid X_0 = x)$ is a probability measure on \mathcal{S} for $x \in S$. In general, the conditional distribution of one random variable, conditioned on a value of another random variable defines a probability kernel.

That is, $P_t(x, \cdot)$ is the conditional distribution of X_t given $X_0 = x$ for $t \in T$ and $x \in S$. By the time homogenous property, $P_t(x, \cdot)$ is also the conditional distribution of X_{s+t} given $X_s = x$ for $s \in T$:

$$P_t(x, A) = \mathbb{P}(X_{s+t} \in A \mid X_s = x), \quad s, t \in T, x \in S, A \in \mathcal{S} \quad (16.1.10)$$

Note that $P_0 = I$, the *identity kernel* on (S, \mathcal{S}) defined by $I(x, A) = \mathbf{1}(x \in A)$ for $x \in S$ and $A \in \mathcal{S}$, so that $I(x, A) = 1$ if $x \in A$ and $I(x, A) = 0$ if $x \notin A$. Recall also that usually there is a natural reference measure λ on (S, \mathcal{S}) . In this case, the transition kernel P_t will often have a *transition density* p_t with respect to λ for $t \in T$. That is,

$$P_t(x, A) = \mathbb{P}(X_t \in A \mid X_0 = x) = \int_A p_t(x, y) \lambda(dy), \quad x \in S, A \in \mathcal{S} \quad (16.1.11)$$

The next theorem gives the *Chapman-Kolmogorov* equation, named for Sydney Chapman and Andrei Kolmogorov, the fundamental relationship between the probability kernels, and the reason for the name *transition kernel*.

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a Markov process on S with transition kernels $\mathbf{P} = \{P_t : t \in T\}$. If $s, t \in T$, then $P_s P_t = P_{s+t}$. That is,

$$P_{s+t}(x, A) = \int_S P_s(x, dy) P_t(y, A), \quad x \in S, A \in \mathcal{S} \quad (16.1.12)$$

Proof

The Markov property and a conditioning argument are the fundamental tools. Recall again that $P_s(x, \cdot)$ is the conditional distribution of X_s given $X_0 = x$ for $x \in S$. Let $A \in \mathcal{S}$. Conditioning on X_s gives

$$P_{s+t}(x, A) = \mathbb{P}(X_{s+t} \in A \mid X_0 = x) = \int_S P_s(x, dy) \mathbb{P}(X_{s+t} \in A \mid X_s = y, X_0 = x) \quad (16.1.13)$$

But by the Markov and time-homogeneous properties,

$$\mathbb{P}(X_{s+t} \in A \mid X_s = y, X_0 = x) = \mathbb{P}(X_t \in A \mid X_0 = y) = P_t(y, A) \quad (16.1.14)$$

Substituting we have

$$P_{s+t}(x, A) = \int_S P_s(x, dy) P_t(y, A) = (P_s P_t)(x, A) \quad (16.1.15)$$

In the language of functional analysis, \mathbf{P} is a *semigroup*. Recall that the commutative property generally does not hold for the product operation on kernels. However the property does hold for the transition kernels of a homogeneous Markov process. That is, $P_s P_t = P_t P_s = P_{s+t}$ for $s, t \in T$. As a simple corollary, if S has a reference measure, the same basic relationship holds for the transition densities.

Suppose that λ is the reference measure on (S, \mathcal{S}) and that $\mathbf{X} = \{X_t : t \in T\}$ is a Markov process on S and with transition densities $\{p_t : t \in T\}$. If $s, t \in T$ then $p_s p_t = p_{s+t}$. That is,

$$p_t(x, z) = \int_S p_s(x, y) p_t(y, z) \lambda(dy), \quad x, z \in S \quad (16.1.16)$$

Proof

The transition kernels satisfy $P_s P_t = P_{s+t}$. But P_s has density p_s , P_t has density p_t , and P_{s+t} has density p_{s+t} . From a basic result on kernel functions, $P_s P_t$ has density $p_s p_t$ as defined in the theorem.

If $T = \mathbb{N}$ (discrete time), then the transition kernels of \mathbf{X} are just the powers of the *one-step* transition kernel. That is, if we let $P = P_1$ then $P_n = P^n$ for $n \in \mathbb{N}$.

Recall that a kernel defines two operations: operating on the left with positive measures on (S, \mathcal{S}) and operating on the right with measurable, real-valued functions. For the transition kernels of a Markov process, both of these operators have natural interpretations.

Suppose that $s, t \in T$. If μ_s is the distribution of X_s then X_{s+t} has distribution $\mu_{s+t} = \mu_s P_t$. That is,

$$\mu_{s+t}(A) = \int_S \mu_s(dx) P_t(x, A), \quad A \in \mathcal{S} \quad (16.1.17)$$

Proof

Let $A \in \mathcal{S}$. Conditioning on X_s gives

$$\mathbb{P}(X_{s+t} \in A) = \mathbb{E}[\mathbb{P}(X_{s+t} \in A \mid X_s)] = \int_S \mu_s(dx) \mathbb{P}(X_{s+t} \in A \mid X_s = x) = \int_S \mu_s(dx) P_t(x, A) = \mu_s P_t(A) \quad (16.1.18)$$

So if \mathcal{P} denotes the collection of probability measures on (S, \mathcal{S}) , then the left operator P_t maps \mathcal{P} back into \mathcal{P} . In particular, if X_0 has distribution μ_0 (the *initial distribution*) then X_t has distribution $\mu_t = \mu_0 P_t$ for every $t \in T$.

A positive measure μ on (S, \mathcal{S}) is *invariant* for \mathbf{X} if $\mu P_t = \mu$ for every $t \in T$.

Hence if μ is a probability measure that is invariant for \mathbf{X} , and X_0 has distribution μ , then X_t has distribution μ for every $t \in T$ so that the process \mathbf{X} is identically distributed. In discrete time, note that if μ is a positive measure and $\mu P = \mu$ then $\mu P^n = \mu$ for every $n \in \mathbb{N}$, so μ is invariant for \mathbf{X} . The operator on the right is given next.

Suppose that $f : S \rightarrow \mathbb{R}$. If $t \in T$ then (assuming that the expected value exists),

$$P_t f(x) = \int_S P_t(x, dy) f(y) = \mathbb{E}[f(X_t) | X_0 = x], \quad x \in S \quad (16.1.19)$$

Proof

This follows directly from the definitions:

$$P_t f(x) = \int_S P_t(x, dy) f(y), \quad x \in S \quad (16.1.20)$$

and $P_t(x, \cdot)$ is the conditional distribution of X_t given $X_0 = x$.

In particular, the right operator P_t is defined on \mathcal{B} , the vector space of bounded, linear functions $f : S \rightarrow \mathbb{R}$, and in fact is a linear operator on \mathcal{B} . That is, if $f, g \in \mathcal{B}$ and $c \in \mathbb{R}$, then $P_t(f + g) = P_t f + P_t g$ and $P_t(cf) = cP_t f$. Moreover, P_t is a *contraction operator* on \mathcal{B} , since $\|P_t f\| \leq \|f\|$ for $f \in \mathcal{B}$. It then follows that P_t is a continuous operator on \mathcal{B} for $t \in T$.

For the right operator, there is a concept that is complementary to the invariance of of a positive measure for the left operator.

A measurable function $f : S \rightarrow \mathbb{R}$ is *harmonic* for \mathbf{X} if $P_t f = f$ for all $t \in T$.

Again, in discrete time, if $Pf = f$ then $P^n f = f$ for all $n \in \mathbb{N}$, so f is harmonic for \mathbf{X} .

Combining two results above, if X_0 has distribution μ_0 and $f : S \rightarrow \mathbb{R}$ is measurable, then (again assuming that the expected value exists), $\mu_0 P_t f = \mathbb{E}[f(X_t)]$ for $t \in T$. That is,

$$\mathbb{E}[f(X_t)] = \int_S \mu_0(dx) \int_S P_t(x, dy) f(y) \quad (16.1.21)$$

The result [above](#) shows how to obtain the distribution of X_t from the distribution of X_0 and the transition kernel P_t for $t \in T$. But we can do more. Recall that one basic way to describe a stochastic process is to give its *finite dimensional distributions*, that is, the distribution of $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ for every $n \in \mathbb{N}_+$ and every $(t_1, t_2, \dots, t_n) \in T^n$. For a Markov process, the initial distribution and the transition kernels determine the finite dimensional distributions. It's easiest to state the distributions in differential form.

Suppose $\mathbf{X} = \{X_t : t \in T\}$ is a Markov process with transition operators $\mathbf{P} = \{P_t : t \in T\}$, and that $(t_1, \dots, t_n) \in T^n$ with $0 < t_1 < \dots < t_n$. If X_0 has distribution μ_0 , then in differential form, the distribution of $(X_0, X_{t_1}, \dots, X_{t_n})$ is

$$\mu_0(dx_0) P_{t_1}(x_0, dx_1) P_{t_2-t_1}(x_1, dx_2) \cdots P_{t_n-t_{n-1}}(x_{n-1}, dx_n) \quad (16.1.22)$$

Proof

This follows from induction and repeated use of the Markov property. For example, if $t \in T$ with $t > 0$, then conditioning on X_0 gives

$$\mathbb{P}(X_0 \in A, X_t \in B) = \int_A \mathbb{P}(X_t \in B | X_0 = x) \mu_0(dx) = \int_A P_t(x, B) \mu_0(dx) = \int_A \int_B P_t(x, dy) \mu_0(dx) \quad (16.1.23)$$

for $A, B \in \mathcal{S}$. So in differential form, the distribution of (X_0, X_t) is $\mu(dx) P_t(x, dy)$. If $s, t \in T$ with $0 < s < t$, then conditioning on (X_0, X_s) and using our previous result gives

$$\mathbb{P}(X_0 \in A, X_s \in B, X_t \in C) = \int_{A \times B} \mathbb{P}(X_t \in C | X_0 = x, X_s = y) \mu_0(dx) P_s(x, dy) \quad (16.1.24)$$

for $A, B, C \in \mathcal{S}$. But by the Markov property,

$$\mathbb{P}(X_t \in C | X_0 = x, X_s = y) = \mathbb{P}(X_t \in C | X_s = y) = P_{t-s}(y, C) = \int_C P_{t-s}(y, dz) \quad (16.1.25)$$

Hence in differential form, the distribution of (X_0, X_s, X_t) is $\mu_0(dx) P_s(x, dy) P_{t-s}(y, dz)$. Continuing in this manner gives the general result.

This result is very important for constructing Markov processes. If we know how to define the transition kernels P_t for $t \in T$ (based on modeling considerations, for example), and if we know the initial distribution μ_0 , then the last result gives a consistent set of finite dimensional distributions. From the Kolmogorov construction theorem, we know that there *exists* a stochastic process that has these finite dimensional distributions. In continuous time, however, two serious problems remain. First, it's not clear how we would construct the transition kernels so that the crucial Chapman-Kolmogorov equations [above](#) are satisfied. Second, we usually want our Markov process to have certain properties (such as continuity properties of the sample paths) that go beyond the finite dimensional distributions. The first problem will be addressed in the next section, and fortunately, the second problem can be resolved for a Feller process.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a Markov process on an LCCB state space (S, \mathcal{S}) with transition operators $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. Then \mathbf{X} is a Feller process if and only if the following conditions hold:

1. *Continuity in space:* If $f \in \mathcal{C}_0$ and $t \in [0, \infty)$ then $P_t f \in \mathcal{C}_0$
2. *Continuity in time:* If $f \in \mathcal{C}_0$ and $x \in S$ then $P_t f(x) \rightarrow f(x)$ as $t \downarrow 0$.

A semigroup of probability kernels $\mathbf{P} = \{P_t : t \in T\}$ that satisfies the properties in this theorem is called a *Feller semigroup*. So the theorem states that the Markov process \mathbf{X} is Feller if and only if the transition semigroup of transition \mathbf{P} is Feller. As before, (a) is automatically satisfied if S is discrete, and (b) is automatically satisfied if T is discrete. Condition (a) means that P_t is an operator on the vector space \mathcal{C}_0 , in addition to being an operator on the larger space \mathcal{B} . Condition (b) actually implies a stronger form of continuity in time.

Suppose that $\mathbf{P} = \{P_t : t \in T\}$ is a Feller semigroup of transition operators. Then $t \mapsto P_t f$ is continuous (with respect to the supremum norm) for $f \in \mathcal{C}_0$.

Additional details

This means that for $f \in \mathcal{C}_0$ and $t \in [0, \infty)$,

$$\|P_{t+s}f - P_t f\| = \sup\{|P_{t+s}f(x) - P_t f(x)| : x \in S\} \rightarrow 0 \text{ as } s \rightarrow 0 \quad (16.1.26)$$

So combining this with the remark above, note that if \mathbf{P} is a Feller semigroup of transition operators, then $f \mapsto P_t f$ is continuous on \mathcal{C}_0 for fixed $t \in T$, and $t \mapsto P_t f$ is continuous on T for fixed $f \in \mathcal{C}_0$. Again, the importance of this is that we often start with the collection of probability kernels \mathbf{P} and want to know that there exists a nice Markov process \mathbf{X} that has these transition operators.

Sampling in Time

If we *sample* a Markov process at an increasing sequence of points in time, we get another Markov process in discrete time. But the discrete time process may not be homogeneous even if the original process is homogeneous.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a Markov process with state space (S, \mathcal{S}) and that (t_0, t_1, t_2, \dots) is a sequence in T with $0 = t_0 < t_1 < t_2 < \dots$. Let $Y_n = X_{t_n}$ for $n \in \mathbb{N}$. Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a Markov process in discrete time.

Proof

For $n \in \mathbb{N}$, let $\mathcal{G}_n = \sigma\{Y_k : k \in \mathbb{N}, k \leq n\}$, so that $\{\mathcal{G}_n : n \in \mathbb{N}\}$ is the natural filtration associated with \mathbf{Y} . Note that $\mathcal{G}_n \subseteq \mathcal{F}_{t_n}$ and $Y_n = X_{t_n}$ is measurable with respect to \mathcal{G}_n for $n \in \mathbb{N}$. Let $k, n \in \mathbb{N}$ and let $A \in \mathcal{S}$. Then

$$\mathbb{P}(Y_{k+n} \in A \mid \mathcal{G}_k) = \mathbb{P}(X_{t_{k+n}} \in A \mid \mathcal{G}_k) = \mathbb{P}(X_{t_{k+n}} \in A \mid X_{t_k}) = \mathbb{P}(Y_{n+k} \in A \mid Y_k) \quad (16.1.27)$$

If we sample a homogeneous Markov process at multiples of a fixed, positive time, we get a homogeneous Markov process in discrete time.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a homogeneous Markov process with state space (S, \mathcal{S}) and transition kernels $\mathbf{P} = \{P_t : t \in T\}$. Fix $r \in T$ with $r > 0$ and define $Y_n = X_{nr}$ for $n \in \mathbb{N}$. Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a homogeneous Markov process in discrete time, with one-step transition kernel Q given by

$$Q(x, A) = P_r(x, A); \quad x \in S, A \in \mathcal{S} \quad (16.1.28)$$

In some cases, sampling a strong Markov process at an increasing sequence of stopping times yields another Markov process in discrete time. The point of this is that discrete-time Markov processes are often found naturally embedded in continuous-time Markov processes.

Enlarging the State Space

Our first result in this discussion is that a non-homogeneous Markov process can be turned into a homogeneous Markov process, but only at the expense of enlarging the state space.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a non-homogeneous Markov process with state space (S, \mathcal{S}) . Suppose also that τ is a random variable taking values in T , independent of \mathbf{X} . Let $\tau_t = \tau + t$ and let $Y_t = (X_{\tau_t}, \tau_t)$ for $t \in T$. Then $\mathbf{Y} = \{Y_t : t \in T\}$ is a homogeneous Markov process with state space $(S \times T, \mathcal{S} \otimes \mathcal{T})$. For $t \in T$, the transition kernel P_t is given by

$$P_t[(x, r), A \times B] = \mathbb{P}(X_{r+t} \in A \mid X_r = x) \mathbf{1}(r+t \in B), \quad (x, r) \in S \times T, A \times B \in \mathcal{S} \otimes \mathcal{T} \quad (16.1.29)$$

Proof

By definition and the substitution rule,

$$\begin{aligned} \mathbb{P}[Y_{s+t} \in A \times B \mid Y_s = (x, r)] &= \mathbb{P}(X_{\tau_{s+t}} \in A, \tau_{s+t} \in B \mid X_{\tau_s} = x, \tau_s = r) \\ &= \mathbb{P}(X_{\tau+s+t} \in A, \tau+s+t \in B \mid X_{\tau+s} = x, \tau+s = r) \\ &= \mathbb{P}(X_{r+t} \in A, r+t \in B \mid X_r = x, \tau+s = r) \end{aligned}$$

But τ is independent of \mathbf{X} , so the last term is

$$\mathbb{P}(X_{r+t} \in A, r+t \in B \mid X_r = x) = \mathbb{P}(X_{r+t} \in A \mid X_r = x) \mathbf{1}(r+t \in B) \quad (16.1.30)$$

The important point is that the last expression does not depend on s , so \mathbf{Y} is homogeneous.

The trick of enlarging the state space is a common one in the study of stochastic processes. Sometimes a process that has a weaker form of “forgetting the past” can be made into a Markov process by enlarging the state space appropriately. Here is an example in discrete time.

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a random process with state space (S, \mathcal{S}) in which the future depends stochastically on the last two states. That is, for $n \in \mathbb{N}$

$$\mathbb{P}(X_{n+2} \in A \mid \mathcal{F}_{n+1}) = \mathbb{P}(X_{n+2} \in A \mid X_n, X_{n+1}), \quad A \in \mathcal{S} \quad (16.1.31)$$

where $\{\mathcal{F}_n : n \in \mathbb{N}\}$ is the natural filtration associated with the process \mathbf{X} . Suppose also that the process is time homogeneous in the sense that

$$\mathbb{P}(X_{n+2} \in A \mid X_n = x, X_{n+1} = y) = Q(x, y, A) \quad (16.1.32)$$

independently of $n \in \mathbb{N}$. Let $Y_n = (X_n, X_{n+1})$ for $n \in \mathbb{N}$. Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a homogeneous Markov process with state space $(S \times S, \mathcal{S} \otimes \mathcal{S})$. The one step transition kernel P is given by

$$P[(x, y), A \times B] = I(y, A)Q(x, y, B); \quad x, y \in S, A, B \in \mathcal{S} \quad (16.1.33)$$

Proof

Note first that for $n \in \mathbb{N}$, $\sigma\{Y_k : k \leq n\} = \sigma\{(X_k, X_{k+1}) : k \leq n\} = \mathcal{F}_{n+1}$ so the natural filtration associated with the process \mathbf{Y} is $\{\mathcal{F}_{n+1} : n \in \mathbb{N}\}$. If $C \in \mathcal{S} \otimes \mathcal{S}$ then

$$\begin{aligned} \mathbb{P}(Y_{n+1} \in C \mid \mathcal{F}_{n+1}) &= \mathbb{P}[(X_{n+1}, X_{n+2}) \in C \mid \mathcal{F}_{n+1}] \\ &= \mathbb{P}[(X_{n+1}, X_{n+2}) \in C \mid X_n, X_{n+1}] = \mathbb{P}(Y_{n+1} \in C \mid Y_n) \end{aligned}$$

by the given assumption on \mathbf{X} . Hence \mathbf{Y} is a Markov process. Next,

$$\begin{aligned} \mathbb{P}[Y_{n+1} \in A \times B \mid Y_n = (x, y)] &= \mathbb{P}[(X_{n+1}, X_{n+2}) \in A \times B \mid (X_n, X_{n+1}) = (x, y)] \\ &= \mathbb{P}(X_{n+1} \in A, X_{n+2} \in B \mid X_n = x, X_{n+1} = y) = \mathbb{P}(y \in A, X_{n+2} \in B \mid X_n = x, X_{n+1} = y) \\ &= I(y, A)Q(x, y, B) \end{aligned}$$

The last result generalizes in a completely straightforward way to the case where the future of a random process in discrete time depends stochastically on the last k states, for some fixed $k \in \mathbb{N}$.

Examples and Applications

Recurrence Relations and Differential Equations

As noted in the introduction, Markov processes can be viewed as stochastic counterparts of deterministic recurrence relations (discrete time) and differential equations (continuous time). Our goal in this discussion is to explore these connections.

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a stochastic process with state space (S, \mathcal{S}) and that \mathbf{X} satisfies the recurrence relation

$$X_{n+1} = g(X_n), \quad n \in \mathbb{N} \quad (16.1.34)$$

where $g : S \rightarrow S$ is measurable. Then \mathbf{X} is a homogeneous Markov process with one-step transition operator P given by $Pf = f \circ g$ for a measurable function $f : S \rightarrow \mathbb{R}$.

Proof

Clearly \mathbf{X} is uniquely determined by the initial state, and in fact $X_n = g^n(X_0)$ for $n \in \mathbb{N}$ where g^n is the n -fold composition power of g . So the only possible source of randomness is in the initial state. The Markov and time homogeneous properties simply follow from the trivial fact that $g^{m+n}(X_0) = g^n[g^m(X_0)]$, so that $X_{m+n} = g^n(X_m)$. That is, the state at time $m+n$ is completely determined by the state at time m (regardless of the previous states) and the time increment n . In particular, $Pf(x) = \mathbb{E}[g(X_1) \mid X_0 = x] = f[g(x)]$ for measurable $f : S \rightarrow \mathbb{R}$ and $x \in S$. Note that for $n \in \mathbb{N}$, the n -step transition operator is given by $P^n f = f \circ g^n$.

In the deterministic world, as in the stochastic world, the situation is more complicated in continuous time. Nonetheless, the same basic analogy applies.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with state space $(\mathbb{R}, \mathcal{B})$ satisfies the first-order differential equation

$$\frac{d}{dt} X_t = g(X_t) \quad (16.1.35)$$

where $g : \mathbb{R} \rightarrow \mathbb{R}$ is Lipschitz continuous. Then \mathbf{X} is a Feller Markov process

Proof

Recall that *Lipschitz continuous* means that there exists a constant $k \in (0, \infty)$ such that $|g(y) - g(x)| \leq k|x - y|$ for $x, y \in \mathbb{R}$. This is a standard condition on g that guarantees the existence and uniqueness of a solution to the differential equation on $[0, \infty)$. So as before, the only source of randomness in the process comes from the initial value X_0 . Let $t \mapsto X_t(x)$ denote the unique solution with $X_0(x) = x$ for $x \in \mathbb{R}$. The Markov and homogenous properties follow from the fact that $X_{t+s}(x) = X_t(X_s(x))$ for $s, t \in [0, \infty)$ and $x \in S$. That is, the state at time $t+s$ depends only on the state at time s and the time increment t . The Feller properties follow from the continuity of $t \mapsto X_t(x)$ and the continuity of $x \mapsto X_t(x)$. The latter is the *continuous dependence on the initial value*, again guaranteed by the assumptions on g . Note that the transition operator is given by $P_t f(x) = f[X_t(x)]$ for a measurable function $f : S \rightarrow \mathbb{R}$ and $x \in S$.

In differential form, the process can be described by $dX_t = g(X_t) dt$. This essentially deterministic process can be extended to a very important class of Markov processes by the addition of a stochastic term related to Brownian motion. Such *stochastic differential equations* are the main tools for constructing Markov processes known as *diffusion processes*.

Processes with Stationary, Independent Increments

For our next discussion, we consider a general class of stochastic processes that are Markov processes. Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a random process with $S \subseteq \mathbb{R}$ as the set of states. The state space can be discrete (countable) or “continuous”. Typically, S is either \mathbb{N} or \mathbb{Z} in the discrete case, and is either $[0, \infty)$ or \mathbb{R} in the continuous case. In any case, S is given the usual σ -algebra \mathcal{S} of Borel subsets of S (which is the power set in the discrete case). Also, the state space (S, \mathcal{S}) has a natural reference measure λ , namely counting measure in the discrete case and Lebesgue measure in the continuous case. Let $\mathcal{F} = \{\mathcal{F}_t : t \in T\}$ denote the natural filtration, so that $\mathcal{F}_t = \sigma\{X_s : s \in T, s \leq t\}$ for $t \in T$.

The process \mathbf{X} has

1. *Independent increments* if $X_{s+t} - X_s$ is independent of \mathcal{F}_s for all $s, t \in T$.
2. *Stationary increments* if the distribution of $X_{s+t} - X_s$ is the same as the distribution of $X_t - X_0$ for all $s, t \in T$.

A difference of the form $X_{s+t} - X_s$ for $s, t \in T$ is an *increment* of the process, hence the names. Sometimes the definition of stationary increments is that $X_{s+t} - X_s$ have the same distribution as X_t . But this forces $X_0 = 0$ with probability 1, and as usual with Markov processes, it's best to keep the initial distribution unspecified. If \mathbf{X} has stationary increments in the sense of our definition, then the process $\mathbf{Y} = \{Y_t = X_t - X_0 : t \in T\}$ has stationary increments in the more restricted sense. For the remainder of this discussion, assume that $\mathbf{X} = \{X_t : t \in T\}$ has stationary, independent increments, and let Q_t denote the distribution of $X_t - X_0$ for $t \in T$.

$$Q_s * Q_t = Q_{s+t} \text{ for } s, t \in T.$$

Proof

For $s, t \in T$, Q_s is the distribution of $X_s - X_0$, and by the stationary property, Q_t is the distribution of $X_{s+t} - X_s$. By the independence property, $X_s - X_0$ and $X_{s+t} - X_s$ are independent. Hence $Q_s * Q_t$ is the distribution of $[X_s - X_0] + [X_{s+t} - X_s] = X_{s+t} - X_0$. But by definition, this variable has distribution Q_{s+t} .

So the collection of distributions $\mathbf{Q} = \{Q_t : t \in T\}$ forms a semigroup, with convolution as the operator. Note that Q_0 is simply point mass at 0.

The process \mathbf{X} is a homogeneous Markov process. For $t \in T$, the transition operator P_t is given by

$$P_t f(x) = \int_S f(x+y) Q_t(dy), \quad f \in \mathcal{B} \quad (16.1.36)$$

Proof

Suppose that $s, t \in T$ and $f \in \mathcal{B}$,

$$\mathbb{E}[f(X_{s+t}) | \mathcal{F}_s] = \mathbb{E}[f(X_{s+t} - X_s + X_s) | \mathcal{F}_s] = \mathbb{E}[f(X_{s+t}) | X_s] \quad (16.1.37)$$

since $X_{s+t} - X_s$ is independent of \mathcal{F}_s . Moreover, by the stationary property,

$$\mathbb{E}[f(X_{s+t}) | X_s = x] = \int_S f(x+y) Q_t(dy), \quad x \in S \quad (16.1.38)$$

Clearly the semigroup property of $\mathbf{P} = \{P_t : t \in T\}$ (with the usual operator product) is equivalent to the semigroup property of $\mathbf{Q} = \{Q_t : t \in T\}$ (with convolution as the product).

Suppose that for positive $t \in T$, the distribution Q_t has probability density function g_t with respect to the reference measure λ . Then the transition density is

$$p_t(x, y) = g_t(y - x), \quad x, y \in S \quad (16.1.39)$$

Of course, from the result above, it follows that $g_s * g_t = g_{s+t}$ for $s, t \in T$, where here $*$ refers to the convolution operation on probability density functions.

If $Q_t \rightarrow Q_0$ as $t \downarrow 0$ then \mathbf{X} is a Feller Markov process.

Thus, by the general theory sketched above, \mathbf{X} is a strong Markov process, and there exists a version of \mathbf{X} that is right continuous and has left limits. Such a process is known as a *Lévy process*, in honor of Paul Lévy.

For a real-valued stochastic process $\mathbf{X} = \{X_t : t \in T\}$, let m and v denote the mean and variance functions, so that

$$m(t) = \mathbb{E}(X_t), \quad v(t) = \text{var}(X_t); \quad t \in T \quad (16.1.40)$$

assuming of course that these exist. The mean and variance functions for a Lévy process are particularly simple.

Suppose again that \mathbf{X} has stationary, independent increments.

1. If $\mu_0 = \mathbb{E}(X_0) \in \mathbb{R}$ and $\mu_1 = \mathbb{E}(X_1) \in \mathbb{R}$ then $m(t) = \mu_0 + (\mu_1 - \mu_0)t$ for $t \in T$.
2. If in addition, $\sigma_0^2 = \text{var}(X_0) \in (0, \infty)$ and $\sigma_1^2 = \text{var}(X_1) \in (0, \infty)$ then $v(t) = \sigma_0^2 + (\sigma_1^2 - \sigma_0^2)t$ for $t \in T$.

Proof

The proofs are simple using the independent and stationary increments properties. For $t \in T$, let $m_0(t) = \mathbb{E}(X_t - X_0) = m(t) - \mu_0$ and $v_0(t) = \text{var}(X_t - X_0) = v(t) - \sigma_0^2$. denote the mean and variance functions for the centered process $\{X_t - X_0 : t \in T\}$. Now let $s, t \in T$.

1. From the additive property of expected value and the stationary property,

$$m_0(t+s) = \mathbb{E}(X_{t+s} - X_0) = \mathbb{E}[(X_{t+s} - X_s) + (X_s - X_0)] = \mathbb{E}(X_{t+s} - X_s) + \mathbb{E}(X_s - X_0) = m_0(t) + m_0(s) \quad (16.1.41)$$

2. From the additive property of variance for independent variables and the stationary property,

$$v_0(t+s) = \text{var}(X_{t+s} - X_0) = \text{var}[(X_{t+s} - X_s) + (X_s - X_0)] = \text{var}(X_{t+s} - X_s) + \text{var}(X_s - X_0) = v_0(t) + v_0(s) \quad (16.1.42)$$

So m_0 and v_0 satisfy the *Cauchy equation*. In discrete time, it's simple to see that there exists $a \in \mathbb{R}$ and $b^2 \in (0, \infty)$ such that $m_0(t) = at$ and $v_0(t) = b^2t$. The same is true in continuous time, given the continuity assumptions that we have on the process \mathbf{X} . Substituting $t = 1$ we have $a = \mu_1 - \mu_0$ and $b^2 = \sigma_1^2 - \sigma_0^2$, so the results follow,

It's easy to describe processes with stationary independent increments in discrete time.

A process $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ has independent increments if and only if there exists a sequence of independent, real-valued random variables (U_0, U_1, \dots) such that

$$X_n = \sum_{i=0}^n U_i \quad (16.1.43)$$

In addition, \mathbf{X} has stationary increments if and only if (U_1, U_2, \dots) are identically distributed.

Proof

Suppose first that $\mathbf{U} = (U_0, U_1, \dots)$ is a sequence of independent, real-valued random variables, and define $X_n = \sum_{i=0}^n U_i$ for $n \in \mathbb{N}$. Note that $\mathcal{F}_n = \sigma\{X_0, \dots, X_n\} = \sigma\{U_0, \dots, U_n\}$ for $n \in \mathbb{N}$. If $k, n \in \mathbb{N}$ with $k \leq n$, then $X_n - X_k = \sum_{i=k+1}^n U_i$ which is independent of \mathcal{F}_k by the independence assumption on \mathbf{U} . Hence \mathbf{X} has independent increments. Suppose in addition that (U_1, U_2, \dots) are identically distributed. Then the increment $X_n - X_k$ above has the same distribution as $\sum_{i=1}^{n-k} U_i = X_{n-k} - X_0$. Hence \mathbf{X} has stationary increments.

Conversely, suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ has independent increments. Let $U_0 = X_0$ and $U_n = X_n - X_{n-1}$ for $n \in \mathbb{N}_+$. Then $X_n = \sum_{i=0}^n U_i$ for $n \in \mathbb{N}$. As before $\mathcal{F}_n = \sigma\{X_0, \dots, X_n\} = \sigma\{U_0, \dots, U_n\}$ for $n \in \mathbb{N}$. Since \mathbf{X} has independent increments, U_n is independent of \mathcal{F}_{n-1} for $n \in \mathbb{N}_+$, so (U_0, U_1, \dots) are mutually independent. If in addition, \mathbf{X} has stationary increments, $U_n = X_n - X_{n-1}$ has the same distribution as $X_1 - X_0 = U_1$ for $n \in \mathbb{N}_+$. Hence (U_1, U_2, \dots) are identically distributed.

Thus suppose that $\mathbf{U} = (U_0, U_1, \dots)$ is a sequence of independent, real-valued random variables, with (U_1, U_2, \dots) identically distributed with common distribution Q . Then from our main result above, the partial sum process $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ associated with \mathbf{U} is a homogeneous Markov process with one step transition kernel P given by

$$P(x, A) = Q(A - x), \quad x \in S, A \in \mathcal{S} \quad (16.1.44)$$

More generally, for $n \in \mathbb{N}$, the n -step transition kernel is $P^n(x, A) = Q^{*n}(A - x)$ for $x \in S$ and $A \in \mathcal{S}$. This Markov process is known as a *random walk* (although unfortunately, the term *random walk* is used in a number of other contexts as well). The idea is that at time n , the walker moves a (directed) distance U_n on the real line, and these steps are independent and identically distributed. If Q has probability density function g with respect to the reference measure λ , then the one-step transition density is

$$p(x, y) = g(y - x), \quad x, y \in S \quad (16.1.45)$$

Consider the random walk on \mathbb{R} with steps that have the standard normal distribution. Give each of the following explicitly:

1. The one-step transition density.
2. The n -step transition density for $n \in \mathbb{N}_+$.

Proof

1. For $x \in \mathbb{R}$, $p(x, \cdot)$ is the normal PDF with mean x and variance 1:

$$p(x, y) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2}(y - x)^2\right]; \quad x, y \in \mathbb{R} \quad (16.1.46)$$

2. For $x \in \mathbb{R}$, $p^n(x, \cdot)$ is the normal PDF with mean x and variance n :

$$p^n(x, y) = \frac{1}{\sqrt{2\pi n}} \exp\left[-\frac{1}{2n}(y - x)^2\right], \quad x, y \in \mathbb{R} \quad (16.1.47)$$

In continuous time, there are two processes that are particularly important, one with the discrete state space \mathbb{N} and one with the continuous state space \mathbb{R} .

For $t \in [0, \infty)$, let g_t denote the probability density function of the Poisson distribution with parameter t , and let $p_t(x, y) = g_t(y - x)$ for $x, y \in \mathbb{N}$. Then $\{p_t : t \in [0, \infty)\}$ is the collection of transition densities for a Feller semigroup on \mathbb{N} .

Proof

Recall that

$$g_t(n) = e^{-t} \frac{t^n}{n!}, \quad n \in \mathbb{N} \quad (16.1.48)$$

We just need to show that $\{g_t : t \in [0, \infty)\}$ satisfies the semigroup property, and that the continuity result holds. But we already know that if U, V are independent variables having Poisson distributions with parameters $s, t \in [0, \infty)$, respectively, then $U + V$ has the Poisson distribution with parameter $s + t$. That is, $g_s * g_t = g_{s+t}$. Moreover, $g_t \rightarrow g_0$ as $t \downarrow 0$.

So a Lévy process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ with these transition densities would be a Markov process with stationary, independent increments and with sample paths that are right continuous and have left limits. We do know of such a process, namely the Poisson process with rate 1.

Open the Poisson experiment and set the rate parameter to 1 and the time parameter to 10. Run the experiment several times in single-step mode and note the behavior of the process.

For $t \in (0, \infty)$, let g_t denote the probability density function of the normal distribution with mean 0 and variance t , and let $p_t(x, y) = g_t(y - x)$ for $x, y \in \mathbb{R}$. Then $\{p_t : t \in [0, \infty)\}$ is the collection of transition densities of a Feller semigroup on \mathbb{R} .

Proof

Recall that for $t \in (0, \infty)$,

$$g_t(z) = \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{z^2}{2t}\right), \quad z \in \mathbb{R} \quad (16.1.49)$$

We just need to show that $\{g_t : t \in [0, \infty)\}$ satisfies the semigroup property, and that the continuity result holds. But we already know that if U, V are independent variables having normal distributions with mean 0 and variances $s, t \in (0, \infty)$, respectively, then $U + V$ has the normal distribution with mean 0 and variance $s + t$. That is, $g_s * g_t = g_{s+t}$. Moreover, we also know that the normal distribution with variance t converges to point mass at 0 as $t \downarrow 0$.

So a Lévy process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on \mathbb{R} with these transition densities would be a Markov process with stationary, independent increments, and whose sample paths are continuous from the right and have left limits. In fact, there exists such a process with *continuous* sample paths. This process is Brownian motion, a process important enough to have its own chapter.

Run the simulation of standard Brownian motion and note the behavior of the process.

This page titled [16.1: Introduction to Markov Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.2: Potentials and Generators for General Markov Processes

Our goal in this section is to continue the broad sketch of the general theory of Markov processes. As with the last section, some of the statements are not completely precise and rigorous, because we want to focus on the main ideas without being overly burdened by technicalities. If you are a new student of probability, or are primarily interested in applications, you may want to skip ahead to the study of discrete-time Markov chains.

Preliminaries

Basic Definitions

As usual, our starting point is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so that Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . The set of times T is either \mathbb{N} , discrete time with the discrete topology, or $[0, \infty)$, continuous time with the usual Euclidean topology. The time set T is given the Borel σ -algebra \mathcal{T} , which is just the power set if $T = \mathbb{N}$, and then the time space (T, \mathcal{T}) is given the usual measure, counting measure in the discrete case and Lebesgue measure in the continuous case. The set of states S has an LCCB topology (locally compact, Hausdorff, with a countable base), and is also given the Borel σ -algebra \mathcal{S} . Recall that to say that the state space is *discrete* means that S is countable with the discrete topology, so that \mathcal{S} is the power set of S . The topological assumptions mean that the state space (S, \mathcal{S}) is nice enough for a rich mathematical theory and general enough to encompass the most important applications. There is often a natural Borel measure λ on (S, \mathcal{S}) , counting measure $\#$ if S is discrete, and for example, Lebesgue measure if $S = \mathbb{R}^k$ for some $k \in \mathbb{N}_+$.

Recall also that there are several spaces of functions on S that are important. Let \mathcal{B} denote the set of bounded, measurable functions $f : S \rightarrow \mathbb{R}$. Let \mathcal{C} denote the set of bounded, continuous functions $f : S \rightarrow \mathbb{R}$, and let \mathcal{C}_0 denote the set of continuous functions $f : S \rightarrow \mathbb{R}$ that *vanish at ∞* in the sense that for every $\epsilon > 0$, there exists a compact set $K \subseteq S$ such that $|f(x)| < \epsilon$ for $x \in K^c$. These are all vector spaces under the usual (pointwise) addition and scalar multiplication, and $\mathcal{C}_0 \subseteq \mathcal{C} \subseteq \mathcal{B}$. The supremum norm, defined by $\|f\| = \sup\{|f(x)| : x \in S\}$ for $f \in \mathcal{B}$ is the norm that is used on these spaces.

Suppose now that $\mathbf{X} = \{X_t : t \in T\}$ is a time-homogeneous Markov process with state space (S, \mathcal{S}) defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. As before, we also assume that we have a *filtration* $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, that is, an increasing family of sub σ -algebras of \mathcal{F} , indexed by the time space, with the properties that X_t is measurable with respect to \mathcal{F}_t for $t \in T$. Intuitively, \mathcal{F}_t is the collection of events up to time $t \in T$.

As usual, we let P_t denote the transition probability kernel for an increase in time of size $t \in T$. Thus

$$P_t(x, A) = \mathbb{P}(X_t \in A \mid X_0 = x), \quad x \in S, A \in \mathcal{S} \quad (16.2.1)$$

Recall that for $t \in T$, the transition kernel P_t defines two operators, on the *left* with measures and on the *right* with functions. So, if μ is a measure on (S, \mathcal{S}) then μP_t is the measure on (S, \mathcal{S}) given by

$$\mu P_t(A) = \int_S \mu(dx) P_t(x, A), \quad A \in \mathcal{S} \quad (16.2.2)$$

If μ is the distribution of X_0 then μP_t is the distribution of X_t for $t \in T$. If $f \in \mathcal{B}$ then $P_t f \in \mathcal{B}$ is defined by

$$P_t f(x) = \int_S P_t(x, dy) f(y) = \mathbb{E}[f(X_t) \mid X_0 = x] \quad (16.2.3)$$

Recall that the collection of transition operators $\mathbf{P} = \{P_t : t \in T\}$ is a semigroup because $P_s P_t = P_{s+t}$ for $s, t \in T$. Just about everything in this section is defined in terms of the semigroup \mathbf{P} , which is one of the main analytic tools in the study of Markov processes.

Feller Markov Processes

We make the same assumptions as in the Introduction. Here is a brief review:

We assume that the Markov process $\mathbf{X} = \{X_t : t \in T\}$ satisfies the following properties (and hence is a *Feller Markov process*):

1. For $t \in T$ and $y \in S$, the distribution of X_t given $X_0 = x$ converges to the distribution of X_t given $X_0 = y$ as $x \rightarrow y$.
2. Given $X_0 = x \in S$, X_t converges in probability to x as $t \downarrow 0$.

Part (a) is an assumption on *continuity in space*, while part (b) is an assumption on *continuity in time*. If S is discrete then (a) automatically holds, and if T is discrete then (b) automatically holds. As we will see, the Feller assumptions are sufficient for a very nice mathematical theory, and yet are general enough to encompass the most important continuous-time Markov processes.

The process $\mathbf{X} = \{X_t : t \in T\}$ has the following properties:

1. There is a version of \mathbf{X} such that $t \mapsto X_t$ is continuous from the right and has left limits.
2. \mathbf{X} is a strong Markov process relative to the \mathfrak{F}_+^0 , the right-continuous refinement of the natural filtration.

The Feller assumptions on the Markov process have equivalent formulations in terms of the transition semigroup.

The transition semigroup $\mathbf{P} = \{P_t : t \in T\}$ has the following properties:

1. If $f \in \mathcal{C}_0$ and $t \in T$ then $P_t f \in \mathcal{C}_0$
2. If $f \in \mathcal{C}_0$ and $x \in S$ then $P_t f(x) \rightarrow f(x)$ as $t \downarrow 0$.

As before, part (a) is a condition on *continuity in space*, while part (b) is a condition on *continuity in time*. Once again, (a) is trivial if S is discrete, and (b) trivial if T is discrete. The first condition means that P_t is a linear operator on \mathcal{C}_0 (as well as being a linear operator on \mathcal{B}). The second condition leads to a stronger continuity result.

For $f \in \mathcal{C}_0$, the mapping $t \mapsto P_t f$ is continuous on T . That is, for $t \in T$,

$$\|P_s f - P_t f\| = \sup\{|P_s f(x) - P_t f(x)| : x \in S\} \rightarrow 0 \text{ as } s \rightarrow t \quad (16.2.4)$$

Our interest in this section is primarily the continuous time case. However, we start with the discrete time case since the concepts are clearer and simpler, and we can avoid some of the technicalities that inevitably occur in continuous time.

Discrete Time

Suppose that $T = \mathbb{N}$, so that time is discrete. Recall that the transition kernels are just powers of the one-step kernel. That is, we let $P = P_1$ and then $P_n = P^n$ for $n \in \mathbb{N}$.

Potential Operators

For $\alpha \in (0, 1]$, the α -potential kernel R_α of \mathbf{X} is defined as follows:

$$R_\alpha(x, A) = \sum_{n=0}^{\infty} \alpha^n P^n(x, A), \quad x \in S, A \in \mathcal{S} \quad (16.2.5)$$

1. The special case $R = R_1$ is simply the *potential kernel* of \mathbf{X} .
2. For $x \in S$ and $A \in \mathcal{S}$, $R(x, A)$ is the expected number of visits of \mathbf{X} to A , starting at x .

Proof

The function $x \mapsto R_\alpha(x, A)$ from S to $[0, \infty)$ is measurable for $A \in \mathcal{S}$ since $x \mapsto P^n(x, A)$ is measurable for each $n \in \mathbb{N}$. The mapping $A \mapsto R_\alpha(x, A)$ is a positive measure on \mathcal{S} for $x \in S$ since $A \mapsto P^n(x, A)$ is a probability measure for each $n \in \mathbb{N}$. Finally, the interpretation of $R(x, A)$ for $x \in S$ and $A \in \mathcal{S}$ comes from interchanging sum and expected value, which is allowed since the terms are nonnegative:

$$R(x, A) = \sum_{n=0}^{\infty} P^n(x, A) = \sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}(X_n \in A) \mid X_0 = x] = \mathbb{E}\left(\sum_{n=0}^{\infty} \mathbf{1}(X_n \in A) \mid X_0 = x\right) = \mathbb{E}[\#\{n \in \mathbb{N} : X_n \in A\} \mid X_0 = x] \quad (16.2.6)$$

Note that it's quite possible that $R(x, A) = \infty$ for some $x \in S$ and $A \in \mathcal{S}$. In fact, knowing when this is the case is of considerable importance in the study of Markov processes. As with all kernels, the potential kernel R_α defines two operators, operating on the right on functions, and operating on the left on positive measures. For the right potential operator, if $f : S \rightarrow \mathbb{R}$ is measurable then

$$R_\alpha f(x) = \sum_{n=0}^{\infty} \alpha^n P^n f(x) = \sum_{n=0}^{\infty} \alpha^n \int_S P^n(x, dy) f(y) = \sum_{n=0}^{\infty} \alpha^n \mathbb{E}[f(X_n) \mid X_0 = x], \quad x \in S \quad (16.2.7)$$

assuming as usual that the expected values and the infinite series make sense. This will be the case, in particular, if f is nonnegative or if $p \in (0, 1)$ and $f \in \mathcal{B}$.

If $\alpha \in (0, 1)$, then $R_\alpha(x, S) = \frac{1}{1-\alpha}$ for all $x \in S$.

Proof

Using geometric series,

$$R_\alpha(x, S) = \sum_{n=0}^{\infty} \alpha^n P^n(x, S) = \sum_{n=0}^{\infty} \alpha^n = \frac{1}{1-\alpha} \quad (16.2.8)$$

It follows that for $\alpha \in (0, 1)$, the right operator R_α is a bounded, linear operator on \mathcal{B} with $\|R_\alpha\| = \frac{1}{1-\alpha}$. It also follows that $(1-\alpha)R_\alpha$ is a probability kernel. There is a nice interpretation of this kernel.

If $\alpha \in (0, 1)$ then $(1-\alpha)R_\alpha(x, \cdot)$ is the conditional distribution of X_N given $X_0 = x \in S$, where N is independent of \mathbf{X} and has the geometric distribution on \mathbb{N} with parameter $1-\alpha$.

Proof

Suppose that $x \in S$ and $A \in \mathcal{S}$. Conditioning on N gives

$$\mathbb{P}(X_N \in A \mid X_0 = x) = \sum_{n=0}^{\infty} \mathbb{P}(N = n) \mathbb{P}(X_n \in A \mid N = n, X_0 = x) \quad (16.2.9)$$

But by the substitution rule and the assumption of independence,

$$\mathbb{P}(X_N \in A \mid N = n, X_0 = x) = \mathbb{P}(X_n \in A \mid N = n, X_0 = x) = \mathbb{P}(X_n \in A \mid X_0 = x) = P^n(x, A) \quad (16.2.10)$$

Since N has the geometric distribution on \mathbb{N} with parameter $1-\alpha$ we have $P(N = n) = (1-\alpha)\alpha^n$ for $n \in \mathbb{N}$. Substituting gives

$$\mathbb{P}(X_N \in A \mid X_0 = x) = \sum_{n=0}^{\infty} (1 - \alpha) \alpha^n P^n(x, A) = (1 - \alpha) R_\alpha(x, A) \quad (16.2.11)$$

So $(1 - \alpha)R_\alpha$ is a transition probability kernel, just as P_n is a transition probability kernel, but corresponding to the *random* time N (with $\alpha \in (0, 1)$ as a parameter), rather than the *deterministic* time $n \in \mathbb{N}$. An interpretation of the potential kernel R_α for $\alpha \in (0, 1)$ can be also given in economic terms. Suppose that $A \in \mathcal{S}$ and that we receive one monetary unit each time the process \mathbf{X} visits A . Then as above, $R(x, A)$ is the expected total amount of money we receive, starting at $x \in S$. However, typically money that we will receive at times distant in the future has less value to us *now* than money that we will receive soon. Specifically suppose that a monetary unit received at time $n \in \mathbb{N}$ has a *present value* of α^n , where $\alpha \in (0, 1)$ is an *inflation factor* (sometimes also called a *discount factor*). Then $R_\alpha(x, A)$ gives the expected, total, discounted amount we will receive, starting at $x \in S$. A bit more generally, if $f \in \mathcal{B}$ is a *reward function*, so that $f(x)$ is the reward (or cost, depending on the sign) that we receive when we visit state $x \in S$, then for $\alpha \in (0, 1)$, $R_\alpha f(x)$ is the expected, total, discounted reward, starting at $x \in S$.

For the left potential operator, if μ is a positive measure on \mathcal{S} then

$$\mu R_\alpha(A) = \sum_{n=0}^{\infty} \alpha^n \mu P^n(A) = \sum_{n=0}^{\infty} \alpha^n \int_S \mu(dx) P^n(x, A), \quad A \in \mathcal{S} \quad (16.2.12)$$

In particular, if μ is a probability measure and X_0 has distribution μ then μP^n is the distribution of X_n for $n \in \mathbb{N}$, so from the last result, $(1 - \alpha)\mu R_\alpha$ is the distribution of X_N where again, N is independent of \mathbf{X} and has the geometric distribution on \mathbb{N} with parameter $1 - \alpha$. The family of potential kernels gives the same information as the family of transition kernels.

The potential kernels $\mathbf{R} = \{R_\alpha : \alpha \in (0, 1)\}$ completely determine the transition kernels $\mathbf{P} = \{P_n : n \in \mathbb{N}\}$.

Proof

Note that for $x \in S$ and $A \in \mathcal{S}$, the function $\alpha \mapsto R_\alpha(x, A)$ is a power series in α with coefficients $n \mapsto P^n(x, A)$. In the language of combinatorics, $\alpha \mapsto R_\alpha(x, A)$ is the *ordinary generating function* of the sequence $n \mapsto P^n(x, A)$. As noted above, this power series has radius of convergence at least 1, so we can extend the domain to $\alpha \in (-1, 1)$. Thus, given the potential kernels, we can recover the transition kernels by taking derivatives and evaluating at 0:

$$P^n(x, A) = \frac{1}{n!} \left[\frac{d^n}{d\alpha^n} R_\alpha(x, A) \right]_{\alpha=0} \quad (16.2.13)$$

Of course, it's really only necessary to determine P , the one step transition kernel, since the other transition kernels are powers of P . In any event, it follows that the kernels $\mathbf{R} = \{R_\alpha : \alpha \in (0, 1)\}$, along with the initial distribution, completely determine the finite dimensional distributions of the Markov process \mathbf{X} . The potential kernels commute with each other and with the transition kernels.

Suppose that $\alpha, \beta \in (0, 1]$ and $k \in \mathbb{N}$. Then (as kernels)

1. $P^k R_\alpha = R_\alpha P^k = \sum_{n=0}^{\infty} \alpha^n P^{n+k}$
2. $R_\alpha R_\beta = R_\beta R_\alpha = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \alpha^m \beta^n P^{m+n}$

Proof

Suppose that $f \in \mathcal{B}$ is nonnegative. The interchange of the sums with the kernel operation is allowed since the kernels are nonnegative. The other tool used is the semigroup property.

1. Directly

$$R_\alpha P^k f = \sum_{n=0}^{\infty} \alpha^n P^n P^k f = \sum_{n=0}^{\infty} \alpha^n P^{n+k} f \quad (16.2.14)$$

The other direction requires an interchange.

$$P^k R_\alpha f = P^k \sum_{n=0}^{\infty} \alpha^n P^n f = \sum_{n=0}^{\infty} \alpha^n P^k P^n f = \sum_{n=0}^{\infty} \alpha^n P^{n+k} f \quad (16.2.15)$$

2. First,

$$R_\alpha R_\beta f = \sum_{m=0}^{\infty} \alpha^m P^m R_\beta f = \sum_{m=0}^{\infty} \alpha^m P^m \left(\sum_{n=0}^{\infty} \beta^n P^n f \right) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \alpha^m \beta^n P^m P^n f = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \alpha^m \beta^n P^{m+n} f \quad (16.2.16)$$

The other direction is similar.

The same identities hold for the right operators on the entire space \mathcal{B} , with the additional restrictions that $\alpha < 1$ and $\beta < 1$. The fundamental equation that relates the potential kernels is given next.

If $\alpha, \beta \in (0, 1]$ with $\alpha \leq \beta$ then (as kernels),

$$\beta R_\beta = \alpha R_\alpha + (\beta - \alpha) R_\alpha R_\beta \quad (16.2.17)$$

Proof

If $\alpha = \beta$ the equation is trivial, so assume $\alpha < \beta$. Suppose that $f \in \mathcal{B}$ is nonnegative. From the previous result,

$$R_\alpha R_\beta f = \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} \alpha^j \beta^k P^{j+k} f \quad (16.2.18)$$

Changing variables to sum over $n = j + k$ and j gives

$$R_\alpha R_\beta f = \sum_{n=0}^{\infty} \sum_{j=0}^n \alpha^j \beta^{n-j} P^n f = \sum_{n=0}^{\infty} \sum_{j=0}^n \left(\frac{\alpha}{\beta}\right)^j \beta^n P^n f = \sum_{n=0}^{\infty} \frac{1 - \left(\frac{\alpha}{\beta}\right)^{n+1}}{1 - \frac{\alpha}{\beta}} \beta^n P^n f \quad (16.2.19)$$

Simplifying gives

$$R_\alpha R_\beta f = \frac{1}{\beta - \alpha} (\beta R_\beta f - \alpha R_\alpha f) \quad (16.2.20)$$

Note that since $\alpha < 1$, $R_\alpha f$ is a finite, so we don't have to worry about the dreaded indeterminate form $\infty - \infty$.

The same identity holds for the right operators on the entire space \mathcal{B} , with the additional restriction that $\beta < 1$.

If $\alpha \in (0, 1]$, then (as kernels), $I + \alpha R_\alpha P = I + \alpha P R_\alpha = R_\alpha$.

Proof

Suppose that $f \in \mathcal{B}$ is nonnegative. From the result above,

$$(I + \alpha R_\alpha P)f = (I + \alpha P R_\alpha)f = f + \sum_{n=0}^{\infty} \alpha^{n+1} P^{n+1} f = \sum_{n=0}^{\infty} \alpha^n P^n f = R_\alpha f \quad (16.2.21)$$

The same identity holds for the right operators on the entire space \mathcal{B} , with the additional restriction that $\alpha < 1$. This leads to the following important result:

If $\alpha \in (0, 1)$, then as operators on the space \mathcal{B} ,

1. $R_\alpha = (I - \alpha P)^{-1}$
2. $P = \frac{1}{\alpha} (I - R_\alpha^{-1})$

Proof

The operators are bounded, so we can subtract. The identity $I + \alpha R_\alpha P = R_\alpha$ leads to $R_\alpha (I - \alpha P) = I$ and the identity $I + \alpha P R_\alpha = R_\alpha$ leads to $(I - \alpha P) R_\alpha = I$. Hence (a) holds. Part (b) follows from (a).

This result shows again that the potential operator R_α determines the transition operator P .

Examples and Applications

Our first example considers the binomial process as a Markov process.

Let $\mathbf{I} = \{I_n : n \in \mathbb{N}_+\}$ be a sequence of Bernoulli trials with success parameter $p \in (0, 1)$. Define the Markov process $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ by $X_n = X_0 + \sum_{k=1}^n I_k$ where X_0 takes values in \mathbb{N} and is independent of \mathbf{I} .

1. For $n \in \mathbb{N}$, show that the transition probability matrix P^n of \mathbf{X} is given by

$$P^n(x, y) = \binom{n}{y-x} p^{y-x} (1-p)^{n-y+x}, \quad x \in \mathbb{N}, y \in \{x, x+1, \dots, x+n\} \quad (16.2.22)$$

2. For $\alpha \in (0, 1]$, show that the potential matrix R_α of \mathbf{X} is given by

$$R_\alpha(x, y) = \frac{1}{1 - \alpha + \alpha p} \left(\frac{\alpha p}{1 - \alpha + \alpha p} \right)^{y-x}, \quad x \in \mathbb{N}, y \in \{x, x+1, \dots\} \quad (16.2.23)$$

3. For $\alpha \in (0, 1)$ and $x \in \mathbb{N}$, identify the probability distribution defined by $(1 - \alpha) R_\alpha(x, \cdot)$.
4. For $x, y \in \mathbb{N}$ with $x \leq y$, interpret $R(x, y)$, the expected time in y starting in x , in the context of the process \mathbf{X} .

Solutions

Recall that \mathbf{X} is a Markov process since it has stationary, independent increments.

1. Note that for $n, x \in \mathbb{N}$, $P^n(x, \cdot)$ is the (discrete) PDF of $x + \sum_{k=1}^n I_k$. The result follows since the sum of the indicator variables has the binomial distribution with parameters n and p .
2. Let $\alpha \in (0, 1]$ and let $x, y \in \mathbb{N}$ with $x \leq y$. Then

$$\begin{aligned} R_\alpha(x, y) &= \sum_{n=0}^{\infty} \alpha^n P^n(x, y) = \sum_{n=y-x}^{\infty} \alpha^n \binom{n}{y-x} p^{y-x} (1-p)^{n-y+x} \\ &= (\alpha p)^{y-x} \sum_{n=y-x}^{\infty} \binom{n}{y-x} [\alpha(1-p)]^{n-y+x} = \frac{(\alpha p)^{y-x}}{[1-\alpha(1-p)]^{n-x+1}} \end{aligned}$$

Simplifying gives the result.

3. For $\alpha \in (0, 1)$,

$$(1-\alpha)R_\alpha(x, y) = \frac{1-\alpha}{1-\alpha+\alpha p} \left(\frac{\alpha p}{1-\alpha+\alpha p} \right)^{y-x} \quad (16.2.24)$$

As a function of y for fixed x , this is the PDF of $x + Y_\alpha$ where Y_α has the geometric distribution on \mathbb{N} with parameter $\frac{1-\alpha}{1-\alpha+\alpha p}$.

4. Note that $R(x, y) = 1/p$ for $x, y \in \mathbb{N}$ with $x \leq y$. Starting in state x , the process eventually reaches y with probability 1. The process remains in state y for a geometrically distributed time, with parameter p . The mean of this distribution is $1/p$.

Continuous Time

With the discrete-time setting as motivation, we now turn the more important continuous-time case where $T = [0, \infty)$.

Potential Kernels

For $\alpha \in [0, \infty)$, the α -potential kernel U_α of \mathbf{X} is defined as follows:

$$U_\alpha(x, A) = \int_0^\infty e^{-\alpha t} P_t(x, A) dt, \quad x \in S, A \in \mathcal{S} \quad (16.2.25)$$

1. The special case $U = U_0$ is simply the *potential kernel* of \mathbf{X} .
2. For $x \in S$ and $A \in \mathcal{S}$, $U(x, A)$ is the expected amount of time that \mathbf{X} spends in A , starting at x .
3. The family of kernels $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$ is known as the *resolvent* of \mathbf{X} .

Proof

Since $\mathbf{P} = \{P_t : t \in T\}$ is a Feller semigroup of transition operators, the mapping $(t, x) \mapsto P_t(x, A)$ from $[0, \infty) \times S$ to $[0, 1]$ is jointly measurable for $A \in \mathcal{S}$. Thus, $U_\alpha(x, A)$ makes sense for $x \in S$ and $A \in \mathcal{S}$ and $x \mapsto U_\alpha(x, A)$ from S to $[0, \infty)$ is measurable for $A \in \mathcal{S}$. That $A \mapsto U_\alpha(x, A)$ is a measure on \mathcal{S} follows from the usual interchange of sum and integral, via Fubini's theorem: Suppose that $\{A_j : j \in J\}$ is a countable collection of disjoint sets in \mathcal{S} , and let $S = \bigcup_{j \in J} A_j$

$$\begin{aligned} U_\alpha(x, A) &= \int_0^\infty e^{-\alpha t} P_t(x, A) dt = \int_0^\infty \left[\sum_{j \in J} e^{-\alpha t} P_t(x, A_j) \right] dt \\ &= \sum_{j \in J} \int_0^\infty e^{-\alpha t} P_t(x, A_j) dt = \sum_{j \in J} U_\alpha(x, A_j) \end{aligned}$$

Finally, the interpretation of $U(x, A)$ for $x \in S$ and $A \in \mathcal{S}$ is another interchange of integrals:

$$U(x, A) = \int_0^\infty P_t(x, A) dt = \int_0^\infty \mathbb{E}[\mathbf{1}(X_t \in A) \mid X_0 = x] dt = \mathbb{E} \left(\int_0^\infty \mathbf{1}(X_t \in A) dt \mid X_0 = x \right) \quad (16.2.26)$$

The inside integral is the Lebesgue measure of $\{t \in [0, \infty) : X_t \in A\}$.

As with discrete time, it's quite possible that $U(x, A) = \infty$ for some $x \in S$ and $A \in \mathcal{S}$, and knowing when this is the case is of considerable interest. As with all kernels, the potential kernel U_α defines two operators, operating on the right on functions, and operating on the left on positive measures. If $f : S \rightarrow \mathbb{R}$ is measurable then, giving the right potential operator in its many forms,

$$\begin{aligned} U_\alpha f(x) &= \int_S U_\alpha(x, dy) f(y) = \int_0^\infty e^{-\alpha t} P_t f(x) dt \\ &= \int_0^\infty e^{-\alpha t} \int_S P_t(x, dy) f(y) = \int_0^\infty e^{-\alpha t} \mathbb{E}[f(X_t) \mid X_0 = x] dt, \quad x \in S \end{aligned}$$

assuming that the various integrals make sense. This will be the case in particular if f is nonnegative, or if $f \in \mathcal{B}$ and $\alpha > 0$.

If $\alpha > 0$, then $U_\alpha(x, S) = \frac{1}{\alpha}$ for all $x \in S$.

Proof

For $x \in S$,

$$U_\alpha(x, S) = \int_0^\infty e^{-\alpha t} P_t(x, S) dt = \int_0^\infty e^{-\alpha t} dt = \frac{1}{\alpha} \quad (16.2.27)$$

It follows that for $\alpha \in (0, \infty)$, the right potential operator U_α is a bounded, linear operator on \mathcal{B} with $\|U_\alpha\| = \frac{1}{\alpha}$. It also follows that αU_α is a probability kernel. This kernel has a nice interpretation.

If $\alpha > 0$ then $\alpha U_\alpha(x, \cdot)$ is the conditional distribution of X_τ where τ is independent of \mathbf{X} and has the exponential distribution on $[0, \infty)$ with parameter α .

Proof

Suppose that $x \in S$ and $A \in \mathcal{S}$. The random time τ has PDF $f(t) = \alpha e^{-\alpha t}$ for $t \in [0, \infty)$. Hence, conditioning on τ gives

$$\mathbb{P}(X_\tau \in A \mid X_0 = x) = \int_0^\infty \alpha e^{-\alpha t} \mathbb{P}(X_\tau \in A \mid \tau = t, X_0 = x) dt \quad (16.2.28)$$

But by the substitution rule and the assumption of independence,

$$\mathbb{P}(X_\tau \in A \mid \tau = t, X_0 = x) = \mathbb{P}(X_t \in A \mid \tau = t, X_0 = x) = \mathbb{P}(X_t \in A \mid X_0 = x) = P_t(x, A) \quad (16.2.29)$$

Substituting gives

$$\mathbb{P}(X_\tau \in A \mid X_0 = x) = \int_0^\infty \alpha e^{-\alpha t} P_t(x, A) dt = \alpha U_\alpha(x, A) \quad (16.2.30)$$

So αU_α is a transition probability kernel, just as P_t is a transition probability kernel, but corresponding to the *random* time τ (with $\alpha \in (0, \infty)$ as a parameter), rather than the deterministic time $t \in [0, \infty)$. As in the discrete case, the potential kernel can also be interpreted in economic terms. Suppose that $A \in \mathcal{S}$ and that we receive money at a rate of one unit per unit time whenever the process \mathbf{X} is in A . Then $U(x, A)$ is the expected total amount of money that we receive, starting in state $x \in S$. But again, money that we receive later is of less value to us *now* than money that we will receive sooner. Specifically, suppose that one monetary unit at time $t \in [0, \infty)$ has a present value of $e^{-\alpha t}$ where $\alpha \in (0, \infty)$ is the *inflation factor* or *discount factor*. The $U_\alpha(x, A)$ is the total, expected, discounted amount that we receive, starting in $x \in S$. A bit more generally, suppose that $f \in \mathcal{B}$ and that $f(x)$ is the reward (or cost, depending on the sign) per unit time that we receive when the process is in state $x \in S$. Then $U_\alpha f(x)$ is the expected, total, discounted reward, starting in state $x \in S$.

For the left potential operator, if μ is a positive measure on \mathcal{S} then

$$\begin{aligned} \mu U_\alpha(A) &= \int_S \mu(dx) U_\alpha(x, A) = \int_0^\infty e^{-\alpha t} \mu P_t(A) dt \\ &= \int_0^\infty e^{-\alpha t} \left[\int_S \mu(dx) P_t(x, A) \right] dt = \int_0^\infty e^{-\alpha t} \left[\int_S \mu(dx) \mathbb{P}(X_t \in A) \right] dt, \quad A \in \mathcal{S} \end{aligned}$$

In particular, suppose that $\alpha > 0$ and that μ is a probability measure and X_0 has distribution μ . Then μP_t is the distribution of X_t for $t \in [0, \infty)$, and hence from the last result, $\alpha \mu U_\alpha$ is the distribution of X_τ , where again, τ is independent of \mathbf{X} and has the exponential distribution on $[0, \infty)$ with parameter α . The family of potential kernels gives the same information as the family of transition kernels.

The resolvent $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$ completely determines the family of transition kernels $\mathbf{P} = \{P_t : t \in (0, \infty)\}$.

Proof

Note that for $x \in S$ and $A \in \mathcal{S}$, the function $\alpha \mapsto U_\alpha(x, A)$ on $(0, \infty)$ is the Laplace transform of the function $t \mapsto P_t(x, A)$ on $[0, \infty)$. The Laplace transform of a function determines the function completely.

It follows that the resolvent $\{U_\alpha : \alpha \in [0, \infty)\}$, along with the initial distribution, completely determine the finite dimensional distributions of the Markov process \mathbf{X} . This is much more important here in the continuous-time case than in the discrete-time case, since the transition kernels P_t cannot be generated from a single transition kernel. The potential kernels commute with each other and with the transition kernels.

Suppose that $\alpha, \beta, t \in [0, \infty)$. Then (as kernels),

1. $P_t U_\alpha = U_\alpha P_t = \int_0^\infty e^{-\alpha s} P_{s+t} ds$
2. $U_\alpha U_\beta = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_{s+t} ds dt$

Proof

Suppose that $f \in \mathcal{B}$ is nonnegative. The interchanges of operators and integrals below are interchanges of integrals, and are justified since the integrands are nonnegative. The other tool used is the semigroup property of $\mathbf{P} = \{P_t : t \in [0, \infty)\}$.

1. Directly,

$$U_\alpha P_t f = \int_0^\infty e^{-\alpha s} P_s P_t f ds = \int_0^\infty e^{-\alpha s} P_{s+t} f ds \quad (16.2.31)$$

The other direction involves an interchange.

$$P_t U_\alpha f = P_t \int_0^\infty e^{-\alpha s} P_s f ds = \int_0^\infty e^{-\alpha s} P_t P_s f ds = \int_0^\infty e^{-\alpha s} P_{s+t} f ds \quad (16.2.32)$$

2. First

$$\begin{aligned} U_\alpha U_\beta f &= \int_0^\infty e^{-\alpha s} P_s U_\beta f ds = \int_0^\infty e^{-\alpha s} P_s \int_0^\infty e^{-\beta t} P_t f dt ds \\ &= \int_0^\infty e^{-\alpha s} \int_0^\infty e^{-\beta t} P_s P_t f ds dt = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_{s+t} f ds dt \end{aligned}$$

The other direction is similar.

The same identities hold for the right operators on the entire space \mathcal{B} under the additional restriction that $\alpha > 0$ and $\beta > 0$. The fundamental equation that relates the potential kernels, known as the *resolvent equation*, is given in the next theorem:

If $\alpha, \beta \in [0, \infty)$ with $\alpha \leq \beta$ then (as kernels) $U_\alpha = U_\beta + (\beta - \alpha)U_\alpha U_\beta$.

Proof

If $\alpha = \beta$ the equation is trivial, so assume $\alpha < \beta$. Suppose that $f \in \mathcal{B}$ is nonnegative. From the previous result,

$$U_\alpha U_\beta f = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_{s+t} f dt ds \quad (16.2.33)$$

The transformation $u = s + t, v = s$ maps $[0, \infty)^2$ one-to-one onto $\{(u, v) \in [0, \infty)^2 : u \geq v\}$. The inverse transformation is $s = v, t = u - v$ with Jacobian -1 . Hence we have

$$\begin{aligned} U_\alpha U_\beta f &= \int_0^\infty \int_0^u e^{-\alpha v} e^{-\beta(u-v)} P_u f dv du = \int_0^\infty \left(\int_0^u e^{-(\beta-\alpha)v} dv \right) e^{-\beta u} P_u f du \\ &= \frac{1}{\beta - \alpha} \int_0^\infty [e^{(\beta-\alpha)u} - 1] e^{-\beta u} P_u f du \\ &= \frac{1}{\beta - \alpha} \left(\int_0^\infty e^{-\alpha u} P_u f du - \int_0^\infty e^{-\beta u} P_u f du \right) = \frac{1}{\beta - \alpha} (U_\alpha f - U_\beta f) \end{aligned}$$

Simplifying gives the result. Note that $U_\beta f$ is finite since $\beta > 0$.

The same identity holds for the right potential operators on the entire space \mathcal{B} , under the additional restriction that $\alpha > 0$. For $\alpha \in (0, \infty)$, U_α is also an operator on the space \mathcal{C}_0 .

If $\alpha \in (0, \infty)$ and $f \in \mathcal{C}_0$ then $U_\alpha f \in \mathcal{C}_0$.

Proof

Suppose that $f \in \mathcal{C}_0$ and that (x_1, x_2, \dots) is a sequence in S . Then $P_t f \in \mathcal{C}_0$ for $t \in [0, \infty)$. Hence if $x_n \rightarrow x \in S$ as $n \rightarrow \infty$ then $e^{-\alpha t} P_t f(x_n) \rightarrow e^{-\alpha t} P_t f(x)$ as $n \rightarrow \infty$ for each $t \in [0, \infty)$. By the dominated convergence theorem,

$$U_\alpha f(x_n) = \int_0^\infty e^{-\alpha t} P_t f(x_n) dt \rightarrow \int_0^\infty e^{-\alpha t} P_t f(x) dt = U_\alpha f(x) \text{ as } n \rightarrow \infty \quad (16.2.34)$$

Hence $U_\alpha f$ is continuous. Next suppose that $x_n \rightarrow \infty$ as $n \rightarrow \infty$. This means that for every compact $C \subseteq S$, there exist $m \in \mathbb{N}_+$ such that $x_n \notin C$ for $n > m$. Then $e^{-\alpha t} P_t f(x_n) \rightarrow 0$ as $n \rightarrow \infty$ for each $t \in [0, \infty)$. Again by the dominated convergence theorem,

$$U_\alpha f(x_n) = \int_0^\infty e^{-\alpha t} P_t f(x_n) dt \rightarrow 0 \text{ as } n \rightarrow \infty \quad (16.2.35)$$

So $U_\alpha f \in \mathcal{C}_0$.

If $f \in \mathcal{C}_0$ then $\alpha U_\alpha f \rightarrow f$ as $\alpha \rightarrow \infty$.

Proof

Convergence is with respect to the supremum norm on \mathcal{C}_0 , of course. Suppose that $f \in \mathcal{C}_0$. Note first that with a change of variables $s = \alpha t$,

$$\alpha U_\alpha f = \int_0^\infty \alpha e^{-\alpha t} P_t f dt = \int_0^\infty e^{-s} P_{s/\alpha} f ds \quad (16.2.36)$$

and hence

$$|\alpha U_\alpha f - f| = \left| \int_0^\infty e^{-s} (P_{s/\alpha} f - f) ds \right| \leq \int_0^\infty e^{-s} |P_{s/\alpha} f - f| ds \leq \int_0^\infty e^{-s} \|P_{s/\alpha} f - f\| ds \quad (16.2.37)$$

So it follows that

$$\|\alpha U_\alpha f - f\| \leq \int_0^\infty e^{-s} \|P_{s/\alpha} f - f\| ds \quad (16.2.38)$$

But $\|P_{s/\alpha} f - f\| \rightarrow 0$ as $\alpha \rightarrow \infty$ and hence by the dominated convergence theorem, $\int_0^\infty e^{-s} \|P_{s/\alpha} f - f\| ds \rightarrow 0$ as $\alpha \rightarrow \infty$.

Infinitesimal Generator

In continuous time, it's not at all clear how we could construct a Markov process with desired properties, say to model a real system of some sort. Stated mathematically, the existential problem is how to construct the family of transition kernels $\{P_t : t \in [0, \infty)\}$ so that the semigroup property $P_s P_t = P_{s+t}$ is satisfied for all $s, t \in [0, \infty)$. The answer, as for similar problems in the deterministic world, comes essentially from calculus, from a type of derivative.

The *infinitesimal generator* of the Markov process \mathbf{X} is the operator $G : \mathcal{D} \rightarrow \mathcal{C}_0$ defined by

$$Gf = \lim_{t \downarrow 0} \frac{P_t f - f}{t} \quad (16.2.39)$$

on the domain $\mathcal{D} \subseteq \mathcal{C}_0$ for which the limit exists.

As usual, the limit is with respect to the supremum norm on \mathcal{C}_0 , so $f \in \mathcal{D}$ and $Gf = g$ means that $f, g \in \mathcal{C}_0$ and

$$\left\| \frac{P_t f - f}{t} - g \right\| = \sup \left\{ \left| \frac{P_t f(x) - f(x)}{t} - g(x) \right| : x \in S \right\} \rightarrow 0 \text{ as } t \downarrow 0 \quad (16.2.40)$$

So in particular,

$$Gf(x) = \lim_{t \downarrow 0} \frac{P_t f(x) - f(x)}{t} = \lim_{t \downarrow 0} \frac{\mathbb{E}[f(X_t) | X_0 = x] - f(x)}{t}, \quad x \in S \quad (16.2.41)$$

The domain \mathcal{D} is a subspace of \mathcal{C}_0 and the generator G is a linear operator on \mathcal{D}

1. If $f \in \mathcal{D}$ and $c \in \mathbb{R}$ then $cf \in \mathcal{D}$ and $G(cf) = cGf$.
2. If $f, g \in \mathcal{D}$ then $f + g \in \mathcal{D}$ and $G(f + g) = Gf + Gg$.

Proof

These are simple results that depend on the linearity of P_t for $t \in [0, \infty)$ and basic results on convergence.

1. If $f \in \mathcal{D}$ then

$$\frac{P_t(cf) - (cf)}{t} = c \frac{P_t f - f}{t} \rightarrow cGf \text{ as } t \downarrow 0 \quad (16.2.42)$$

2. If $f, g \in \mathcal{D}$ then

$$\frac{P_t(f + g) - (f + g)}{t} = \frac{P_t f - f}{t} + \frac{P_t g - g}{t} \rightarrow Gf + Gg \text{ as } t \downarrow 0 \quad (16.2.43)$$

Note G is the (right) derivative at 0 of the function $t \mapsto P_t f$. Because of the semigroup property, this differentiability property at 0 implies differentiability at arbitrary $t \in [0, \infty)$. Moreover, the infinitesimal operator and the transition operators commute:

If $f \in \mathcal{D}$ and $t \in [0, \infty)$, then $P_t f \in \mathcal{D}$ and the following derivative rules hold with respect to the supremum norm.

1. $P'_t f = P_t Gf$, the *Kolmogorov forward equation*
2. $P'_t f = G P_t f$, the *Kolmogorov backward equation*

Proof

Let $f \in \mathcal{D}$. All limits and statements about derivatives and continuity are with respect to the supremum norm.

1. By assumption,

$$\frac{1}{h} (P_h f - f) \rightarrow Gf \text{ as } h \downarrow 0 \quad (16.2.44)$$

Since P_t is a bounded, linear operator on the space \mathcal{C}_0 , it preserves limits, so

$$\frac{1}{h} (P_t P_h f - P_t f) = \frac{1}{h} (P_{t+h} f - P_t f) \rightarrow P_t Gf \text{ as } h \downarrow 0 \quad (16.2.45)$$

This proves the result for the derivative from the right. But since $t \mapsto P_t f$ is continuous, the result is also true for the two-sided derivative.

2. From part (a), we now know that

$$\frac{1}{h} (P_h P_t f - P_t f) = \frac{1}{h} (P_{t+h} f - P_t f) \rightarrow P_t Gf \text{ as } h \rightarrow 0 \quad (16.2.46)$$

By definition, this means that $P_t f \in \mathcal{D}$ and $G P_t f = P_t Gf = P'_t f$.

The last result gives a possible solution to the dilemma that motivated this discussion in the first place. If we want to construct a Markov process with desired properties, to model a real system for example, we can start by constructing an appropriate generator G and then solve the initial value problem

$$P'_t = G P_t, \quad P_0 = I \quad (16.2.47)$$

to obtain the transition operators $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. The next theorem gives the relationship between the potential operators and the infinitesimal operator, which in some ways is better. This relationship is analogous to the relationship between the potential operators and the one-step operator given above in discrete time

Suppose $\alpha \in (0, \infty)$.

1. If $f \in \mathcal{D}$ the $Gf \in \mathcal{C}_0$ and $f + U_\alpha Gf = \alpha U_\alpha f$
2. If $f \in \mathcal{C}_0$ then $U_\alpha f \in \mathcal{D}$ and $f + GU_\alpha f = \alpha U_\alpha f$.

Proof

1. By definition, if $f \in \mathcal{D}$ then $Gf \in \mathcal{C}_0$. Hence using the previous result,

$$f + U_\alpha Gf = f + \int_0^\infty e^{-\alpha t} G P_t f dt = f + \int_0^\infty e^{-\alpha t} P_t' f dt \quad (16.2.48)$$

Integrating by parts (with $u = e^{-\alpha t}$ and $dv = P_t' f dt$) gives

$$f + GU_\alpha f = f - e^{-\alpha t} P_t f \Big|_0^\infty + \alpha \int_0^\infty e^{-\alpha t} P_t f dt \quad (16.2.49)$$

But $e^{-\alpha t} P_t f \rightarrow 0$ as $t \rightarrow \infty$ while $P_0 f = f$. The last term is $\alpha U_\alpha f$.

2. Suppose that $f \in \mathcal{C}_0$. From the result above and the substitution $u = s + t$,

$$P_t U_\alpha f = \int_0^\infty e^{-\alpha s} P_{s+t} f ds = \int_t^\infty e^{-\alpha(u-t)} P_u f du = e^{\alpha t} \int_t^\infty e^{-\alpha u} P_u f du \quad (16.2.50)$$

Hence

$$\frac{P_t U_\alpha f - U_\alpha f}{t} = \frac{1}{t} \left[e^{\alpha t} \int_t^\infty e^{-\alpha u} P_u f du - U_\alpha f \right] \quad (16.2.51)$$

Adding and subtracting $e^{\alpha u} U_\alpha f$ and combining integrals gives

$$\begin{aligned} \frac{P_t U_\alpha f - U_\alpha f}{t} &= \frac{1}{t} \left[e^{\alpha t} \int_t^\infty e^{-\alpha u} P_u f du - e^{\alpha t} \int_0^\infty e^{-\alpha u} P_u f du \right] + \frac{e^{\alpha t} - 1}{t} U_\alpha f \\ &= -e^{\alpha t} \frac{1}{t} \int_0^t e^{-\alpha s} P_s f ds + \frac{e^{\alpha t} - 1}{t} U_\alpha f \end{aligned}$$

Since $s \mapsto P_s f$ is continuous, the first term converges to $-f$ as $t \downarrow 0$. The second term converges to $\alpha U_\alpha f$ as $t \downarrow 0$.

For $\alpha > 0$, the operators U_α and G have an inverse relationship.

Suppose again that $\alpha \in (0, \infty)$.

1. $U_\alpha = (\alpha I - G)^{-1} : \mathcal{C}_0 \rightarrow \mathcal{D}$
2. $G = \alpha I - U_\alpha^{-1} : \mathcal{D} \rightarrow \mathcal{C}_0$

Proof

Recall that $U_\alpha : \mathcal{C}_0 \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}_0$

1. By part(a) the previous result we have $\alpha U_\alpha - U_\alpha G = I$ so $U_\alpha(\alpha I - G) = I$. By part (b) we have $\alpha U_\alpha - GU_\alpha = I$ so $(\alpha I - G)U_\alpha = I$.
2. This follows from (a).

So, from the generator G we can determine the potential operators $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$, which in turn determine the transition operators $\mathbf{P} = \{P_t : t \in (0, \infty)\}$. In continuous time, transition operators $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ can be obtained from the single, infinitesimal operator G in a way that is reminiscent of the fact that in discrete time, the transition operators $\mathbf{P} = \{P^n : n \in \mathbb{N}\}$ can be obtained from the single, one-step operator P .

Examples and Applications

Our first example is essentially deterministic.

Consider the Markov process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on \mathbb{R} satisfying the ordinary differential equation

$$\frac{d}{dt} X_t = g(X_t), \quad t \in [0, \infty) \quad (16.2.52)$$

where $g : \mathbb{R} \rightarrow \mathbb{R}$ is Lipschitz continuous. The infinitesimal operator G is given by $Gf(x) = f'(x)g(x)$ for $x \in \mathbb{R}$ on the domain \mathcal{D} of functions $f : \mathbb{R} \rightarrow \mathbb{R}$ where $f \in \mathcal{C}_0$ and $f' \in \mathcal{C}_0$.

Proof

Recall that the only source of randomness in this process is the initial state X_0 . By the continuity assumptions on g , there exists a unique solution $X_t(x)$ to the differential equation with initial value $X_0 = x$, defined for all $t \in [0, \infty)$. The transition operator P_t for $t \in [0, \infty)$ is defined on \mathcal{B} by $P_t f(x) = f[X_t(x)]$ for $x \in \mathbb{R}$. By the ordinary chain rule, if f is differentiable,

$$\frac{P_t f(x) - f(x)}{t} = \frac{f[X_t(x)] - f(x)}{t} \rightarrow f'(x)g(x) \text{ as } t \downarrow 0 \quad (16.2.53)$$

Our next example considers the Poisson process as a Markov process. Compare this with the [binomial process](#) above.

Let $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ denote the Poisson process on \mathbb{N} with rate $\beta \in (0, \infty)$. Define the Markov process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ by $X_t = X_0 + N_t$ where X_0 takes values in \mathbb{N} and is independent of \mathbf{N} .

1. For $t \in [0, \infty)$, show that the probability transition matrix P_t of \mathbf{X} is given by

$$P_t(x, y) = e^{-\beta t} \frac{(\beta t)^{y-x}}{(y-x)!}, \quad x, y \in \mathbb{N}, y \geq x \quad (16.2.54)$$

2. For $\alpha \in [0, \infty)$, show that the potential matrix U_α of \mathbf{X} is given by

$$U_\alpha(x, y) = \frac{1}{\alpha + \beta} \left(\frac{\beta}{\alpha + \beta} \right)^{y-x}, \quad x, y \in \mathbb{N}, y \geq x \quad (16.2.55)$$

3. For $\alpha > 0$ and $x \in \mathbb{N}$, identify the probability distribution defined by $\alpha U_\alpha(x, \cdot)$.

4. Show that the infinitesimal matrix G of \mathbf{X} is given by $G(x, x) = -\beta$, $G(x, x+1) = \beta$ for $x \in \mathbb{N}$.

Solutions

1. Note that for $t \in [0, \infty)$ and $x \in \mathbb{N}$, $P_t(x, \cdot)$ is the (discrete) PDF of $x + N_t$ since N_t has the Poisson distribution with parameter βt .

2. Let $\alpha \in [0, \infty)$ and let $x, y \in \mathbb{N}$ with $x \leq y$. Then

$$\begin{aligned} U_\alpha(x, y) &= \int_0^\infty e^{-\alpha t} P_t(x, y) dt = \int_0^\infty e^{-\alpha t} e^{-\beta t} \frac{(\beta t)^{y-x}}{(y-x)!} dt \\ &= \frac{\beta^{y-x}}{(y-x)!} \int_0^\infty e^{-(\alpha+\beta)t} t^{y-x} dt \end{aligned}$$

The change of variables $s = (\alpha + \beta)t$ gives

$$U_\alpha(x, y) = \frac{\beta^{y-x}}{(y-x)! (\alpha + \beta)^{y-x+1}} \int_0^\infty e^{-s} s^{y-x} ds \quad (16.2.56)$$

But the last integral is $\Gamma(y-x+1) = (y-x)!$. Simplifying gives the result.

3. For $\alpha > 0$,

$$\alpha U_\alpha(x, y) = \frac{\alpha}{\alpha + \beta} \left(\frac{\beta}{\alpha + \beta} \right)^{y-x}, \quad x, y \in \mathbb{N}, y \geq x \quad (16.2.57)$$

As a function of y for fixed x , this is the PDF of $x + Y_\alpha$ where Y_α has the geometric distribution with parameter $\frac{\alpha}{\alpha + \beta}$.

4. Note that for $x, y \in \mathbb{N}$, $G(x, y) = \frac{d}{dt} P_t(x, y) \big|_{t=0}$. By simple calculus, this is $-\beta$ if $y = x$, β if $y = x + 1$, and 0 otherwise.

This page titled [16.2: Potentials and Generators for General Markov Processes](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.3: Introduction to Discrete-Time Chains

In this and the next several sections, we consider a Markov process with the discrete time space \mathbb{N} and with a discrete (countable) state space. Recall that a Markov process with a discrete state space is called a *Markov chain*, so we are studying *discrete-time Markov chains*.

Review

We will review the basic definitions and concepts in the general introduction. With both time and space discrete, many of these definitions and concepts simplify considerably. As usual, our starting point is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so Ω is the sample space, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on (Ω, \mathcal{F}) . Let $\mathbf{X} = (X_0, X_1, X_2, \dots)$ be a stochastic process defined on the probability space, with time space \mathbb{N} and with countable state space S . In the context of the general introduction, S is given the power set $\mathcal{P}(S)$ as the σ -algebra, so all subsets of S are measurable, as are all functions from S into another measurable space. Counting measure $\#$ is the natural measure on S , so integrals over S are simply sums. The same comments apply to the time space \mathbb{N} : all subsets of \mathbb{N} are measurable and counting measure $\#$ is the natural measure on \mathbb{N} .

The vector space \mathcal{B} consisting of bounded functions $f : S \rightarrow \mathbb{R}$ will play an important role. The norm that we use is the *supremum norm* defined by

$$\|f\| = \sup\{|f(x)| : x \in S\}, \quad f \in \mathcal{B} \quad (16.3.1)$$

For $n \in \mathbb{N}$, let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$, the σ -algebra generated by the process up to time n . Thus $\mathcal{F} = \{\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots\}$ is the natural filtration associated with \mathbf{X} . We also let $\mathcal{G}_n = \sigma\{X_n, X_{n+1}, \dots\}$, the σ -algebra generated by the process from time n on. So if $n \in \mathbb{N}$ represents the present time, then \mathcal{F}_n contains the events in the past and \mathcal{G}_n the events in the future.

Definitions

We start with the basic definition of the Markov property: the past and future are conditionally independent, given the present.

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a Markov chain if $\mathbb{P}(A \cap B \mid X_n) = \mathbb{P}(A \mid X_n)\mathbb{P}(B \mid X_n)$ for every $n \in \mathbb{N}$, $A \in \mathcal{F}_n$ and $B \in \mathcal{G}_n$.

There are a number of equivalent formulations of the Markov property for a discrete-time Markov chain. We give a few of these.

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a Markov chain if either of the following equivalent conditions is satisfied:

1. $\mathbb{P}(X_{n+1} = x \mid \mathcal{F}_n) = \mathbb{P}(X_{n+1} = x \mid X_n)$ for every $n \in \mathbb{N}$ and $x \in S$.
2. $\mathbb{E}[f(X_{n+1}) \mid \mathcal{F}_n] = \mathbb{E}[f(X_{n+1}) \mid X_n]$ for every $n \in \mathbb{N}$ and $f \in \mathcal{B}$.

Part (a) states that for $n \in \mathbb{N}$, the conditional probability density function of X_{n+1} given \mathcal{F}_n is the same as the conditional probability density function of X_{n+1} given X_n . Part (b) also states, in terms of expected value, that the conditional distribution of X_{n+1} given \mathcal{F}_n is the same as the conditional distribution of X_{n+1} given X_n . Both parts are the Markov property looking just one time step in the future. But with discrete time, this is equivalent to the Markov property at general future times.

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a Markov chain if either of the following equivalent conditions is satisfied:

1. $\mathbb{P}(X_{n+k} = x \mid \mathcal{F}_n) = \mathbb{P}(X_{n+k} = x \mid X_n)$ for every $n, k \in \mathbb{N}$ and $x \in S$.
2. $\mathbb{E}[f(X_{n+k}) \mid \mathcal{F}_n] = \mathbb{E}[f(X_{n+k}) \mid X_n]$ for every $n, k \in \mathbb{N}$ and $f \in \mathcal{B}$.

Part (a) states that for $n, k \in \mathbb{N}$, the conditional probability density function of X_{n+k} given \mathcal{F}_n is the same as the conditional probability density function of X_{n+k} given X_n . Part (b) also states, in terms of expected value, that the conditional distribution of X_{n+k} given \mathcal{F}_n is the same as the conditional distribution of X_{n+k} given X_n . In discrete time and space, the Markov property can also be stated without explicit reference to σ -algebras. If you are not familiar with measure theory, you can take this as the starting definition.

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a Markov chain if for every $n \in \mathbb{N}$ and every sequence of states $(x_0, x_1, \dots, x_{n-1}, x, y)$

$$\mathbb{P}(X_{n+1} = y \mid X_0 = x_0, X_1 = x_1, \dots, X_{n-1} = x_{n-1}, X_n = x) = \mathbb{P}(X_{n+1} = y \mid X_n = x) \quad (16.3.2)$$

The theory of discrete-time Markov chains is simplified considerably if we add an additional assumption.

A Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is *time homogeneous* if

$$\mathbb{P}(X_{n+k} = y \mid X_k = x) = \mathbb{P}(X_n = y \mid X_0 = x) \quad (16.3.3)$$

for every $k, n \in \mathbb{N}$ and every $x, y \in S$.

That is, the conditional distribution of X_{n+k} given $X_k = x$ depends only on n . So if \mathbf{X} is homogeneous (we usually don't bother with the *time* adjective), then the chain $\{X_{k+n} : n \in \mathbb{N}\}$ given $X_k = x$ is equivalent (in distribution) to the chain $\{X_n : n \in \mathbb{N}\}$ given $X_0 = x$. For this reason, the initial distribution is often unspecified in the study of Markov chains—if the chain is in state $x \in S$ at a particular time $k \in \mathbb{N}$, then it doesn't really matter how the chain got to state x ; the process essentially “starts over”, independently of the past. The term *stationary* is sometimes used instead of homogeneous.

From now on, we will usually assume that our Markov chains are homogeneous. This is not as big of a loss of generality as you might think. A non-homogeneous Markov chain can be turned into a homogeneous Markov process by enlarging the state space, as shown in the introduction to general Markov processes, but at the cost of creating an uncountable state space. For a homogeneous Markov chain, if $k, n \in \mathbb{N}$, $x \in S$, and $f \in \mathcal{B}$, then

$$\mathbb{E}[f(X_{k+n}) \mid X_k = x] = \mathbb{E}[f(X_n) \mid X_0 = x] \quad (16.3.4)$$

Stopping Times and the Strong Markov Property

Consider again a stochastic process $\mathbf{X} = (X_0, X_1, X_2, \dots)$ with countable state space S , and with the natural filtration $\mathfrak{F} = (\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots)$ as given above. Recall that a random variable τ taking values in $\mathbb{N} \cup \{\infty\}$ is a *stopping time* or a *Markov time* for \mathbf{X} if $\{\tau = n\} \in \mathcal{F}_n$ for each $n \in \mathbb{N}$. Intuitively, we can tell whether or not $\tau = n$ by observing the chain up to time n . In a sense, a stopping time is a random time that does not require that we see into the future. The following result gives the quintessential examples of stopping times.

Suppose again $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a discrete-time Markov chain with state space S as defined above. For $A \subseteq S$, the following random times are stopping times:

1. $\rho_A = \inf\{n \in \mathbb{N} : X_n \in A\}$, the *entrance time* to A .
2. $\tau_A = \inf\{n \in \mathbb{N}_+ : X_n \in A\}$, the *hitting time* to A .

Proof

For $n \in \mathbb{N}$

1. $\{\rho_A = n\} = \{X_0 \notin A, X_1 \notin A, \dots, X_{n-1} \notin A, X_n \in A\} \in \mathcal{F}_n$
2. $\{\tau_A = n\} = \{X_1 \notin A, X_2 \notin A, \dots, X_{n-1} \notin A, X_n \in A\} \in \mathcal{F}_n$

An example of a random time that is generally *not* a stopping time is the *last* time that the process is in A :

$$\zeta_A = \max\{n \in \mathbb{N}_+ : X_n \in A\} \quad (16.3.5)$$

We cannot tell if $\zeta_A = n$ without looking into the future: $\{\zeta_A = n\} = \{X_n \in A, X_{n+1} \notin A, X_{n+2} \notin A, \dots\}$ for $n \in \mathbb{N}$.

If τ is a stopping time for \mathbf{X} , the σ -algebra associated with τ is

$$\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau = n\} \in \mathcal{F}_n \text{ for all } n \in \mathbb{N}\} \quad (16.3.6)$$

Intuitively, \mathcal{F}_τ contains the events that can be described by the process up to the random time τ , in the same way that \mathcal{F}_n contains the events that can be described by the process up to the deterministic time $n \in \mathbb{N}$. For more information see the section on filtrations and stopping times.

The *strong Markov property* states that the future is independent of the past, given the present, when the present time is a stopping time. For a discrete-time Markov chain, the ordinary Markov property implies the strong Markov property.

If $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a discrete-time Markov chain then \mathbf{X} has the strong Markov property. That is, if τ is a finite stopping time for \mathbf{X} then

1. $\mathbb{P}(X_{\tau+k} = x \mid \mathcal{F}_\tau) = \mathbb{P}(X_{\tau+k} = x \mid X_\tau)$ for every $k \in \mathbb{N}$ and $x \in S$.
2. $\mathbb{E}[f(X_{\tau+k}) \mid \mathcal{F}_\tau] = \mathbb{E}[f(X_{\tau+k}) \mid X_\tau]$ for every $k \in \mathbb{N}$ and $f \in \mathcal{B}$.

Part (a) states that the conditional probability density function of $X_{\tau+k}$ given \mathcal{F}_τ is the same as the conditional probability density function of $X_{\tau+k}$ given just X_τ . Part (b) also states, in terms of expected value, that the conditional distribution of $X_{\tau+k}$ given \mathcal{F}_τ is the same as the conditional distribution of $X_{\tau+k}$ given just X_τ . Assuming homogeneity as usual, the Markov chain $\{X_{\tau+n} : n \in \mathbb{N}\}$ given $X_\tau = x$ is equivalent in distribution to the chain $\{X_n : n \in \mathbb{N}\}$ given $X_0 = x$.

Transition Matrices

Suppose again that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a homogeneous, discrete-time Markov chain with state space S . With a discrete state space, the transition kernels studied in the general introduction become transition *matrices*, with rows and columns indexed by S (and so perhaps of infinite size). The kernel operations become familiar matrix operations. The results in this section are special cases of the general results, but we sometimes give independent proofs for completeness, and because the proofs are simpler. You may want to review the section on kernels in the chapter on expected value.

For $n \in \mathbb{N}$ let

$$P_n(x, y) = \mathbb{P}(X_n = y \mid X_0 = x), \quad (x, y) \in S \times S \quad (16.3.7)$$

The matrix P_n is the n -step *transition probability matrix* for \mathbf{X} .

Thus, $y \mapsto P_n(x, y)$ is the probability density function of X_n given $X_0 = x$. In particular, P_n is a *probability matrix* (or *stochastic matrix*) since $P_n(x, y) \geq 0$ for $(x, y) \in S^2$ and $\sum_{y \in S} P_n(x, y) = 1$ for $x \in S$. As with any nonnegative matrix on S , P_n defines a kernel on S for $n \in \mathbb{N}$:

$$P_n(x, A) = \sum_{y \in A} P_n(x, y) = \mathbb{P}(X_n \in A \mid X_0 = x), \quad x \in S, A \subseteq S \quad (16.3.8)$$

So $A \mapsto P_n(x, A)$ is the probability distribution of X_n given $X_0 = x$. The next result is the *Chapman-Kolmogorov equation*, named for Sydney Chapman and Andrei Kolmogorov. It gives the basic relationship between the transition matrices.

If $m, n \in \mathbb{N}$ then $P_m P_n = P_{m+n}$

Proof

This follows from the Markov and time-homogeneous properties and a basic conditioning argument. If $x, z \in S$ then

$$P_{m+n}(x, z) = \mathbb{P}(X_{m+n} = z \mid X_0 = x) = \sum_{y \in S} \mathbb{P}(X_{m+n} = z \mid X_0 = x, X_m = y) \mathbb{P}(X_m = y \mid X_0 = x) \quad (16.3.9)$$

But by the Markov property and time-homogeneous properties

$$\mathbb{P}(X_{m+n} = z \mid X_0 = x, X_m = y) = \mathbb{P}(X_n = z \mid X_0 = y) = P_n(y, z) \quad (16.3.10)$$

Of course also $\mathbb{P}(X_m = y \mid X_0 = x) = P_m(x, y)$. Hence we have

$$P_{m+n}(x, z) = \sum_{y \in S} P_m(x, y) P_n(y, z) \quad (16.3.11)$$

The right side, by definition, is $P_m P_n(x, z)$.

It follows immediately that the transition matrices are just the matrix powers of the one-step transition matrix. That is, letting $P = P_1$ we have $P_n = P^n$ for all $n \in \mathbb{N}$. Note that $P^0 = I$, the identity matrix on S given by $I(x, y) = 1$ if $x = y$ and 0 otherwise. The right operator corresponding to P^n yields an expected value.

Suppose that $n \in \mathbb{N}$ and that $f : S \rightarrow \mathbb{R}$. Then, assuming that the expected value exists,

$$P^n f(x) = \sum_{y \in S} P^n(x, y) f(y) = \mathbb{E}[f(X_n) \mid X_0 = x], \quad x \in S \quad (16.3.12)$$

Proof

This follows easily from the definitions:

$$P^n f(x) = \sum_{y \in S} P^n(x, y) f(y) = \sum_{y \in S} \mathbb{P}(X_n = y \mid X_0 = x) f(y) = \mathbb{E}[f(X_n) \mid X_0 = x], \quad x \in S \quad (16.3.13)$$

The existence of the expected value is only an issue if S is infinite. In particular, the result holds if f is nonnegative or if $f \in \mathcal{B}$ (which in turn would always be the case if S is finite). In fact, P^n is a linear *contraction operator* on the space \mathcal{B} for $n \in \mathbb{N}$. That is, if $f \in \mathcal{B}$ then $P^n f \in \mathcal{B}$ and $\|P^n f\| \leq \|f\|$. The left operator corresponding to P^n is defined similarly. For $f : S \rightarrow \mathbb{R}$

$$f P^n(y) = \sum_{x \in S} f(x) P^n(x, y), \quad y \in S \quad (16.3.14)$$

assuming again that the sum makes sense (as before, only an issue when S is infinite). The left operator is often restricted to nonnegative functions, and we often think of such a function as the density function (with respect to $\#$) of a positive measure on S . In this sense, the left operator maps a density function to another density function.

A function $f : S \rightarrow \mathbb{R}$ is *invariant* for P (or for the chain \mathbf{X}) if $fP = f$.

Clearly if f is invariant, so that $fP = f$ then $fP^n = f$ for all $n \in \mathbb{N}$. If f is a probability density function, then so is fP .

If X_0 has probability density function f , then X_n has probability density function fP^n for $n \in \mathbb{N}$.

Proof

Again, this follows easily from the definitions and a conditioning argument.

$$\mathbb{P}(X_n = y) = \sum_{x \in S} \mathbb{P}(X_0 = x) \mathbb{P}(X_n = y \mid X_0 = x) = \sum_{x \in S} f(x) P^n(x, y) = f P^n(y), \quad y \in S \quad (16.3.15)$$

In particular, if X_0 has probability density function f , and f is invariant for \mathbf{X} , then X_n has probability density function f for all $n \in \mathbb{N}$, so the sequence of variables $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is identically distributed. Combining two results above, suppose that X_0 has probability density function f and that $g : S \rightarrow \mathbb{R}$. Assuming the expected value exists, $\mathbb{E}[g(X_n)] = f P^n g$. Explicitly,

$$\mathbb{E}[g(X_n)] = \sum_{x \in S} \sum_{y \in S} f(x) P^n(x, y) g(y) \quad (16.3.16)$$

It also follows from the last theorem that the distribution of X_0 (the *initial distribution*) and the one-step transition matrix determine the distribution of X_n for each $n \in \mathbb{N}$. Actually, these basic quantities determine the *finite dimensional distributions* of the process, a stronger result.

Suppose that X_0 has probability density function f_0 . For any sequence of states $(x_0, x_1, \dots, x_n) \in S^n$,

$$\mathbb{P}(X_0 = x_0, X_1 = x_1, \dots, X_n = x_n) = f_0(x_0) P(x_0, x_1) P(x_1, x_2) \cdots P(x_{n-1}, x_n) \quad (16.3.17)$$

Proof

This follows directly from the Markov property and the multiplication rule of conditional probability:

$$\begin{aligned} \mathbb{P}(X_0 = x_0, X_1 = x_1, \dots, X_n = x_n) &= \mathbb{P}(X_0 = x_0) \mathbb{P}(X_1 = x_1 \mid X_0 = x_0) \mathbb{P}(X_2 = x_2 \mid X_0 = x_0, X_1 = x_1) \\ &\quad \cdots \mathbb{P}(X_n = x_n \mid X_0 = x_0, \dots, X_{n-1} = x_{n-1}) \end{aligned} \quad (16.3.18)$$

But by the Markov property, this reduces to

$$\begin{aligned}\mathbb{P}(X_0 = x_0, X_1 = x_1, \dots, X_n = x_n) &= \mathbb{P}(X_0 = x_0) \mathbb{P}(X_1 = x_1 | X_0 = x_0) \mathbb{P}(X_2 = x_2 | X_1 = x_1) \cdots \mathbb{P}(X_n = x_n | X_{n-1} = x_{n-1}) \\ &= f_0(x_0) P(x_0, x_1) P(x_1, x_2) \cdots P(x_{n-1}, x_n)\end{aligned}$$

Computations of this sort are the reason for the term *chain* in the name *Markov chain*. From this result, it follows that given a probability matrix P on S and a probability density function f on S , we can construct a Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ such that X_0 has probability density function f and the chain has one-step transition matrix P . In applied problems, we often know the one-step transition matrix P from modeling considerations, and again, the initial distribution is often unspecified.

There is a natural graph (in the combinatorial sense) associated with a homogeneous, discrete-time Markov chain.

Suppose again that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a Markov chain with state space S and transition probability matrix P . The *state graph* of \mathbf{X} is the directed graph with vertex set S and edge set $E = \{(x, y) \in S^2 : P(x, y) > 0\}$.

That is, there is a directed edge from x to y if and only if state x leads to state y in one step. Note that the graph may well have loops, since a state can certainly lead back to itself in one step. More generally, we have the following result:

Suppose again that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a Markov chain with state space S and transition probability matrix P . For $x, y \in S$ and $n \in \mathbb{N}_+$, there is a directed path of length n in the state graph from x to y if and only if $P^n(x, y) > 0$.

Proof

This follows since $P^n(x, y) > 0$ if and only if there exists a sequence of states $(x_1, x_2, \dots, x_{n-1})$ with $P(x, x_1) > 0, P(x_1, x_2) > 0, \dots, P(x_{n-1}, y) > 0$. This is also precisely the condition for the existence of a directed path $(x, x_1, \dots, x_{n-1}, y)$ of length n from x to y in the state graph.

Potential Matrices

For $\alpha \in (0, 1]$, the α -potential matrix R_α of \mathbf{X} is

$$R_\alpha = \sum_{n=0}^{\infty} \alpha^n P^n, \quad (x, y) \in S^2 \quad (16.3.19)$$

1. $R = R_1$ is simply the *potential matrix* of \mathbf{X} .
2. $R(x, y)$ is the expected number of visits by \mathbf{X} to $y \in S$, starting at $x \in S$.

Proof

First the definition of R_α as an infinite series of matrices makes sense since P^n is a nonnegative matrix for each n . The interpretation of $R(x, y)$ for $(x, y) \in S^2$ comes from interchanging sum and expected value, again justified since the terms are nonnegative.

$$R(x, y) = \sum_{n=0}^{\infty} P^n(x, y) = \sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}(X_n = y) | X_0 = x] = \mathbb{E}\left(\sum_{n=0}^{\infty} \mathbf{1}(X_n = y) \mid X_0 = x\right) = \mathbb{E}[\#\{n \in \mathbb{N} : X_n = y\} | X_0 = x] \quad (16.3.20)$$

Note that it's quite possible that $R(x, y) = \infty$ for some $(x, y) \in S^2$. In fact, knowing when this is the case is of considerable importance in recurrence and transience, which we study in the next section. As with any nonnegative matrix, the α -potential matrix defines a kernel and defines left and right operators. For the kernel,

$$R_\alpha(x, A) = \sum_{y \in A} R_\alpha(x, y) = \sum_{n=0}^{\infty} \alpha^n P^n(x, A), \quad x \in S, A \subseteq S \quad (16.3.21)$$

In particular, $R(x, A)$ is the expected number of visits by the chain to A starting in x :

$$R(x, A) = \sum_{y \in A} R(x, y) = \sum_{n=0}^{\infty} P^n(x, A) = \mathbb{E}\left[\sum_{n=0}^{\infty} \mathbf{1}(X_n \in A)\right], \quad x \in S, A \subseteq S \quad (16.3.22)$$

If $\alpha \in (0, 1)$, then $R_\alpha(x, S) = \frac{1}{1-\alpha}$ for all $x \in S$.

Proof

Using geometric series,

$$R_\alpha(x, S) = \sum_{n=0}^{\infty} \alpha^n P^n(x, S) = \sum_{n=0}^{\infty} \alpha^n = \frac{1}{1-\alpha} \quad (16.3.23)$$

Hence R_α is a bounded matrix for $\alpha \in (0, 1)$ and $(1-\alpha)R_\alpha$ is a probability matrix. There is a simple interpretation of this matrix.

If $\alpha \in (0, 1)$ then $(1-\alpha)R_\alpha(x, y) = \mathbb{P}(X_N = y | X_0 = x)$ for $(x, y) \in S^2$, where N is independent of \mathbf{X} and has the geometric distribution on \mathbb{N} with parameter $1-\alpha$.

Proof

Let $(x, y) \in S^2$. Conditioning on N gives

$$\mathbb{P}(X_N = y \mid X_0 = x) = \sum_{n=0}^{\infty} \mathbb{P}(N = n) \mathbb{P}(X_N = y \mid X_0 = x, N = n) \quad (16.3.24)$$

But by the substitution rule and the assumption of independence,

$$\mathbb{P}(X_N = y \mid N = n, X_0 = x) = \mathbb{E}(X_n = y \mid N = n, X_0 = x) = \mathbb{P}(X_n = y \mid X_0 = x) = P^n(x, y) \quad (16.3.25)$$

Since N has the geometric distribution on N with parameter $1 - \alpha$ we have $\mathbb{P}(N = n) = (1 - \alpha)\alpha^n$. Hence

$$\mathbb{P}(X_N = y \mid X_0 = x) = \sum_{n=0}^{\infty} (1 - \alpha)\alpha^n P^n(x, y) = (1 - \alpha)R_\alpha(x, y) \quad (16.3.26)$$

So $(1 - \alpha)R_\alpha$ can be thought of as a *transition matrix* just as P^n is a transition matrix, but corresponding to the *random time* N (with α as a parameter) rather than the *deterministic time* n . An interpretation of the potential matrix R_α for $\alpha \in (0, 1)$ can also be given in economic terms. Suppose that we receive one monetary unit each time the chain visits a fixed state $y \in S$. Then $R(x, y)$ is the expected total reward, starting in state $x \in S$. However, typically money that we will receive at times distant in the future have less value to us *now* than money that we will receive soon. Specifically suppose that a monetary unit at time $n \in \mathbb{N}$ has a *present value* of α^n , so that α is an *inflation factor* (sometimes also called a *discount factor*). Then $R_\alpha(x, y)$ gives the expected total discounted reward, starting at $x \in S$.

The potential kernels $\mathbf{R} = \{R_\alpha : \alpha \in (0, 1)\}$ completely determine the transition kernels $\mathbf{P} = \{P_n : n \in \mathbb{N}\}$.

Proof

Note that for $(x, y) \in S^2$, the function $\alpha \mapsto R_\alpha(x, y)$ is a power series in α with coefficients $n \mapsto P^n(x, y)$. In the language of combinatorics, $\alpha \mapsto R_\alpha(x, y)$ is the *ordinary generating function* of the sequence $n \mapsto P^n(x, y)$. As noted above, this power series has radius of convergence at least 1, so we can extend the domain to $\alpha \in (-1, 1)$. Thus, given the potential matrices, we can recover the transition matrices by taking derivatives and evaluating at 0:

$$P^n(x, y) = \frac{1}{n!} \left[\frac{d^n}{d\alpha^n} R_\alpha(x, y) \right]_{\alpha=0} \quad (16.3.27)$$

Of course, it's really only necessary to determine P , the one step transition kernel, since the other transition kernels are powers of P . In any event, it follows that the matrices $\mathbf{R} = \{R_\alpha : \alpha \in (0, 1)\}$, along with the initial distribution, completely determine the finite dimensional distributions of the Markov chain \mathbf{X} . The potential matrices commute with each other and with the transition matrices.

If $\alpha, \beta \in (0, 1]$ and $k \in \mathbb{N}$, then

1. $P^k R_\alpha = R_\alpha P^k = \sum_{n=0}^{\infty} \alpha^n P^{n+k}$
2. $R_\alpha R_\beta = R_\beta R_\alpha = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \alpha^m \beta^n P^{m+n}$

Proof

Distributing matrix products through matrix sums is allowed since the matrices are nonnegative.

1. Directly

$$R_\alpha P^k = \sum_{n=0}^{\infty} \alpha^n P^n P^k = \sum_{n=0}^{\infty} \alpha^n P^{n+k} \quad (16.3.28)$$

The other direction requires an interchange.

$$P^k R_\alpha = P^k \sum_{n=0}^{\infty} \alpha^n P^n = \sum_{n=0}^{\infty} \alpha^n P^k P^n = \sum_{n=0}^{\infty} \alpha^n P^{n+k} \quad (16.3.29)$$

2. First,

$$R_\alpha R_\beta = \sum_{m=0}^{\infty} \alpha^m P^m R_\beta = \sum_{m=0}^{\infty} \alpha^m P^m \left(\sum_{n=0}^{\infty} \beta^n P^n \right) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \alpha^m \beta^n P^m P^n = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} \alpha^m \beta^n P^{m+n} \quad (16.3.30)$$

The other direction is similar.

The fundamental equation that relates the potential matrices is given next.

If $\alpha, \beta \in (0, 1]$ with $\alpha \geq \beta$ then

$$\alpha R_\alpha = \beta R_\beta + (\alpha - \beta) R_\alpha R_\beta \quad (16.3.31)$$

Proof

If $\alpha = \beta$ the equation is trivial, so assume $\alpha > \beta$. From the previous result,

$$R_\alpha R_\beta = \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} \alpha^j \beta^k P^{j+k} \quad (16.3.32)$$

Changing variables to sum over $n = j + k$ and k gives

$$R_\alpha R_\beta = \sum_{n=0}^{\infty} \sum_{k=0}^n \alpha^{n-k} \beta^k P^n = \sum_{n=0}^{\infty} \sum_{k=0}^n \left(\frac{\beta}{\alpha}\right)^k \alpha^n P^n = \sum_{n=0}^{\infty} \frac{1 - \left(\frac{\beta}{\alpha}\right)^{n+1}}{1 - \frac{\beta}{\alpha}} \alpha^n P^n \quad (16.3.33)$$

Simplifying gives

$$R_\alpha R_\beta = \frac{1}{\alpha - \beta} [\alpha R_\alpha - \beta R_\beta] \quad (16.3.34)$$

Note that since $\beta < 1$, the matrix R_β has finite values, so we don't have to worry about the dreaded indeterminate form $\infty - \infty$.

If $\alpha \in (0, 1]$ then $I + \alpha R_\alpha P = I + \alpha P R_\alpha = R_\alpha$.

Proof

From the result above,

$$I + \alpha R_\alpha P = I + \alpha P R_\alpha = I + \sum_{n=0}^{\infty} \alpha^{n+1} P^{n+1} = \sum_{n=0}^{\infty} \alpha^n P^n = R_\alpha \quad (16.3.35)$$

This leads to an important result: when $\alpha \in (0, 1)$, there is an inverse relationship between P and R_α .

If $\alpha \in (0, 1)$, then

1. $R_\alpha = (I - \alpha P)^{-1}$
2. $P = \frac{1}{\alpha} (I - R_\alpha^{-1})$

Proof

The matrices have finite values, so we can subtract. The identity $I + \alpha R_\alpha P = R_\alpha$ leads to $R_\alpha (I - \alpha P) = I$ and the identity $I + \alpha P R_\alpha = R_\alpha$ leads to $(I - \alpha P) R_\alpha = I$. Hence (a) holds. Part (b) follows from (a).

This result shows again that the potential matrix R_α determines the transition operator P .

Sampling in Time

If we sample a Markov chain at multiples of a fixed time k , we get another (homogeneous) chain.

Suppose that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is an Markov chain with state space S and transition probability matrix P . For fixed $k \in \mathbb{N}_+$, the sequence $\mathbf{X}_k = (X_0, X_k, X_{2k}, \dots)$ is a Markov chain on S with transition probability matrix P^k .

If we sample a Markov chain at a general increasing sequence of time points $0 < n_1 < n_2 < \dots$ in \mathbb{N} , then the resulting stochastic process $\mathbf{Y} = (Y_0, Y_1, Y_2, \dots)$, where $Y_k = X_{n_k}$ for $k \in \mathbb{N}$, is still a Markov chain, but is not time homogeneous in general.

Recall that if A is a nonempty subset of S , then P_A is the matrix P restricted to $A \times A$. So P_A is a sub-stochastic matrix, since the row sums may be less than 1. Recall also that P_A^n means $(P_A)^n$, not $(P^n)_A$; in general these matrices are different.

If A is a nonempty subset of S then for $n \in \mathbb{N}$,

$$P_A^n(x, y) = \mathbb{P}(X_1 \in A, X_2 \in A, \dots, X_{n-1} \in A, X_n = y \mid X_0 = x), \quad (x, y) \in A \times A \quad (16.3.36)$$

That is, $P_A^n(x, y)$ is the probability of going from state x to y in n steps, remaining in A all the while. In terms of the state graph of \mathbf{X} , it is the sum of products of probabilities along paths of length n from x to y that stay inside A .

Examples and Applications

Computational Exercises

Let $\mathbf{X} = (X_0, X_1, \dots)$ be the Markov chain on $S = \{a, b, c\}$ with transition matrix

$$P = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{4} & 0 & \frac{3}{4} \\ 1 & 0 & 0 \end{bmatrix} \quad (16.3.37)$$

For the Markov chain \mathbf{X} ,

1. Draw the state graph.
2. Find $\mathbb{P}(X_1 = a, X_2 = b, X_3 = c \mid X_0 = a)$

- Find P^2
- Suppose that $g: S \rightarrow \mathbb{R}$ is given by $g(a) = 1$, $g(b) = 2$, $g(c) = 3$. Find $\mathbb{E}[g(X_2) | X_0 = x]$ for $x \in S$.
- Suppose that X_0 has the uniform distribution on S . Find the probability density function of X_2 .

Answer

- The edge set is $E = \{(a, a), (a, b), (b, a), (b, c), (c, a)\}$
- $P(a, a)P(a, b)P(b, c) = \frac{3}{16}$
- By standard matrix multiplication,

$$P^2 = \begin{bmatrix} \frac{3}{8} & \frac{1}{4} & \frac{3}{8} \\ \frac{7}{8} & \frac{1}{8} & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 \end{bmatrix} \quad (16.3.38)$$

- In matrix form,

$$g = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \quad P^2 g = \begin{bmatrix} 2 \\ \frac{9}{8} \\ \frac{3}{2} \end{bmatrix} \quad (16.3.39)$$

- In matrix form, X_0 has PDF $f = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$, and X_2 has PDF $fP^2 = \begin{bmatrix} \frac{7}{12} & \frac{7}{24} & \frac{1}{8} \end{bmatrix}$.

Let $A = \{a, b\}$. Find each of the following:

- P_A
- P_A^2
- $(P^2)_A$

Proof

- $P_A = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{4} & 0 \\ \frac{3}{8} & \frac{1}{4} \end{bmatrix}$
- $P_A^2 = \begin{bmatrix} \frac{3}{8} & \frac{1}{4} \\ \frac{1}{8} & \frac{1}{8} \end{bmatrix}$
- $(P^2)_A = \begin{bmatrix} \frac{3}{8} & \frac{1}{4} \\ \frac{7}{8} & \frac{1}{8} \end{bmatrix}$

Find the invariant probability density function of \mathbf{X}

Answer

Solving $fP = f$ subject to the condition that f is a PDF gives $f = \begin{bmatrix} \frac{8}{15} & \frac{4}{15} & \frac{3}{15} \end{bmatrix}$

Compute the α -potential matrix R_α for $\alpha \in (0, 1)$.

Answer

Computing $R_\alpha = (I - \alpha P)^{-1}$ gives

$$R_\alpha = \frac{1}{(1 - \alpha)(8 + 4\alpha + 3\alpha^2)} \begin{bmatrix} 8 & 4\alpha & 3\alpha^2 \\ 2\alpha + 6\alpha^2 & 8 - 4\alpha & 6\alpha - 3\alpha^2 \\ 8\alpha & 4\alpha^2 & 8 - 4\alpha - \alpha^2 \end{bmatrix} \quad (16.3.40)$$

As a check on our work, note that the row sums are $\frac{1}{1 - \alpha}$.

The Two-State Chain

Perhaps the simplest, non-trivial Markov chain has two states, say $S = \{0, 1\}$ and the transition probability matrix given below, where $p \in (0, 1)$ and $q \in (0, 1)$ are parameters.

$$P = \begin{bmatrix} 1 - p & p \\ q & 1 - q \end{bmatrix} \quad (16.3.41)$$

For $n \in \mathbb{N}$,

$$P^n = \frac{1}{p + q} \begin{bmatrix} q + p(1 - p - q)^n & p - p(1 - p - q)^n \\ q - q(1 - p - q)^n & p + q(1 - p - q)^n \end{bmatrix} \quad (16.3.42)$$

Proof

The eigenvalues of P are 1 and $1 - p - q$. Next, $B^{-1}PB = D$ where

$$B = \begin{bmatrix} 1 & -p \\ 1 & q \end{bmatrix}, \quad D = \begin{bmatrix} 1 & 0 \\ 0 & 1 - p - q \end{bmatrix} \quad (16.3.43)$$

Hence $P^n = BD^nB^{-1}$, which gives the expression above.

As $n \rightarrow \infty$,

$$P^n \rightarrow \frac{1}{p+q} \begin{bmatrix} q & p \\ q & p \end{bmatrix} \quad (16.3.44)$$

Proof

Note that $0 < p+q < 2$ and so $-1 < 1 - (p+q) < 1$. Hence $(1 - p - q)^n \rightarrow 0$ as $n \rightarrow \infty$.

Open the simulation of the two-state, discrete-time Markov chain. For various values of p and q , and different initial states, run the simulation 1000 times. Compare the relative frequency distribution to the limiting distribution, and in particular, note the rate of convergence. Be sure to try the case $p = q = 0.01$

The only invariant probability density function for the chain is

$$f = \left[\frac{q}{p+q} \quad \frac{p}{p+q} \right] \quad (16.3.45)$$

Proof

Let $f = [a \quad b]$. The matrix equation $fP = f$ leads to $-pa + qb = 0$ so $b = a\frac{p}{q}$. The condition $a + b = 1$ for f to be a PDF then gives $a = \frac{q}{p+q}$, $b = \frac{p}{p+q}$

For $\alpha \in (0, 1)$, the α -potential matrix is

$$R_\alpha = \frac{1}{(p+q)(1-\alpha)} \begin{bmatrix} q & p \\ q & p \end{bmatrix} + \frac{1}{(p+q)^2(1-\alpha)} \begin{bmatrix} p & -p \\ -q & q \end{bmatrix} \quad (16.3.46)$$

Proof

In this case, R_α can be computed directly as $\sum_{n=0}^{\infty} \alpha^n P^n$ using geometric series.

In spite of its simplicity, the two state chain illustrates some of the basic limiting behavior and the connection with invariant distributions that we will study in general in a later section.

Independent Variables and Random Walks

Suppose that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a sequence of independent random variables taking values in a countable set S , and that (X_1, X_2, \dots) are identically distributed with (discrete) probability density function f .

\mathbf{X} is a Markov chain on S with transition probability matrix P given by $P(x, y) = f(y)$ for $(x, y) \in S \times S$. Also, f is invariant for P .

Proof

As usual, let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$ for $n \in \mathbb{N}$. Since the sequence \mathbf{X} is independent,

$$\mathbb{P}(X_{n+1} = y \mid \mathcal{F}_n) = \mathbb{P}(X_{n+1} = y) = f(y), \quad y \in S \quad (16.3.47)$$

Also,

$$fP(y) = \sum_{x \in S} f(x)P(x, y) = \sum_{x \in S} f(x)f(y) = f(y), \quad y \in S \quad (16.3.48)$$

As a Markov chain, the process \mathbf{X} is not very interesting, although of course it is very interesting in other ways. Suppose now that $S = \mathbb{Z}$, the set of integers, and consider the *partial sum process* (or *random walk*) \mathbf{Y} associated with \mathbf{X} :

$$Y_n = \sum_{i=0}^n X_i, \quad n \in \mathbb{N} \quad (16.3.49)$$

\mathbf{Y} is a Markov chain on \mathbb{Z} with transition probability matrix Q given by $Q(x, y) = f(y - x)$ for $(x, y) \in \mathbb{Z} \times \mathbb{Z}$.

Proof

Again, let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$ for $n \in \mathbb{N}$. Then also, $\mathcal{F}_n = \sigma\{Y_0, Y_1, \dots, Y_n\}$ for $n \in \mathbb{N}$. Hence

$$\mathbb{P}(Y_{n+1} = y \mid \mathcal{F}_n) = \mathbb{P}(Y_n + X_{n+1} = y \mid \mathcal{F}_n) = \mathbb{P}(Y_n + X_{n+1} = y \mid Y_n), \quad y \in \mathbb{Z} \quad (16.3.50)$$

since the sequence \mathbf{X} is independent. In particular,

$$\mathbb{P}(Y_{n+1} = y \mid Y_n = x) = \mathbb{P}(x + X_{n+1} = y \mid Y_n = x) = \mathbb{P}(X_{n+1} = y - x) = f(y - x), \quad (x, y) \in \mathbb{Z}^2 \quad (16.3.51)$$

Thus the probability density function f governs the distribution of a step size of the random walker on \mathbb{Z} .

Consider the special case of the random walk on \mathbb{Z} with $f(1) = p$ and $f(-1) = 1 - p$, where $p \in (0, 1)$.

1. Give the transition matrix Q explicitly.
2. Give Q^n explicitly for $n \in \mathbb{N}$.

Answer

1. $Q(x, x - 1) = 1 - p$, $Q(x, x + 1) = p$ for $x \in \mathbb{Z}$.
2. For $k \in \{0, 1, \dots, n\}$

$$Q^n(x, x + 2k - n) = \binom{n}{k} p^k (1 - p)^{n-k} \quad (16.3.52)$$

This corresponds to k steps to the right and $n - k$ steps to the left.

This special case is the *simple random walk* on \mathbb{Z} . When $p = \frac{1}{2}$ we have the *simple, symmetric random walk*. The simple random walk on \mathbb{Z} is studied in more detail in the section on random walks on graphs. The simple symmetric random walk is studied in more detail in the chapter on Bernoulli Trials.

Doubly Stochastic Matrices

A matrix P on S is *doubly stochastic* if it is nonnegative and if the row and columns sums are 1:

$$\sum_{u \in S} P(x, u) = 1, \quad \sum_{u \in S} P(u, y) = 1, \quad (x, y) \in S \times S \quad (16.3.53)$$

Suppose that \mathbf{X} is a Markov chain on a finite state space S with doubly stochastic transition matrix P . Then the uniform distribution on S is invariant.

Proof

Constant functions are left invariant. Suppose that $f(x) = c$ for $x \in S$. Then

$$fP(y) = \sum_{x \in S} f(x)P(x, y) = c \sum_{x \in S} P(x, y) = c, \quad y \in S \quad (16.3.54)$$

Hence if S is finite, the uniform PDF f given by $f(x) = 1/\#(S)$ for $x \in S$ is invariant.

If P and Q are doubly stochastic matrices on S , then so is PQ .

Proof

For $y \in S$,

$$\sum_{x \in S} PQ(x, y) = \sum_{x \in S} \sum_{z \in S} P(x, z)Q(z, y) = \sum_{z \in S} Q(z, y) \sum_{x \in S} P(x, z) = \sum_{z \in S} Q(z, y) = 1 \quad (16.3.55)$$

The interchange of sums is valid since the terms are nonnegative.

It follows that if P is doubly stochastic then so is P^n for $n \in \mathbb{N}$.

Suppose that $\mathbf{X} = (X_0, X_1, \dots)$ is the Markov chain with state space $S = \{-1, 0, 1\}$ and with transition matrix

$$P = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix} \quad (16.3.56)$$

1. Draw the state graph.
2. Show that P is doubly stochastic
3. Find P^2 .
4. Show that the uniform distribution on S is the only invariant distribution for \mathbf{X} .
5. Suppose that X_0 has the uniform distribution on S . For $n \in \mathbb{N}$, find $\mathbb{E}(X_n)$ and $\text{var}(X_n)$.
6. Find the α -potential matrix R_α for $\alpha \in (0, 1)$.

Proof

1. The edge set is $E = \{(-1, -1), (-1, 0), (0, 0), (0, 1), (1, -1), (1, 1)\}$
2. Just note that the row sums and the column sums are 1.
3. By matrix multiplication,

$$P^2 = \begin{bmatrix} \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{bmatrix} \quad (16.3.57)$$

4. Let $f = [p \quad q \quad r]$. Solving the equation $fP = f$ gives $p = q = r$. The requirement that f be a PDF then forces the common value to be $\frac{1}{3}$.
5. If X_0 has the uniform distribution on S , then so does X_n for every $n \in \mathbb{N}$, so $\mathbb{E}(X_n) = 0$ and $\text{var}(X_n) = \mathbb{E}(X_0^2) = \frac{2}{3}$.

6.

$$R_\alpha = (I - \alpha P)^{-1} = \frac{1}{(1 - \alpha)(4 - 2\alpha + \alpha^2)} \begin{bmatrix} 4 - 4\alpha + \alpha^2 & 2\alpha - \alpha^2 & \alpha^2 \\ \alpha^2 & 4 - 4\alpha + \alpha^2 & 2\alpha - \alpha^2 \\ 2\alpha - \alpha^2 & \alpha^2 & 4 - 4\alpha + \alpha^2 \end{bmatrix} \quad (16.3.58)$$

Recall that a matrix M indexed by a countable set S is *symmetric* if $M(x, y) = M(y, x)$ for all $x, y \in S$.

If P is a symmetric, stochastic matrix then P is doubly stochastic.

Proof

This is trivial since

$$\sum_{x \in S} P(x, y) = \sum_{x \in S} P(y, x) = 1, \quad y \in S \quad (16.3.59)$$

The converse is not true. The doubly stochastic matrix in the [exercise above](#) is not symmetric. But since a symmetric, stochastic matrix on a finite state space is doubly stochastic, the uniform distribution is invariant.

Suppose that $\mathbf{X} = (X_0, X_1, \dots)$ is the Markov chain with state space $S = \{-1, 0, 1\}$ and with transition matrix

$$P = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{4} & \frac{3}{4} \\ 0 & \frac{3}{4} & \frac{1}{4} \end{bmatrix} \quad (16.3.60)$$

1. Draw the state graph.
2. Show that P is symmetric
3. Find P^2 .
4. Find all invariant probability density functions for \mathbf{X} .
5. Find the α -potential matrix R_α for $\alpha \in (0, 1)$.

Proof

1. The edge set is $E = \{(-1, -1), (0, 0), (0, 1), (1, 0), (1, 1)\}$
2. Just note that P is symmetric with respect to the main diagonal.
3. By matrix multiplication,

$$P^2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{5}{8} & \frac{3}{8} \\ 0 & \frac{3}{8} & \frac{5}{8} \end{bmatrix} \quad (16.3.61)$$

4. Let $f = [p \quad q \quad r]$. Solving the equation $fP = f$ gives simply $r = q$. The requirement that f be a PDF forces $p = 1 - 2q$. Thus the invariant PDFs are $f = [1 - 2q \quad q \quad q]$ where $q \in [0, \frac{1}{2}]$. The special case $q = \frac{1}{3}$ gives the uniform distribution on S .

5.

$$R_\alpha = (I - \alpha P)^{-1} = \frac{1}{2(1 - \alpha)^2(2 + \alpha)} \begin{bmatrix} 4 - 2\alpha - 2\alpha^2 & 0 & 0 \\ 0 & 4 - 5\alpha + \alpha^2 & 3\alpha - 3\alpha^2 \\ 0 & 3\alpha - 3\alpha^2 & 4 - 5\alpha + \alpha^2 \end{bmatrix} \quad (16.3.62)$$

Special Models

The Markov chains in the following exercises model interesting processes that are studied in separate sections.

Read the introduction to the Ehrenfest chains.

Read the introduction to the Bernoulli-Laplace chain.

Read the introduction to the reliability chains.

Read the introduction to the branching chain.

Read the introduction to the queuing chains.

Read the introduction to random walks on graphs.

Read the introduction to birth-death chains.

This page titled [16.3: Introduction to Discrete-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.4: Transience and Recurrence for Discrete-Time Chains

The study of discrete-time Markov chains, particularly the limiting behavior, depends critically on the random times between visits to a given state. The nature of these random times leads to a fundamental dichotomy of the states.

Basic Theory

As usual, our starting point is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so that Ω is the sample space, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on (Ω, \mathcal{F}) . Suppose now that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a (homogeneous) discrete-time Markov chain with (countable) state space S and transition probability matrix P . So by definition,

$$P(x, y) = \mathbb{P}(X_{n+1} = y \mid X_n = x) \quad (16.4.1)$$

for $x, y \in S$ and $n \in \mathbb{N}$. Let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$, the σ -algebra of events defined by the chain up to time $n \in \mathbb{N}$, so that $\mathcal{F} = (\mathcal{F}_0, \mathcal{F}_1, \dots)$ is the natural filtration associated with \mathbf{X} .

Hitting Times and Probabilities

Let A be a nonempty subset of S . Recall that the *hitting time* to A is the random variable that gives the first positive time that the chain is in A :

$$\tau_A = \min\{n \in \mathbb{N}_+ : X_n \in A\} \quad (16.4.2)$$

Since the chain may *never* enter A , the random variable τ_A takes values in $\mathbb{N}_+ \cup \{\infty\}$ (recall our convention that the minimum of the empty set is ∞). Recall also that τ_A is a stopping time for \mathbf{X} . That is, $\{\tau_A = n\} \in \mathcal{F}_n$ for $n \in \mathbb{N}_+$. Intuitively, this means that we can tell if $\tau_A = n$ by observing the chain up to time n . This is clearly the case, since explicitly

$$\{\tau_A = n\} = \{X_1 \notin A, \dots, X_{n-1} \notin A, X_n \in A\} \quad (16.4.3)$$

When $A = \{x\}$ for $x \in S$, we will simplify the notation to τ_x . This random variable gives the first positive time that the chain is in state x . When the chain enters a set of states A for the first time, the chain must visit some state in A for the first time, so it's clear that

$$\tau_A = \min\{\tau_x : x \in A\}, \quad A \subseteq S \quad (16.4.4)$$

Next we define two functions on S that are related to the hitting times.

For $x \in S$, $A \subseteq S$ (nonempty), and $n \in \mathbb{N}_+$ define

1. $H_n(x, A) = \mathbb{P}(\tau_A = n \mid X_0 = x)$
2. $H(x, A) = \mathbb{P}(\tau_A < \infty \mid X_0 = x)$

So $H(x, A) = \sum_{n=1}^{\infty} H_n(x, A)$.

Note that $n \mapsto H_n(x, A)$ is the probability density function of τ_A , given $X_0 = x$, except that the density function may be *defective* in the sense that the sum $H(x, A)$ may be less than 1, in which case of course, $1 - H(x, A) = \mathbb{P}(\tau_A = \infty \mid X_0 = x)$. Again, when $A = \{y\}$, we will simplify the notation to $H_n(x, y)$ and $H(x, y)$, respectively. In particular, $H(x, x)$ is the probability, starting at x , that the chain eventually returns to x . If $x \neq y$, $H(x, y)$ is the probability, starting at x , that the chain eventually reaches y . Just knowing when $H(x, y)$ is 0, positive, and 1 will turn out to be of considerable importance in the overall structure and limiting behavior of the chain. As a function on S^2 , we will refer to H as the *hitting matrix* of \mathbf{X} . Note however, that unlike the transition matrix P , we do not have the structure of a kernel. That is, $A \mapsto H(x, A)$ is not a measure, so in particular, it is generally not true that $H(x, A) = \sum_{y \in A} H(x, y)$. The same remarks apply to H_n for $n \in \mathbb{N}_+$. However, there are interesting relationships between the transition matrix and the hitting matrix.

$H(x, y) > 0$ if and only if $P^n(x, y) > 0$ for some $n \in \mathbb{N}_+$.

Proof

Note that $\{X_n = y\} \subseteq \{\tau_y < \infty\}$ for all $n \in \mathbb{N}_+$, and $\{\tau_y < \infty\} = \{X_k = y \text{ for some } k \in \mathbb{N}_+\}$. From the increasing property of probability and Boole's inequality it follows that for each $n \in \mathbb{N}_+$,

$$P^n(x, y) \leq H(x, y) \leq \sum_{k=1}^{\infty} P^k(x, y) \quad (16.4.5)$$

The following result gives a basic relationship between the sequence of hitting probabilities and the sequence of transition probabilities.

Suppose that $(x, y) \in S^2$. Then

$$P^n(x, y) = \sum_{k=1}^n H_k(x, y) P^{n-k}(y, y), \quad n \in \mathbb{N}_+ \quad (16.4.6)$$

Proof

This result follows from conditioning on τ_y . Starting in state x , the chain is in state y at time n if and only if the chain hits y for the first time at some previous time k , and then returns to y in the remaining $n - k$ steps. More formally,

$$P^n(x, y) = \mathbb{P}(X_n = y \mid X_0 = x) = \sum_{k=0}^{\infty} \mathbb{P}(X_n = y \mid \tau_y = k, X_0 = x) \mathbb{P}(\tau_y = k \mid X_0 = x) \quad (16.4.7)$$

But the event $\tau_y = k$ implies $X_k = y$ and is in \mathcal{F}_k . Hence by the Markov property,

$$\mathbb{P}(X_n = y \mid \tau_y = k, X_0 = x) = \mathbb{P}(X_n = y \mid X_k = y, \tau_y = k, X_0 = x) = \mathbb{P}(X_n = y \mid X_k = y) = P^k(x, y) \quad (16.4.8)$$

Of course, by definition, $\mathbb{P}(\tau_y = k \mid X_0 = x) = H_k(x, y)$, so the result follows by substitution.

Suppose that $x \in S$ and $A \subseteq S$. Then

1. $H_{n+1}(x, A) = \sum_{y \notin A} P(x, y) H_n(y, A)$ for $n \in \mathbb{N}_+$
2. $H(x, A) = P(x, A) + \sum_{y \notin A} P(x, y) H(y, A)$

Proof

These results follow from conditioning on X_1 .

1. Starting in state x , the chain first enters A at time $n + 1$ if and only if the chain goes to some state $y \notin A$ at time 1, and then from state y , first enters A in n steps.

$$H_{n+1}(x, A) = \mathbb{P}(\tau_A = n + 1 \mid X_0 = x) = \sum_{y \in S} \mathbb{P}(\tau_A = n + 1 \mid X_0 = x, X_1 = y) \mathbb{P}(X_1 = y \mid X_0 = x) \quad (16.4.9)$$

But $\mathbb{P}(\tau_A = n + 1 \mid X_0 = x, X_1 = y) = 0$ for $y \in A$. By the Markov and time homogeneous properties,

$\mathbb{P}(\tau_A = n + 1 \mid X_0 = x, X_1 = y) = \mathbb{P}(\tau_A = n \mid X_0 = y) = H_n(x, A)$ for $y \notin A$. Of course

$\mathbb{P}(X_1 = y \mid X_0 = x) = P(x, y)$. So the result follows by substitution.

2. Starting in state x , the chain eventually enters A if and only if it either enters A at the first step, or moves to some other state $y \notin A$ at the first step, and then eventually enters A from y .

$$H(x, A) = \mathbb{P}(\tau_A < \infty \mid X_0 = x) = \sum_{y \in S} \mathbb{P}(\tau_A < \infty \mid X_1 = y, X_0 = x) \mathbb{P}(X_1 = y \mid X_0 = x) \quad (16.4.10)$$

But $\mathbb{P}(\tau_A < \infty \mid X_1 = y, X_0 = x) = 1$ for $y \in A$. By the Markov and homogeneous properties,

$\mathbb{P}(\tau_A < \infty \mid X_1 = y, X_0 = x) = \mathbb{P}(\tau_A < \infty \mid X_0 = y) = H(y, A)$ for $y \notin A$. Substituting we have

$$H(x, A) = \sum_{y \in A} P(x, y) + \sum_{y \notin A} P(x, y) H(y, A) = P(x, A) + \sum_{y \notin A} P(x, y) H(y, A) \quad (16.4.11)$$

The following definition is fundamental for the study of Markov chains.

Let $x \in S$.

1. State x is *recurrent* if $H(x, x) = 1$.
2. State x is *transient* if $H(x, x) < 1$.

Thus, starting in a recurrent state, the chain will, with probability 1, eventually return to the state. As we will see, the chain will return to the state infinitely often with probability 1, and the times of the visits will form the arrival times of a renewal process. This will turn out to be the critical observation in the study of the limiting behavior of the chain. By contrast, if the chain starts in a transient state, then there is a positive probability that the chain will never return to the state.

Counting Variables and Potentials

Again, suppose that A is a nonempty set of states. A natural complement to the hitting time to A is the *counting variable* that gives the number of visits to A (at positive times). Thus, let

$$N_A = \sum_{n=1}^{\infty} \mathbf{1}(X_n \in A) \quad (16.4.12)$$

Note that N_A takes value in $\mathbb{N} \cup \{\infty\}$. We will mostly be interested in the special case $A = \{x\}$ for $x \in S$, and in this case, we will simplify the notation to N_x .

Let $G(x, A) = \mathbb{E}(N_A \mid X_0 = x)$ for $x \in S$ and $A \subseteq S$. Then G is a kernel on S and

$$G(x, A) = \sum_{n=1}^{\infty} P^n(x, A) \quad (16.4.13)$$

Proof

Note that

$$G(x, A) = \mathbb{E} \left(\sum_{n=1}^{\infty} \mathbf{1}(X_n \in A) \mid X_0 = x \right) = \sum_{n=1}^{\infty} \mathbb{P}(X_n \in A \mid X_0 = x) = \sum_{n=1}^{\infty} P^n(x, A) \quad (16.4.14)$$

The interchange of sum and expected value is justified since the terms are nonnegative. For fixed $x \in S$, $A \mapsto G(x, A)$ is a positive measure on S since $A \mapsto P^n(x, A)$ is a probability measure on S for each $n \in \mathbb{N}_+$. Note also that $A \mapsto N_A$ is a random, counting measure on S and hence $A \mapsto G(x, A)$ is a (deterministic) positive measure on S .

Thus $G(x, A)$ is the expected number of visits to A at positive times. As usual, when $A = \{y\}$ for $y \in S$ we simplify the notation to $G(x, y)$, and then more generally we have $G(x, A) = \sum_{y \in A} G(x, y)$ for $A \subseteq S$. So, as a matrix on S , $G = \sum_{n=1}^{\infty} P^n$. The matrix G is closely related to the potential matrix R of \mathbf{X} , given by $R = \sum_{n=0}^{\infty} P^n$. So $R = I + G$, and $R(x, y)$ gives the expected number of visits to $y \in S$ at all times (not just positive times), starting at $x \in S$. The matrix G is more useful for our purposes in this section.

The distribution of N_y has a simple representation in terms of the hitting probabilities. Note that because of the Markov property and time homogeneous property, whenever the chain reaches state y , the future behavior is independent of the past and is stochastically the same as the chain starting in state y at time 0. This is the critical observation in the proof of the following theorem.

If $x, y \in S$ then

1. $\mathbb{P}(N_y = 0 \mid X_0 = x) = 1 - H(x, y)$
2. $\mathbb{P}(N_y = n \mid X_0 = x) = H(x, y)[H(y, y)]^{n-1}[1 - H(y, y)]$ for $n \in \mathbb{N}_+$

Proof




Figure 16.4.1: Visits to state y starting in state x

The essence of the proof is illustrated in the graphic above. The thick lines are intended as reminders that these are not one step transitions, but rather represent all paths between the given vertices. Note that in the special case that $x = y$ we have

$$\mathbb{P}(N_x = n \mid X_0 = x) = [H(x, x)]^n [1 - H(x, x)], \quad n \in \mathbb{N} \quad (16.4.15)$$

In all cases, the counting variable N_y has essentially a geometric distribution, but the distribution may well be *defective*, with some of the probability mass at ∞ . The behavior is quite different depending on whether y is transient or recurrent.

If $x, y \in S$ and y is transient then

1. $\mathbb{P}(N_y < \infty \mid X_0 = x) = 1$
2. $G(x, y) = H(x, y) / [1 - H(y, y)]$
3. $H(x, y) = G(x, y) / [1 + G(y, y)]$

Proof

1. If y is transient then $H(y, y) < 1$. Hence using the result [above](#) and geometric series,

$$\mathbb{P}(N_y \in \mathbb{N}_+ \mid X_0 = x) = \sum_{n=1}^{\infty} \mathbb{P}(N_y = n \mid X_0 = x) = H(x, y)[1 - H(y, y)] \sum_{n=1}^{\infty} [H(y, y)]^{n-1} = H(x, y) \quad (16.4.16)$$

Hence

$$\begin{aligned} \mathbb{P}(N_y < \infty \mid X_0 = x) &= \mathbb{P}(N_y \in \mathbb{N} \mid X_0 = x) = \mathbb{P}(N_y = 0 \mid X_0 = x) + \mathbb{P}(N_y \in \mathbb{N}_+ \mid X_0 = x) \\ &= [1 - H(x, y)] + H(x, y) = 1 \end{aligned}$$

2. Using the derivative of the geometric series,

$$\begin{aligned} G(x, y) &= \mathbb{E}(N_y \mid X_0 = x) = \sum_{n=1}^{\infty} n \mathbb{P}(N_y = n \mid X_0 = x) \\ &= H(x, y)[1 - H(y, y)] \sum_{n=1}^{\infty} n [H(y, y)]^{n-1} = \frac{H(x, y)}{1 - H(y, y)} \end{aligned}$$

3. From (b), $G(y, y) = H(y, y)/[1 - H(y, y)]$ so solving for $H(y, y)$ gives $H(y, y) = G(y, y)/[1 + G(y, y)]$. Substituting this back into (b) gives $G(x, y) = H(x, y)[1 + G(y, y)]$.

if $x, y \in S$ and y is recurrent then

1. $\mathbb{P}(N_y = 0 \mid X_0 = x) = 1 - H(x, y)$ and $\mathbb{P}(N_y = \infty \mid X_0 = x) = H(x, y)$
2. $G(x, y) = 0$ if $H(x, y) = 0$ and $G(x, y) = \infty$ if $H(x, y) > 0$
3. $\mathbb{P}(N_y = \infty \mid X_0 = y) = 1$ and $G(y, y) = \infty$

Proof

1. If y is recurrent, $H(y, y) = 1$ and so from the result [above](#), $\mathbb{P}(N_y = n \mid X_0 = x) = 0$ for all $n \in \mathbb{N}_+$. Hence $\mathbb{P}(N_y = \infty \mid X_0 = x) = 1 - \mathbb{P}(N_y = 0 \mid X_0 = x) = 1 - H(x, y)$.
2. If $H(x, y) = 0$ then $\mathbb{P}(N_y = 0 \mid X_0 = x) = 1$, so $\mathbb{E}(N_y \mid X_0 = x) = 0$. If $H(x, y) > 0$ then $\mathbb{P}(N_y = \infty \mid X_0 = x) > 0$ so $\mathbb{E}(N_y \mid X_0 = x) = \infty$.
3. From the result above, $\mathbb{P}(N_y = n \mid X_0 = y) = 0$ for all $n \in \mathbb{N}$, so $\mathbb{P}(N_y = \infty \mid X_0 = y) = 1$.

Note that there is an invertible relationship between the matrix H and the matrix G ; if we know one we can compute the other. In particular, we can characterize the transience or recurrence of a state in terms of G . Here is our summary so far:

Let $x \in S$.

1. State x is transient if and only if $H(x, x) < 1$ if and only if $G(x, x) < \infty$.
2. State x is recurrent if and only if $H(x, x) = 1$ if and only if $G(x, x) = \infty$.

Of course, the classification also holds for the potential matrix $R = I + G$. That is, state $x \in S$ is transient if and only if $R(x, x) < \infty$ and state x is recurrent if and only if $R(x, x) = \infty$.

Relations

The hitting probabilities suggest an important relation on the state space S .

For $(x, y) \in S^2$, we say that x leads to y and we write $x \rightarrow y$ if either $x = y$ or $H(x, y) > 0$.

It follows immediately from the result [above](#) that $x \rightarrow y$ if and only if $P^n(x, y) > 0$ for some $n \in \mathbb{N}$. In terms of the state graph of the chain, $x \rightarrow y$ if and only if $x = y$ or there is a directed path from x to y . Note that the leads to relation is *reflexive* by definition: $x \rightarrow x$ for every $x \in S$. The relation has another important property as well.

The leads to relation is *transitive*: For $x, y, z \in S$, if $x \rightarrow y$ and $y \rightarrow z$ then $x \rightarrow z$.

Proof

If $x \rightarrow y$ and $y \rightarrow z$, then there exist $j, k \in \mathbb{N}$ such that $P^j(x, y) > 0$ and $P^k(y, z) > 0$. But then $P^{j+k}(x, z) \geq P^j(x, y)P^k(y, z) > 0$ so $x \rightarrow z$.

The *leads to* relation naturally suggests a couple of other definitions that are important.

Suppose that $A \subseteq S$ is nonempty.

1. A is *closed* if $x \in A$ and $x \rightarrow y$ implies $y \in A$.
2. A is *irreducible* if A is closed and has no proper closed subsets.

Suppose that $A \subseteq S$ is closed. Then

1. P_A , the restriction of P to $A \times A$, is a transition probability matrix on A .
2. \mathbf{X} restricted to A is a Markov chain with transition probability matrix P_A .
3. $(P^n)_A = (P_A)^n$ for $n \in \mathbb{N}$.

Proof

1. If $x \in A$ and $y \notin A$, then x does not lead to y so in particular $P(x, y) = 0$. It follows that $\sum_{y \in A} P(x, y) = 1$ for $x \in A$ so P_A is a transition probability matrix.
2. This follows from (a). If the chain starts in A , then the chain remains in A for all time, and of course, the Markov property still holds.
3. Again, this follows from (a).

Of course, the entire state space S is closed by definition. If it is also irreducible, we say the Markov chain \mathbf{X} itself is *irreducible*. Recall that for a nonempty subset A of S and for $n \in \mathbb{N}$, the notation P_A^n refers to $(P_A)^n$ and not $(P^n)_A$. In general, these are not the same, and in fact for $x, y \in A$,

$$P_A^n(x, y) = \mathbb{P}(X_1 \in A, \dots, X_{n-1} \in A, X_n = y \mid X_0 = x) \quad (16.4.17)$$

the probability of going from x to y in n steps, remaining in A all the while. But if A is closed, then as noted in part (c), this is just $P^n(x, y)$.

Suppose that A is a nonempty subset of S . Then $\text{cl}(A) = \{y \in S : x \rightarrow y \text{ for some } x \in A\}$ is the smallest closed set containing A , and is called the *closure* of A . That is,

1. $\text{cl}(A)$ is closed.
2. $A \subseteq \text{cl}(A)$.
3. If B is closed and $A \subseteq B$ then $\text{cl}(A) \subseteq B$.

Proof

1. Suppose that $x \in \text{cl}(A)$ and that $x \rightarrow y$. Then there exists $a \in A$ such that $a \rightarrow x$. By the transitive property, $a \rightarrow y$ and hence $y \in \text{cl}(A)$.
2. If $x \in A$ then $x \rightarrow x$ so $x \in \text{cl}(A)$.
3. Suppose that B is closed and that $A \subseteq B$. If $x \in \text{cl}(A)$, then there exists $a \in A$ such that $a \rightarrow x$. Hence $a \in B$ and $a \rightarrow x$. Since B is closed, it follows that $x \in B$. Hence $\text{cl}(A) \subseteq B$.

Recall that for a fixed positive integer k , P^k is also a transition probability matrix, and in fact governs the *k-step* Markov chain $(X_0, X_k, X_{2k}, \dots)$. It follows that we could consider the *leads to* relation for this chain, and all of the results above would still hold (relative, of course, to the *k-step* chain). Occasionally we will need to consider this relation, which we will denote by $\xrightarrow[k]$, particularly in our study of periodicity.

Suppose that $j, k \in \mathbb{N}_+$. If $x \xrightarrow[k]{} y$ and $j \mid k$ then $x \xrightarrow[j]{} y$.

Proof

If $x \xrightarrow[k]{} y$ then there exists $n \in \mathbb{N}$ such that $P^{nk}(x, y) > 0$. If $j \mid k$, there exists $m \in \mathbb{N}_+$ such that $k = mj$. Hence $P^{nmj}(x, y) > 0$ so $x \xrightarrow[j]{} y$.

By combining the *leads to* relation \rightarrow with its inverse, the *comes from* relation \leftarrow , we can obtain another very useful relation.

For $(x, y) \in S^2$, we say that x *to and from* y and we write $x \leftrightarrow y$ if $x \rightarrow y$ and $y \rightarrow x$.

By definition, this relation is *symmetric*: if $x \leftrightarrow y$ then $y \leftrightarrow x$. From our work above, it is also reflexive and transitive. Thus, the *to and from* relation is an equivalence relation. Like all equivalence relations, it partitions the space into mutually disjoint *equivalence classes*. We will denote the equivalence class of a state $x \in S$ by

$$[x] = \{y \in S : x \leftrightarrow y\} \quad (16.4.18)$$

Thus, for any two states $x, y \in S$, either $[x] = [y]$ or $[x] \cap [y] = \emptyset$, and moreover, $\bigcup_{x \in S} [x] = S$.




Figure 16.4.2: The equivalence relation partitions S into mutually disjoint equivalence classes

Two negative results:

1. A closed set is not necessarily an equivalence class.
2. An equivalence class is not necessarily closed.

Example

Consider the trivial Markov chain with state space $S = \{0, 1\}$ and transition matrix $P = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$. So state 0 leads deterministically to 1 in one step, while state 1 is absorbing. For the *leads to* relation, the only relationships are $0 \rightarrow 0$, $0 \rightarrow 1$, and $1 \rightarrow 1$. Thus, the equivalence classes are $\{0\}$ and $\{1\}$.

1. The entire state space S is closed, but is not an equivalence class.
2. $\{0\}$ is an equivalence class but is not closed.

On the other hand, we have the following result:

If $A \subseteq S$ is irreducible, then A is an equivalence class.

Proof

Fix $x \in A$ (recall that closed sets are nonempty by definition). Since A is closed it follows that $[x] \subseteq A$. Since A is irreducible, $\text{cl}(y) = A$ for each $y \in A$ and in particular, $\text{cl}(x) = A$. It follows that $x \leftrightarrow y$ for each $y \in A$. Hence $A \subseteq [x]$.

The *to and from* equivalence relation is very important because many interesting state properties turn out in fact to be *class properties*, shared by all states in a given equivalence class. In particular, the recurrence and transience properties are class properties.

Transient and Recurrent Classes

Our next result is of fundamental importance: a recurrent state can only lead to other recurrent states.

If x is a recurrent state and $x \rightarrow y$ then y is recurrent and $H(x, y) = H(y, x) = 1$.

Proof

The result trivially holds if $x = y$, so we assume $x \neq y$. Let $\alpha(x, y)$ denote the probability, starting at x , that the chain reaches y without an intermediate return to x . It must be the case that $\alpha(x, y) > 0$ since $x \rightarrow y$. In terms of the graph of \mathbf{X} , if there is a path from x to y , then there is a path from x to y without cycles. Starting at x , the chain could fail to return to x by first reaching y without an intermediate return to x , and then from y never reaching x . From the Markov and time homogeneous properties, it follows that $1 - H(x, x) \geq \alpha(x, y)[1 - H(y, x)] \geq 0$. But $H(x, x) = 1$ so it follows that $H(y, x) = 1$. So we now know that there exist positive integers j, k such that $P^j(x, y) > 0$ and $P^k(y, x) > 0$. Hence for every $n \in \mathbb{N}$,

$$P^{j+k+n}(y, y) \geq P^k(y, x)P^n(x, x)P^j(x, y) \quad (16.4.19)$$

Recall that $G(x, x) = \infty$ since x is recurrent. Thus, summing over n in the displayed equation gives $G(y, y) = \infty$. Hence y is recurrent. Finally, reversing the roles of x and y , it follows that $H(x, y) = 1$.

From the last theorem, note that if x is recurrent, then all states in $[x]$ are also recurrent. Thus, for each equivalence class, either all states are transient or all states are recurrent. We can therefore refer to *transient or recurrent classes* as well as states.

If A is a recurrent equivalence class then A is irreducible.

Proof

Suppose that $x \in A$ and that $x \rightarrow y$. Since x is recurrent, y is also recurrent and $y \rightarrow x$. Hence $x \leftrightarrow y$ and so $y \in A$ since A is an equivalence class. Suppose that $B \subseteq A$ is closed. Since B is nonempty by definition, there exists $x \in B$ and so $x \in A$ also. For every $y \in A$, $x \leftrightarrow y$ so $y \in B$ since B is closed. Thus $A = B$ so A is irreducible.

If A is finite and closed then A has a recurrent state.

Proof

Fix $x \in A$. Since A is closed, it follows that $\mathbb{P}(N_A = \infty \mid X_0 = x) = 1$. Since A is finite, it follows that $\mathbb{P}(N_y = \infty \mid X_0 = x) > 0$ for some $y \in A$. But then y is recurrent.

If A is finite and irreducible then A is a recurrent equivalence class.

Proof

Note that A is an equivalence class by a result [above](#), and A has a recurrent state by [previous result](#). It follows that all states in A are recurrent.

Thus, the Markov chain \mathbf{X} will have a collection (possibly empty) of recurrent equivalence classes $\{A_j : j \in J\}$ where J is a countable index set. Each A_j is irreducible. Let B denote the set of all transient states. The set B may be empty or may consist of a number of equivalence classes, but the class structure of B is usually not important to us. If the chain starts in A_j for some $j \in J$ then the chain remains in A_j forever, visiting each state infinitely often with probability 1. If the chain starts in B , then the chain may stay in B forever (but only if B is infinite) or may enter one of the recurrent classes A_j , never to escape. However, in either case, the chain will visit a given transient state only finitely many time with probability 1. This basic structure is known as the *canonical decomposition* of the chain, and is shown in graphical form below. The edges from B are in gray to indicate that these transitions may not exist.




Figure 16.4.3: The canonical decomposition of the state space

Staying Probabilities and a Classification Test

Suppose that A is a proper subset of S . Then

1. $P_A^n(x, A) = \mathbb{P}(X_1 \in A, X_2 \in A, \dots, X_n \in A \mid X_0 = x)$ for $x \in A$
2. $\lim_{n \rightarrow \infty} P_A^n(x, A) = \mathbb{P}(X_1 \in A, X_2 \in A, \dots \mid X_0 = x)$ for $x \in A$

Proof

Recall that P_A^n means $(P_A)^n$ where P_A is the restriction of P to $A \times A$.

1. This is a consequence of the Markov property, and is the probability that the chain stays in A at least through time n , starting in $x \in A$.
2. This follows from (a) and the continuity theorem for decreasing events. This is the probability that the chain stays in A forever, starting in $x \in A$.

Let g_A denote the function defined by part (b), so that

$$g_A(x) = \mathbb{P}(X_1 \in A, X_2 \in A, \dots \mid X_0 = x), \quad x \in A \quad (16.4.20)$$

The *staying probability function* g_A is an interesting complement to the hitting matrix studied above. The following result characterizes this function and provides a method that can be used to compute it, at least in some cases.

For $A \subset S$, g_A is the largest function on A that takes values in $[0, 1]$ and satisfies $g = P_A g$. Moreover, either $g_A = \mathbf{0}_A$ or $\sup\{g_A(x) : x \in A\} = 1$.

Proof

Note that $P_A^{n+1} \mathbf{1}_A = P_A P_A^n \mathbf{1}_A$ for $n \in \mathbb{N}$. Taking the limit as $n \rightarrow \infty$ and using the bounded convergence theorem gives $g_A = P_A g_A$. Suppose now that g is a function on A that takes values in $[0, 1]$ and satisfies $g = P_A g$. Then $g \leq \mathbf{1}_A$ and hence $g \leq P_A^n \mathbf{1}_A$ for all $n \in \mathbb{N}$. Letting $n \rightarrow \infty$ it follows that $g \leq g_A$. Next, let $c = \sup\{g_A(x) : x \in A\}$. Then $g_A \leq c \mathbf{1}_A$ and hence $g_A \leq c P_A^n \mathbf{1}_A$ for each $n \in \mathbb{N}$. Letting $n \rightarrow \infty$ gives $g_A \leq c g_A$. It follows that either $g_A = \mathbf{0}_A$ or $c = 1$.

Note that the characterization in the last result includes a *zero-one law* of sorts: either the probability that the chain stays in A forever is 0 for every initial state $x \in A$, or we can find states in A for which the probability is arbitrarily close to 1. The next two results explore the relationship between the staying function and recurrence.

Suppose that \mathbf{X} is an irreducible, recurrent chain with state space S . Then $g_A = \mathbf{0}_A$ for every proper subset A of S .

Proof

Fix $y \notin A$ and note that $0 \leq g_A(x) \leq 1 - H(x, y)$ for every $x \in A$. But $H(x, y) = 1$ since the chain is irreducible and recurrent. Hence $g_A(x) = 0$ for $x \in A$.

Suppose that \mathbf{X} is an irreducible Markov chain with state space S and transition probability matrix P . If there exists a state x such that $g_A = \mathbf{0}_A$ where $A = S \setminus \{x\}$, then \mathbf{X} is recurrent.

Proof

With A as defined above, note that $1 - H(x, x) = \sum_{y \in A} P(x, y) g_A(y)$. Hence $H(x, x) = 1$, so x is recurrent. Since the \mathbf{X} is irreducible, it follows that \mathbf{X} is recurrent.

More generally, suppose that \mathbf{X} is a Markov chain with state space S and transition probability matrix P . The last two theorems can be used to test whether an irreducible equivalence class C is recurrent or transient. We fix a state $x \in C$ and set $A = C \setminus \{x\}$. We then try to solve the equation $g = P_A g$ on A . If the only solution taking values in $[0, 1]$ is $\mathbf{0}_A$, then the class C is recurrent by the [previous result](#). If there are nontrivial solutions, then C is transient. Often we try to choose x to make the computations easy.

Computing Hitting Probabilities and Potentials

We now know quite a bit about Markov chains, and we can often classify the states and compute quantities of interest. However, we do not yet know how to compute:

- $G(x, y)$ when x and y are transient
- $H(x, y)$ when x is transient and y is transient or recurrent.

These problems are related, because of the general inverse relationship between the matrix H and the matrix G noted in our discussion above. As usual, suppose that \mathbf{X} is a Markov chain with state space S , and let B denote the set of transient states. The next result shows how to compute G_B , the matrix G restricted to the transient states. Recall that the values of this matrix are finite.

G_B satisfies the equation $G_B = P_B + P_B G_B$ and is the smallest nonnegative solution. If B is finite then $G_B = (I_B - P_B)^{-1} P_B$.

Proof

First note the $(P^n)_B = (P_B)^n$ since a path between two transient states can only pass through other transient states. Thus $G_B = \sum_{n=1}^{\infty} P_B^n$. From the monotone convergence theorem it follows that $P_B G_B = G_B - P_B$. Suppose now that U is a nonnegative matrix on B satisfying $U = P_B + P_B U$. Then $U = \sum_{k=1}^n P_B^k + P_B^{n+1} U$ for each $n \in \mathbb{N}_+$. Hence $U \geq \sum_{k=1}^n P_B^k$ for every $n \in \mathbb{N}_+$ and therefore $U \geq G_B$. It follows that $(I_B - P_B)(I_B + G_B) = I_B$. If B is finite, the matrix $I_B - P_B$ is invertible.

Now that we can compute G_B , we can also compute H_B using the result [above](#). All that remains is for us to compute the hitting probability $H(x, y)$ when x is transient and y is recurrent. The first thing to notice is that the hitting probability is a class property.

Suppose that x is transient and that A is a recurrent class. Then $H(x, y) = H(x, A)$ for $y \in A$.

That is, starting in the transient state $x \in S$, the hitting probability to y is constant for $y \in A$, and is just the hitting probability to the class A . As before, let B denote the set of transient states and suppose that A is a recurrent equivalence class. Let h_A denote the

function on B that gives the hitting probability to class A , and let p_A denote the function on B that gives the probability of entering A on the first step:

$$h_A(x) = H(x, A), \quad p_A(x) = P(x, A), \quad x \in B \quad (16.4.21)$$

$$h_A = p_A + G_B p_A.$$

Proof

First note that $\mathbb{P}(\tau_A = n \mid X_0 = x) = (P_B^{n-1} p_A)(x)$ for $n \in \mathbb{N}_+$. The result then follows by summing over n .

This result is adequate if we have already computed G_B (using the result in [above](#), for example). However, we might just want to compute h_A directly.

h_A satisfies the equation $h_A = p_A + P_B h_A$ and is the smallest nonnegative solution. If B is finite, $h_A = (I_B - P_B)^{-1} p_A$.

Proof

First, conditioning on X_1 gives $h_A = p_A + P_B h_A$. Next suppose that h is nonnegative and satisfies $h = p_A + P_B h$. Then $h = p_A + \sum_{k=1}^{n-1} P_B^k p_A + P_B^n h$ for each $n \in \mathbb{N}_+$. Hence $h \geq p_A + \sum_{k=1}^{n-1} P_B^k p_A$. Letting $n \rightarrow \infty$ gives $h \geq h_A$. The representation when B is finite follows from the [result above](#).

Examples and Applications

Finite Chains

Consider a Markov chain with state space $S = \{a, b, c, d\}$ and transition matrix P given below:

$$P = \begin{bmatrix} \frac{1}{2} & \frac{2}{3} & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{bmatrix} \quad (16.4.22)$$

1. Draw the state graph.
2. Find the equivalent classes and classify each as transient or recurrent.
3. Compute the matrix G .
4. Compute the matrix H .

Answer

1.

State graph


2. $\{a, b\}$ recurrent; $\{c\}$ recurrent; $\{d\}$ transient.

$$3. G = \begin{bmatrix} \infty & \infty & 0 & 0 \\ \infty & \infty & 0 & 0 \\ 0 & 0 & \infty & 0 \\ \infty & \infty & \infty & \frac{1}{3} \end{bmatrix}$$

$$4. H = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{2}{3} & \frac{2}{3} & \frac{1}{3} & \frac{1}{4} \end{bmatrix}$$

Consider a Markov chain with state space $S = \{1, 2, 3, 4, 5, 6\}$ and transition matrix P given below:

$$P = \begin{bmatrix} 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ \frac{1}{4} & 0 & \frac{1}{2} & 0 & \frac{1}{4} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 & \frac{2}{3} & 0 \\ 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} \end{bmatrix} \quad (16.4.23)$$

1. Sketch the state graph.
2. Find the equivalence classes and classify each as recurrent or transient.
3. Compute the matrix G .
4. Compute the matrix H .

Answer

1.

State graph



2. $\{1, 3, 5\}$ recurrent; $\{2, 6\}$ transient; $\{4\}$ recurrent.

$$3. G = \begin{bmatrix} \infty & 0 & \infty & 0 & \infty & 0 \\ \infty & \frac{1}{2} & \infty & \infty & \infty & 2 \\ \infty & 0 & \infty & 0 & \infty & 0 \\ 0 & 0 & 0 & \infty & 0 & 0 \\ \infty & 0 & \infty & 0 & \infty & 0 \\ \infty & \frac{1}{2} & \infty & \infty & \infty & 1 \end{bmatrix}$$

$$4. H = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

Consider a Markov chain with state space $S = \{1, 2, 3, 4, 5, 6\}$ and transition matrix P given below:

$$P = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\ \frac{1}{4} & \frac{3}{4} & 0 & 0 & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{4} & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix} \quad (16.4.24)$$

1. Sketch the state graph.
2. Find the equivalence classes and classify each as recurrent or transient.
3. Compute the matrix G .
4. Compute the matrix H .

Answer

1.

State graph



2. $\{1, 2\}$ recurrent; $\{3, 4\}$ transient; $\{5, 6\}$ recurrent.

$$3. G = \begin{bmatrix} \infty & \infty & 0 & 0 & 0 & 0 \\ \infty & \infty & 0 & 0 & 0 & 0 \\ \infty & \infty & \frac{7}{5} & \frac{4}{5} & \infty & \infty \\ \infty & \infty & \frac{4}{5} & \frac{3}{5} & \infty & \infty \\ 0 & 0 & 0 & 0 & \infty & \infty \\ 0 & 0 & 0 & 0 & \infty & \infty \end{bmatrix}$$

$$4. H = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ \frac{4}{5} & \frac{4}{5} & \frac{7}{12} & \frac{1}{2} & \frac{1}{5} & \frac{1}{5} \\ \frac{3}{5} & \frac{3}{5} & \frac{1}{3} & \frac{3}{8} & \frac{2}{5} & \frac{2}{5} \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}$$

Special Models

Read again the definitions of the Ehrenfest chains and the Bernoulli-Laplace chains. Note that since these chains are irreducible and have finite state spaces, they are recurrent.

Read the discussion on recurrence in the section on the reliability chains.

Read the discussion on random walks on \mathbb{Z}^k in the section on the random walks on graphs.

Read the discussion on extinction and explosion in the section on the branching chain.

Read the discussion on recurrence and transience in the section on queuing chains.

Read the discussion on recurrence and transience in the section on birth-death chains.

This page titled [16.4: Transience and Recurrence for Discrete-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.5: Periodicity of Discrete-Time Chains

A state in a discrete-time Markov chain is periodic if the chain can return to the state only at multiples of some integer larger than 1. Periodic behavior complicates the study of the limiting behavior of the chain. As we will see in this section, we can eliminate the periodic behavior by considering the d -step chain, where $d \in \mathbb{N}_+$ is the period, but only at the expense of introducing additional equivalence classes. Thus, in a sense, we can trade one form of complexity for another.

Basic Theory

Definitions and Basic Results

As usual, our starting point is a (time homogeneous) discrete-time Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ with (countable) state space S and transition probability matrix P .

The *period* of state $x \in S$ is

$$d(x) = \gcd\{n \in \mathbb{N}_+ : P^n(x, x) > 0\} \quad (16.5.1)$$

State x is *aperiodic* if $d(x) = 1$ and *periodic* if $d(x) > 1$.

Thus, starting in x , the chain can return to x only at multiples of the period d , and d is the largest such integer. Perhaps the most important result is that period, like recurrence and transience, is a *class property*, shared by all states in an equivalence class under the *to and from* relation.

If $x \leftrightarrow y$ then $d(x) = d(y)$.

Proof

Suppose that $x \leftrightarrow y$. The result is trivial if $x = y$, so let's assume that $x \neq y$. Recall that there exist $j, k \in \mathbb{N}_+$ such that $P^j(x, y) > 0$ and $P^k(y, x) > 0$. But then $P^{j+k}(x, x) \geq P^j(x, y)P^k(y, x) > 0$ and hence $d(x) \mid (j+k)$. Suppose now that n is a positive integer with $P^n(y, y) > 0$. Then $P^{j+k+n}(x, x) \geq P^j(x, y)P^n(y, y)P^k(y, x) > 0$ and hence $d(x) \mid (j+k+n)$. It follows that $d(x) \mid n$. From the definition of period, $d(y) \mid d(x)$. Reversing the roles of x and y we also have $d(x) \mid d(y)$. Hence $d(x) = d(y)$.

Thus, the definitions of *period*, *periodic*, and *aperiodic* apply to equivalence classes as well as individual states. When the chain is irreducible, we can apply these terms to the entire chain.

Suppose that $x \in S$. If $P(x, x) > 0$ then x (and hence the equivalence class of x) is aperiodic.

Proof

By assumption, $1 \in \{n \in \mathbb{N}_+ : P^n(x, x) > 0\}$ and hence the greatest common divisor of this set is 1.

The converse is not true, of course. A simple counterexample is given [below](#).

The Cyclic Classes

Suppose now that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is irreducible and is periodic with period d . There is no real loss in generality in assuming that the chain is irreducible, for if this were not the case, we could simply restrict our attention to one of the irreducible equivalence classes. Our exposition will be easier and cleaner if we recall the congruence equivalence relation modulo d on \mathbb{Z} , which in turn is based on the division partial order. For $n, m \in \mathbb{Z}$, $n \equiv_d m$ if and only if $d \mid (n - m)$, equivalently $n - m$ is an integer multiple of d , equivalently m and n have the same remainder after division by d . The basic fact that we will need is that \equiv_d is preserved under sums and differences. That is, if $m, n, p, q \in \mathbb{Z}$ and if $m \equiv_d n$ and $p \equiv_d q$, then $m + p \equiv_d n + q$ and $m - p \equiv_d n - q$.

Now, we fix a reference state $u \in S$, and for $k \in \{0, 1, \dots, d-1\}$, define

$$A_k = \{x \in S : P^{nd+k}(u, x) > 0 \text{ for some } n \in \mathbb{N}\} \quad (16.5.2)$$

That is, $x \in A_k$ if and only if there exists $m \in \mathbb{N}$ with $m \equiv_d k$ and $P^m(u, x) > 0$.

Suppose that $x, y \in S$.

1. If $x \in A_j$ and $y \in A_k$ for $j, k \in \{0, 1, \dots, d-1\}$ then $P^n(x, y) > 0$ for some $n \equiv_d k - j$
2. Conversely, if $P^n(x, y) > 0$ for some $n \in \mathbb{N}$, then there exists $j, k \in \{0, 1, \dots, d-1\}$ such that $x \in A_j, y \in A_k$ and $n \equiv_d k - j$.
3. The sets $(A_0, A_1, \dots, A_{d-1})$ partition S .

Proof

Suppose first that $x \in A_k$. By definition, u leads to x in m steps for some $m \equiv_d k$. Since the chain is irreducible, x leads back to u , say in p steps. It follows that u leads back to u in $m + p$ steps. But because u has period d , we must have $m + p \equiv_d 0$ and hence $p \equiv_d -k$.

Next suppose that $x \in A_j$ and $y \in A_k$. We know that u leads to x in m steps for some $m \equiv_d j$, and we now know that y leads to u in p steps for some $p \equiv_d -k$. Since the chain is irreducible, x leads to y , say in n steps. It follows that u leads back to u in $m + n + p$ steps. Again, since u has period d , $m + n + p \equiv_d 0$ and it follows that $n \equiv_d k - j$.

Next, note that since the chain is irreducible and since $\{0, 1, \dots, d-1\}$ is the set of all remainders modulo d , we must have $\bigcup_{i=0}^{d-1} A_i = S$. Suppose that $x, y \in S$ and $P^n(x, y) > 0$ for some $n \in \mathbb{N}$. Then $x \in A_j$ and $y \in A_k$ for some $j, k \in \{0, 1, \dots\}$. By the same argument as the last paragraph, we must have $n \equiv_d k - j$.

All that remains is to show that the sets are disjoint, and that amounts to the same old argument yet again. Suppose that $x \in A_j \cap A_k$ for some $j, k \in \{0, 1, \dots, d-1\}$. Then u leads to x in m steps for some $m \equiv_d j$ and x leads to u in n steps for some $n \equiv_d -k$. Hence u leads back to u in $m + n$ steps. Since the chain has period d we have $m + n \equiv_d 0$ and therefore $j - k \equiv_d 0$. Since $j, k \in \{0, 1, \dots, d-1\}$ it follows that $j = k$.

$(A_0, A_1, \dots, A_{d-1})$ are the equivalence classes for the d -step to and from relation \leftrightarrow_d that governs the d -step chain $(X_0, X_d, X_{2d}, \dots)$ that has transition matrix P^d .

The sets $(A_0, A_1, \dots, A_{d-1})$ are known as the *cyclic classes*. The basic structure of the chain is shown in the state diagram below:




Figure 16.5.1: The cyclic classes of a chain with period d

Examples and Special Cases

Finite Chains

Consider the Markov chain with state space $S = \{a, b, c\}$ and transition matrix P given below:

$$P = \begin{bmatrix} 0 & \frac{1}{3} & \frac{2}{3} \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \quad (16.5.3)$$

1. Sketch the state graph and show that the chain is irreducible.
2. Show that the chain is aperiodic.
3. Note that $P(x, x) = 0$ for all $x \in S$.

Answer

1. The state graph has edge set $E = \{(a, b), (a, c), (b, c), (b, a)\}$.
2. Note that $P^2(a, a) > 0$ and $P^3(a, a) > 0$. Hence $d(a) = 1$ since 2 and 3 are relatively prime. Since the chain is irreducible, it is aperiodic.

Consider the Markov chain with state space $S = \{1, 2, 3, 4, 5, 6, 7\}$ and transition matrix P given below:

$$P = \begin{bmatrix} 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 & \frac{2}{3} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{2}{3} \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & \frac{3}{4} & \frac{1}{4} \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{4} & \frac{3}{4} & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (16.5.4)$$

1. Sketch the state graph and show that the chain is irreducible.
2. Find the period d .
3. Find P^d .
4. Identify the cyclic classes.

Answer

1.

State graph



2. Period 3

$$3. P^3 = \begin{bmatrix} \frac{71}{192} & \frac{121}{192} & 0 & 0 & 0 & 0 & 0 \\ \frac{29}{72} & \frac{43}{72} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{7}{18} & \frac{1}{12} & \frac{19}{36} & 0 & 0 \\ 0 & 0 & \frac{19}{48} & \frac{3}{32} & \frac{49}{96} & 0 & 0 \\ 0 & 0 & \frac{13}{32} & \frac{7}{64} & \frac{31}{64} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{157}{299} & \frac{131}{288} \\ 0 & 0 & 0 & 0 & 0 & \frac{37}{64} & \frac{27}{64} \end{bmatrix}$$

4. Cyclic classes: $\{1, 2\}, \{3, 4, 5\}, \{6, 7\}$

Special Models

Review the definition of the basic Ehrenfest chain. Show that this chain has period 2, and find the cyclic classes.

Review the definition of the modified Ehrenfest chain. Show that this chain is aperiodic.

Review the definition of the simple random walk on \mathbb{Z}^k . Show that the chain is periodic with period 2, and find the cyclic classes.

This page titled [16.5: Periodicity of Discrete-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.6: Stationary and Limiting Distributions of Discrete-Time Chains

In this section, we study some of the deepest and most interesting parts of the theory of discrete-time Markov chains, involving two different but complementary ideas: stationary distributions and limiting distributions. The theory of renewal processes plays a critical role.

Basic Theory

As usual, our starting point is a (time homogeneous) discrete-time Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ with (countable) state space S and transition probability matrix P . In the background, of course, is a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ so that Ω is the sample space, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on (Ω, \mathcal{F}) . For $n \in \mathbb{N}$, let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$, the σ -algebra of events determined by the chain up to time n , so that $\mathcal{F} = \{\mathcal{F}_0, \mathcal{F}_1, \dots\}$ is the natural filtration associated with \mathbf{X} .

The Embedded Renewal Process

Let $y \in S$ and $n \in \mathbb{N}_+$. We will denote the number of visits to y during the first n positive time units by

$$N_{y,n} = \sum_{i=1}^n \mathbf{1}(X_i = y) \quad (16.6.1)$$

Note that $N_{y,n} \rightarrow N_y$ as $n \rightarrow \infty$, where

$$N_y = \sum_{i=1}^{\infty} \mathbf{1}(X_i = y) \quad (16.6.2)$$

is the total number of visits to y at positive times, one of the important random variables that we studied in the section on transience and recurrence. For $n \in \mathbb{N}_+$, we denote the time of the n th visit to y by

$$\tau_{y,n} = \min\{k \in \mathbb{N}_+ : N_{y,k} = n\} \quad (16.6.3)$$

where as usual, we define $\min(\emptyset) = \infty$. Note that $\tau_{y,1}$ is the time of the first visit to y , which we denoted simply by τ_y in the section on transience and recurrence. The times of the visits to y are stopping times for \mathbf{X} . That is, $\{\tau_{y,n} = k\} \in \mathcal{F}_k$ for $n \in \mathbb{N}_+$ and $k \in \mathbb{N}$. Recall also the definition of the hitting probability to state y starting in state x :

$$H(x, y) = \mathbb{P}(\tau_y < \infty \mid X_0 = x), \quad (x, y) \in S^2 \quad (16.6.4)$$

Suppose that $x, y \in S$, and that y is recurrent and $X_0 = x$.

1. If $x = y$, then the successive visits to y form a renewal process.
2. If $x \neq y$ but $x \rightarrow y$, then the successive visits to y form a delayed renewal process.

Proof

Let $\tau_{y,0} = 0$ for convenience.

1. Given $X_0 = y$, the sequence $(\tau_{y,1}, \tau_{y,2}, \dots)$ is the sequence of arrival times of a renewal process. Every time the chain reaches state y , the process starts over, independently of the past, by the Markov property. Thus the interarrival times $\tau_{y,n+1} - \tau_{y,n}$ for $n \in \mathbb{N}$ are conditionally independent, and are identically distributed, given $X_0 = y$.
2. If $x \neq y$ but $x \rightarrow y$, then given $X_0 = x$, the sequence $(\tau_{y,1}, \tau_{y,2}, \dots)$ is the sequence of arrival times of a delayed renewal process. By the same argument as in (a), the interarrival times $\tau_{y,n+1} - \tau_{y,n}$ for $n \in \mathbb{N}$ are conditionally independent, given $X_0 = x$, and all but $\tau_{y,1}$ have the same distribution.

As noted in the proof, $(\tau_{y,1}, \tau_{y,2}, \dots)$ is the sequence of arrival times and $(N_{y,1}, N_{y,2}, \dots)$ is the associated sequence of counting variables for the embedded renewal process associated with the recurrent state y . The corresponding renewal function, given $X_0 = x$, is the function $n \mapsto G_n(x, y)$ where

$$G_n(x, y) = \mathbb{E}(N_{y,n} \mid X_0 = x) = \sum_{k=1}^n P^k(x, y), \quad n \in \mathbb{N} \quad (16.6.5)$$

Thus $G_n(x, y)$ is the expected number of visits to y in the first n positive time units, starting in state x . Note that $G_n(x, y) \rightarrow G(x, y)$ as $n \rightarrow \infty$ where G is the potential matrix that we studied previously. This matrix gives the expected total number visits to state $y \in S$, at positive times, starting in state $x \in S$:

$$G(x, y) = \mathbb{E}(N_y \mid X_0 = x) = \sum_{k=1}^{\infty} P^k(x, y) \quad (16.6.6)$$

Limiting Behavior

The limit theorems of renewal theory can now be used to explore the limiting behavior of the Markov chain. Let $\mu(y) = \mathbb{E}(\tau_y \mid X_0 = y)$ denote the mean return time to state y , starting in y . In the following results, it may be the case that $\mu(y) = \infty$, in which case we interpret $1/\mu(y)$ as 0.

If $x, y \in S$ and y is recurrent then

$$\mathbb{P}\left(\frac{1}{n}N_{y,n} \rightarrow \frac{1}{\mu(y)} \text{ as } n \rightarrow \infty \mid X_0 = x\right) = H(x, y) \quad (16.6.7)$$

Proof

This result follows from the strong law of large numbers for renewal processes.

Note that $\frac{1}{n}N_{y,n} = \frac{1}{n} \sum_{k=1}^n \mathbf{1}(X_k = y)$ is the average number of visits to y in the first n positive time units.

If $x, y \in S$ and y is recurrent then

$$\frac{1}{n}G_n(x, y) = \frac{1}{n} \sum_{k=1}^n P^k(x, y) \rightarrow \frac{H(x, y)}{\mu(y)} \text{ as } n \rightarrow \infty \quad (16.6.8)$$

Proof

This result follows from the elementary renewal theorem for renewal processes.

Note that $\frac{1}{n}G_n(x, y) = \frac{1}{n} \sum_{k=1}^n P^k(x, y)$ is the expected average number of visits to y during the first n positive time units, starting at x .

If $x, y \in S$ and y is recurrent and aperiodic then

$$P^n(x, y) \rightarrow \frac{H(x, y)}{\mu(y)} \text{ as } n \rightarrow \infty \quad (16.6.9)$$

Proof

This result follows from the renewal theorem for renewal processes.

Note that $H(y, y) = 1$ by the very definition of a recurrent state. Thus, when $x = y$, the law of large numbers [above](#) gives convergence with probability 1, and the [first](#) and [second](#) renewal theory limits above are simply $1/\mu(y)$. By contrast, we already know the corresponding limiting behavior when y is transient.

If $x, y \in S$ and y is transient then

1. $\mathbb{P}\left(\frac{1}{n}N_{y,n} \rightarrow 0 \text{ as } n \rightarrow \infty \mid X_0 = x\right) = 1$
2. $\frac{1}{n}G_n(x, y) = \frac{1}{n} \sum_{k=1}^n P^k(x, y) \rightarrow 0 \text{ as } n \rightarrow \infty$
3. $P^n(x, y) \rightarrow 0 \text{ as } n \rightarrow \infty$

Proof

1. Note that $0 \leq \frac{1}{n}N_{y,n} \leq \frac{1}{n}N_y$. But if y is transient, $\mathbb{P}(N_y < \infty \mid X_0 = x) = 1$ and hence $\mathbb{P}\left(\frac{1}{n}N_y \rightarrow 0 \text{ as } n \rightarrow \infty \mid X_0 = x\right) = 1$ so the result follows from the squeeze theorem for limits.
2. Similarly, note that

$$0 \leq \frac{1}{n} \sum_{k=1}^n P^k(x, y) \leq \frac{1}{n} \sum_{k=1}^{\infty} P^k(x, y) \quad (16.6.10)$$

If y is transient, $G(x, y) = \sum_{k=1}^{\infty} P^k(x, y) < \infty$ and hence $\frac{1}{n} G(x, y) \rightarrow 0$ as $n \rightarrow \infty$. Again the result follows from the squeeze theorem for limits.

3. Once more, if y is transient, $G(x, y) = \sum_{k=1}^{\infty} P^k(x, y) < \infty$ and therefore $P^n(x, y) \rightarrow 0$ as $n \rightarrow \infty$.

On the other hand, if y is transient then $\mathbb{P}(\tau_y = \infty \mid X_0 = y) > 0$ by the very definition of a transience. Thus $\mu(y) = \infty$, and so the results in parts (b) and (c) agree with the corresponding results above for a recurrent state. Here is a summary.

For $x, y \in S$,

$$\frac{1}{n} G_n(x, y) = \frac{1}{n} \sum_{k=1}^n P^k(x, y) \rightarrow \frac{H(x, y)}{\mu(y)} \text{ as } n \rightarrow \infty \quad (16.6.11)$$

If y is transient or if y is recurrent and aperiodic,

$$P^n(x, y) \rightarrow \frac{H(x, y)}{\mu(y)} \text{ as } n \rightarrow \infty \quad (16.6.12)$$

Positive and Null Recurrence

Clearly there is a fundamental dichotomy in terms of the limiting behavior of the chain, depending on whether the mean return time to a given state is finite or infinite. Thus the following definition is natural.

Let $x \in S$.

1. State x is *positive recurrent* if $\mu(x) < \infty$.
2. If x is recurrent but $\mu(x) = \infty$ then state x is *null recurrent*.

Implicit in the definition is the following simple result:

If $x \in S$ is positive recurrent, then x is recurrent.

Proof

Recall that if $\mathbb{E}(\tau_x \mid X_0 = x) < \infty$ then $\mathbb{P}(\tau_x < \infty \mid X_0 = x) = 1$.

On the other hand, it is possible to have $\mathbb{P}(\tau_x < \infty \mid X_0 = x) = 1$, so that x is recurrent, and also $\mathbb{E}(\tau_x \mid X_0 = x) = \infty$, so that x is null recurrent. Simply put, a random variable can be finite with probability 1, but can have infinite expected value. A classic example is the Pareto distribution with shape parameter $a \in (0, 1)$.

Like recurrence/transience, and period, the null/positive recurrence property is a *class property*.

If x is positive recurrent and $x \rightarrow y$ then y is positive recurrent.

Proof

Suppose that x is positive recurrent and $x \rightarrow y$. Recall that y is recurrent and $y \rightarrow x$. Hence there exist $i, j \in \mathbb{N}_+$ such that $P^i(x, y) > 0$ and $P^j(y, x) > 0$. Thus for every $k \in \mathbb{N}_+$, $P^{i+j+k}(y, y) \geq P^j(y, x) P^k(x, x) P^i(x, y)$. Averaging over k from 1 to n gives

$$\frac{G_n(y, y)}{n} - \frac{G_{i+j}(y, y)}{n} \geq P^j(y, x) \frac{G_n(x, x)}{n} P^i(x, y) \quad (16.6.13)$$

Letting $n \rightarrow \infty$ and using renewal theory limit [above](#) gives

$$\frac{1}{\mu(y)} \geq P^j(y, x) \frac{1}{\mu(x)} P^i(x, y) > 0 \quad (16.6.14)$$

Therefore $\mu(y) < \infty$ and so y is also positive recurrent.

Thus, the terms *positive recurrent* and *null recurrent* can be applied to equivalence classes (under the *to and from* equivalence relation), as well as individual states. When the chain is irreducible, the terms can be applied to the chain as a whole.

Recall that a nonempty set of states A is *closed* if $x \in A$ and $x \rightarrow y$ implies $y \in A$. Here are some simple results for a finite, closed set of states.

If $A \subseteq S$ is finite and closed, then A contains a positive recurrent state.

Proof

Fix a state $x \in A$ and note that $P^k(x, A) = \sum_{y \in A} P^k(x, y) = 1$ for every $k \in \mathbb{N}_+$ since A is closed. Averaging over k from 1 to n gives

$$\sum_{y \in A} \frac{G_n(x, y)}{n} = 1 \quad (16.6.15)$$

for every $n \in \mathbb{N}_+$. Note that the change in the order of summation is justified since both sums are finite. Assume now that all states in A are transient or null recurrent. Letting $n \rightarrow \infty$ in the displayed equation gives the contradiction $0 = 1$. Again, the interchange of sum and limit is justified by the fact that A is finite.

If $A \subseteq S$ is finite and closed, then A contains no null recurrent states.

Proof

Let $x \in A$. Note that $[x] \subseteq A$ since A is closed. Suppose that x is recurrent. Note that $[x]$ is also closed and finite and hence must have a positive recurrent state by the [previous result](#). Hence the equivalence class $[x]$ is positive recurrent and thus so is x .

If $A \subseteq S$ is finite and irreducible, then A is a positive recurrent equivalence class.

Proof

We already know that A is a recurrent equivalence class, from our study of transience and recurrence. From the previous theorem, A is positive recurrent.

In particular, a Markov chain with a finite state space cannot have null recurrent states; every state must be transient or positive recurrent.

Limiting Behavior, Revisited

Returning to the limiting behavior, suppose that the chain \mathbf{X} is irreducible, so that either all states are transient, all states are null recurrent, or all states are positive recurrent. From the basic limit theorem [above](#), if the chain is transient or if the chain is recurrent and aperiodic, then

$$P^n(x, y) \rightarrow \frac{1}{\mu(y)} \text{ as } n \rightarrow \infty \text{ for every } x \in S \quad (16.6.16)$$

Note in particular that the limit is independent of the initial state x . Of course in the transient case and in the null recurrent and aperiodic case, the limit is 0. Only in the positive recurrent, aperiodic case is the limit positive, which motivates our next definition.

A Markov chain \mathbf{X} that is irreducible, positive recurrent, and aperiodic, is said to be *ergodic*.

In the ergodic case, as we will see, X_n has a limiting distribution as $n \rightarrow \infty$ that is independent of the initial distribution.

The behavior when the chain is periodic with period $d \in \{2, 3, \dots\}$ is a bit more complicated, but we can understand this behavior by considering the d -step chain $\mathbf{X}_d = (X_0, X_d, X_{2d}, \dots)$ that has transition matrix P^d . Essentially, this allows us to trade *periodicity* (one form of complexity) for *reducibility* (another form of complexity). Specifically, recall that the d -step chain is aperiodic but has d equivalence classes $(A_0, A_1, \dots, A_{d-1})$; and these are the *cyclic classes* of original chain \mathbf{X} .




Figure 16.6.1: The cyclic classes of a chain with period d

The mean return time to state x for the d -step chain \mathbf{X}_d is $\mu_d(x) = \mu(x)/d$.

Proof

Note that every single step for the d -step chain corresponds to d steps for the original chain.

Let $i, j, k \in \{0, 1, \dots, d-1\}$,

1. $P^{nd+k}(x, y) \rightarrow d/\mu(y)$ as $n \rightarrow \infty$ if $x \in A_i$ and $y \in A_j$ and $j = (i+k) \bmod d$.
2. $P^{nd+k}(x, y) \rightarrow 0$ as $n \rightarrow \infty$ in all other cases.

Proof

These results follow from the previous theorem and the cyclic behavior of the chain.

If $y \in S$ is null recurrent or transient then regardless of the period of y , $P^n(x, y) \rightarrow 0$ as $n \rightarrow \infty$ for every $x \in S$.

Invariant Distributions

Our next goal is to see how the limiting behavior is related to invariant distributions. Suppose that f is a probability density function on the state space S . Recall that f is *invariant* for P (and for the chain \mathbf{X}) if $fP = f$. It follows immediately that $fP^n = f$ for every $n \in \mathbb{N}$. Thus, if X_0 has probability density function f then so does X_n for each $n \in \mathbb{N}$, and hence \mathbf{X} is a sequence of identically distributed random variables. A bit more generally, suppose that $g: S \rightarrow [0, \infty)$ is invariant for P , and let $C = \sum_{x \in S} g(x)$. If $0 < C < \infty$ then f defined by $f(x) = g(x)/C$ for $x \in S$ is an invariant probability density function.

Suppose that $g: S \rightarrow [0, \infty)$ is invariant for P and satisfies $\sum_{x \in S} g(x) < \infty$. Then

$$g(y) = \frac{1}{\mu(y)} \sum_{x \in S} g(x)H(x, y), \quad y \in S \quad (16.6.17)$$

Proof

Recall again that $gP^k = g$ for every $k \in \mathbb{N}$ since g is invariant for P . Averaging over k from 1 to n gives $gG_n/n = g$ for each $n \in \mathbb{N}_+$. Explicitly,

$$\sum_{x \in S} g(x) \frac{G_n(x, y)}{n} = g(y), \quad y \in S \quad (16.6.18)$$

Letting $n \rightarrow \infty$ and using the limit theorem [above](#) gives the result. The dominated convergence theorem justifies interchanging the limit with the sum, since the terms are positive, $\frac{1}{n}G_n(x, y) \leq 1$, and $\sum_{x \in S} g(x) < \infty$.

Note that if y is transient or null recurrent, then $g(y) = 0$. Thus, an invariant function with finite sum, and in particular an invariant probability density function must be concentrated on the positive recurrent states.

Suppose now that the chain \mathbf{X} is irreducible. If \mathbf{X} is transient or null recurrent, then from the previous result, the only nonnegative functions that are invariant for P are functions that satisfy $\sum_{x \in S} g(x) = \infty$ and the function that is identically 0: $g = \mathbf{0}$. In particular, the chain does not have an invariant distribution. On the other hand, if the chain is positive recurrent, then $H(x, y) = 1$

for all $x, y \in S$. Thus, from the previous result, the only possible invariant probability density function is the function f given by $f(x) = 1/\mu(x)$ for $x \in S$. Any other nonnegative function g that is invariant for P and has finite sum, is a multiple of f (and indeed the multiple is sum of the values). Our next goal is to show that f really is an invariant probability density function.

If \mathbf{X} is an irreducible, positive recurrent chain then the function f given by $f(x) = 1/\mu(x)$ for $x \in S$ is an invariant probability density function for \mathbf{X} .

Proof

Let $f(x) = 1/\mu(x)$ for $x \in S$, and let A be a finite subset of S . Then $\sum_{y \in A} \frac{1}{n} G_n(x, y) \leq 1$ for every $x \in S$. Letting $n \rightarrow \infty$ using the basic limit [above](#) gives $\sum_{y \in A} f(y) \leq 1$. The interchange of limit and sum is justified since A is finite. Since this is true for every finite $A \subseteq S$, it follows that $C \leq 1$ where $C = \sum_{y \in S} f(y)$. Note also that $C > 0$ since the chain is positive recurrent. Next note that

$$\sum_{y \in A} \frac{1}{n} G_n(x, y) P(y, z) \leq \frac{1}{n} G_{n+1}(x, z) \quad (16.6.19)$$

for every $x, z \in S$. Letting $n \rightarrow \infty$ gives $\sum_{y \in A} f(y) P(y, z) \leq f(z)$ for every $z \in S$. It then follows that $\sum_{y \in S} f(y) P(y, z) \leq f(z)$ for every $z \in S$. Suppose that strict inequality holds for some for some $z \in S$. Then

$$\sum_{z \in S} \sum_{y \in S} f(y) P(y, z) < \sum_{z \in S} f(z) \quad (16.6.20)$$

Interchanging the order of summation on the left in the displayed inequality yields the contradiction $C < C$. Thus f is invariant for P . Hence f/C is an invariant probability density function. By the uniqueness result noted earlier, it follows that $f/C = f$ so that in fact $C = 1$.

In summary, an irreducible, positive recurrent Markov chain \mathbf{X} has a unique invariant probability density function f given by $f(x) = 1/\mu(x)$ for $x \in S$. We also now have a *test* for positive recurrence. An irreducible Markov chain \mathbf{X} is positive recurrent if and only if there exists a positive function g on S that is invariant for P and satisfies $\sum_{x \in S} g(x) < \infty$ (and then, of course, normalizing g would give f).

Consider now a general Markov chain \mathbf{X} on S . If \mathbf{X} has no positive recurrent states, then as noted earlier, there are no invariant distributions. Thus, suppose that \mathbf{X} has a collection of positive recurrent equivalence classes $(A_i : i \in I)$ where I is a nonempty, countable index set. The chain restricted to A_i is irreducible and positive recurrent for each $i \in I$, and hence has a unique invariant probability density function f_i on A_i given by

$$f_i(x) = \frac{1}{\mu(x)}, \quad x \in A_i \quad (16.6.21)$$

We extend f_i to S by defining $f_i(x) = 0$ for $x \notin A_i$, so that f_i is a probability density function on S . All invariant probability density functions for \mathbf{X} are mixtures of these functions:

f is an invariant probability density function for \mathbf{X} if and only if f has the form

$$f(x) = \sum_{i \in I} p_i f_i(x), \quad x \in S \quad (16.6.22)$$

where $(p_i : i \in I)$ is a probability density function on the index set I . That is, $f(x) = p_i f_i(x)$ for $i \in I$ and $x \in A_i$, and $f(x) = 0$ otherwise.

Proof

Let $A = \bigcup_{i \in I} A_i$, the set of positive recurrent states. Suppose that f has the form given in the theorem. Since $f(x) = 0$ for $x \notin A$ we have

$$(fP)(y) = \sum_{x \in S} f(x) P(x, y) = \sum_{i \in I} \sum_{x \in A_i} p_i f_i(x) P(x, y) \quad (16.6.23)$$

Suppose that $y \in A_j$ for some $j \in I$. Since $P(x, y) = 0$ if $x \in A_i$ and $i \neq j$, the last sum becomes

$$(fP)(y) = p_j \sum_{x \in A_j} f_j(x) P(x, y) = p_j f_j(y) = f(y) \quad (16.6.24)$$

because f_j is invariant for the P restricted to A_j . If $y \notin A$ then $P(x, y) = 0$ for $x \in A$ so the sum above becomes $(fP)(y) = 0 = f(y)$. Hence f is invariant. Moreover,

$$\sum_{x \in S} f(x) = \sum_{i \in I} \sum_{x \in A_i} f(x) = \sum_{i \in I} p_i \sum_{x \in A_i} f_i(x) = \sum_{i \in I} p_i = 1 \quad (16.6.25)$$

so f is a PDF on S . Conversely, suppose that f is an invariant PDF for \mathbf{X} . We know that f is concentrated on the positive recurrent states, so $f(x) = 0$ for $x \notin A$. For $i \in I$ and $y \in A_i$

$$\sum_{x \in A_i} f(x) P(x, y) = \sum_{x \in S} f(x) P(x, y) = f(y) \quad (16.6.26)$$

since f is invariant for P and since, as noted before, $f(x) P(x, y) = 0$ if $x \notin A_i$. It follows that f restricted to A_i is invariant for the chain restricted to A_i for each $i \in I$. Let $p_i = \sum_{x \in A_i} f(x)$, the normalizing constant for f restricted to A_i . By uniqueness, the restriction of f/p_i to A_i must be f_i , so f has the form given in the theorem.

Invariant Measures

Suppose that \mathbf{X} is irreducible. In this section we are interested in general functions $g: S \rightarrow [0, \infty)$ that are invariant for \mathbf{X} , so that $gP = g$. A function $g: S \rightarrow [0, \infty)$ defines a positive measure ν on S by the simple rule

$$\nu(A) = \sum_{x \in A} g(x), \quad A \subseteq S \quad (16.6.27)$$

so in this sense, we are interested in invariant positive measures for \mathbf{X} that may not be probability measures. Technically, g is the density function of ν with respect to counting measure $\#$ on S .

From our work above, We know the situation if \mathbf{X} is positive recurrent. In this case, there exists a unique invariant probability density function f that is positive on S , and every other nonnegative invariant function g is a nonnegative multiples of f . In particular, either $g = 0$, the zero function on S , or g is positive on S and satisfies $\sum_{x \in S} g(x) < \infty$.

We can generalize to chains that are simply recurrent, either null or positive. We will show that there exists a positive invariant function that is unique, up to multiplication by positive constants. To set up the notation, recall that $\tau_x = \min\{k \in \mathbb{N}_+ : X_k = x\}$ is the first positive time that the chain is in state $x \in S$. In particular, if the chain starts in x then τ_x is the time of the first return to x . For $x \in S$ we define the function γ_x by

$$\gamma_x(y) = \mathbb{E} \left(\sum_{n=0}^{\tau_x-1} \mathbf{1}(X_n = y) \mid X_0 = x \right), \quad y \in S \quad (16.6.28)$$

so that $\gamma_x(y)$ is the expected number of visits to y before the first return to x , starting in x . Here is the existence result.

Suppose that \mathbf{X} is recurrent. For $x \in S$,

1. $\gamma_x(x) = 1$
2. γ_x is invariant for \mathbf{X}
3. $\gamma_x(y) \in (0, \infty)$ for $y \in S$.

Proof

1. By definition, given $X_0 = x$, we have $X_0 = x$ but $X_n \neq x$ for $n \in \{1, \dots, \tau_x - 1\}$. Hence $\gamma_x(x) = 1$.
2. Since the chain is recurrent, with probability 1 we have $\tau_x < \infty$ and $X_{\tau_x} = x$. Hence for $y \in S$,

$$\gamma_x(y) = \mathbb{E} \left(\sum_{n=0}^{\tau_x-1} \mathbf{1}(X_n = y) \mid X_0 = x \right) = \mathbb{E} \left(\sum_{n=1}^{\tau_x} \mathbf{1}(X_n = y) \mid X_0 = x \right) \quad (16.6.29)$$

(Note that if $x = y$ then with probability 1, the $n = 0$ term in the first sum and the $n = \tau_x$ term in the second sum are 1 and the remaining terms are 0. If $x \neq y$, the $n = 0$ term in the first sum and the $n = \tau_x$ term in the second sum are 0 with probability 1, so again the two sums are the same.) Hence

$$\gamma_x(y) = \mathbb{E} \left(\sum_{n=1}^{\infty} \mathbf{1}(X_n = y, \tau_x \geq n) \mid X_0 = x \right) = \sum_{n=1}^{\infty} \mathbb{P}(X_n = y, \tau_x \geq n \mid X_0 = x) \quad (16.6.30)$$

Next we partition on the values of X_{n-1} in the sum to obtain

$$\begin{aligned} \gamma_x(y) &= \sum_{n=1}^{\infty} \sum_{z \in S} \mathbb{P}(X_n = y, X_{n-1} = z, \tau_x \geq n \mid X_0 = x) \\ &= \sum_{n=1}^{\infty} \sum_{z \in S} \mathbb{P}(X_n = y \mid X_{n-1} = z, \tau_x \geq n, X_0 = x) \mathbb{P}(X_{n-1} = z, \tau_x \geq n \mid X_0 = x) \end{aligned}$$

But $\{X_0 = x, \tau_x \geq n\} \in \mathcal{F}_{n-1}$ (that is, the events depend only on (X_0, \dots, X_{n-1})). Hence by the Markov property, the first factor in the last displayed equation is simply $\mathbb{P}(X_n = y \mid X_{n-1} = z) = P(z, y)$. Substituting and re-indexing the sum gives

$$\begin{aligned} \gamma_x(y) &= \sum_{n=1}^{\infty} \sum_{z \in S} P(z, y) \mathbb{P}(X_{n-1} = z, \tau_x \geq n \mid X_0 = x) = \sum_{z \in S} P(z, y) \mathbb{E} \left(\sum_{n=1}^{\tau_x} \mathbf{1}(X_{n-1} = z) \mid X_0 = x \right) \\ &= \sum_{z \in S} P(z, y) \mathbb{E} \left(\sum_{m=0}^{\tau_x-1} \mathbf{1}(X_m = z) \mid X_0 = x \right) = \sum_{z \in S} P(z, y) \gamma_x(z) = \gamma_x P(y) \end{aligned}$$

3. By the invariance in part (b), $\gamma_x = \gamma_x P^n$ for every $n \in \mathbb{N}$. Let $y \in S$. Since the chain is irreducible, there exists $j \in \mathbb{N}$ such that $P^j(x, y) > 0$. Hence

$$\gamma_x(y) = \gamma_x P^j(y) \geq \gamma_x(x) P^j(x, y) = P^j(x, y) > 0 \quad (16.6.31)$$

Similarly, there exists $k \in \mathbb{N}$ such that $P^k(y, x) > 0$. Hence

$$1 = \gamma_x(x) = \gamma_x P^k(x) \geq \gamma_x(y) P^k(y, x) \quad (16.6.32)$$

and therefore $\gamma_x(y) \leq 1/P^k(y, x) < \infty$.

Next is the uniqueness result.

Suppose again that \mathbf{X} is recurrent and that $g: S \rightarrow [0, \infty)$ is invariant for \mathbf{X} . For fixed $x \in S$,

$$g(y) = g(x) \gamma_x(y), \quad y \in S \quad (16.6.33)$$

Proof

Let $S_x = S - \{x\}$ and let $y \in S$. Since g is invariant,

$$g(y) = gP(y) = \sum_{z \in S} g(z) P(z, y) = \sum_{z \in S_x} g(z) P(z, y) + g(x) P(x, y) \quad (16.6.34)$$

Note that the last term is $g(x) \mathbb{P}(X_1 = y, \tau_x \geq 1 \mid X_0 = x)$. Repeating the argument for $g(z)$ in the sum above gives

$$g(y) = \sum_{z \in S_x} \sum_{w \in S_x} g(w) P(w, z) P(z, y) + g(x) \sum_{z \in S_x} P(x, z) P(z, y) + g(x) P(x, y) \quad (16.6.35)$$

The last two terms are

$$g(x) [\mathbb{P}(X_2 = y, \tau_x \geq 2 \mid X_0 = x) + \mathbb{P}(X_1 = y, \tau_x \geq 1 \mid X_0 = x)] \quad (16.6.36)$$

Continuing in this way shows that for each $n \in \mathbb{N}_+$,

$$g(y) \geq g(x) \sum_{k=1}^n \mathbb{P}(X_k = y, \tau_x \geq k \mid X_0 = x) \quad (16.6.37)$$

Letting $n \rightarrow \infty$ then shows that $g(y) \geq g(x) \gamma_x(y)$. Next, note that the function $h = g - g(x) \gamma_x$ is invariant, since it is a difference of two invariant functions, and as just shown, is nonnegative. Also, $h(x) = g(x) - g(x) \gamma_x(x) = 0$. Let $y \in S$. Since the chain is irreducible, there exists $j \in \mathbb{N}$ such that $P^j(y, x) > 0$. Hence

$$0 = h(x) = hP^j(x) \geq h(y)P^j(y, x) \geq 0 \quad (16.6.38)$$

Since $P^j(y, x) > 0$ it follows that $h(y) = 0$.

Thus, suppose that \mathbf{X} is null recurrent. Then there exists an invariant function g that is positive on S and satisfies $\sum_{x \in S} g(x) = \infty$. Every other nonnegative invariant function is a nonnegative multiple of g . In particular, either $g = \mathbf{0}$, the zero function on S , or g is positive on S and satisfies $\sum_{x \in S} g(x) = \infty$. The section on reliability chains gives an example of the invariant function for a null recurrent chain.

The situation is complicated when \mathbf{X} is transient. In this case, there may or may not exist nonnegative invariant functions that are not identically 0. When they do exist, they may not be unique (up to multiplication by nonnegative constants). But we still know that there are no invariant probability density functions, so if g is a nonnegative function that is invariant for \mathbf{X} then either $g = \mathbf{0}$ or $\sum_{x \in S} g(x) = \infty$. The section on random walks on graphs provides lots of examples of transient chains with nontrivial invariant functions. In particular, the non-symmetric random walk on \mathbb{Z} has a two-dimensional space of invariant functions.

Examples and Applications

Finite Chains

Consider again the general two-state chain on $S = \{0, 1\}$ with transition probability matrix given below, where $p \in (0, 1)$ and $q \in (0, 1)$ are parameters.

$$P = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix} \quad (16.6.39)$$

1. Find the invariant distribution.
2. Find the mean return time to each state.
3. Find $\lim_{n \rightarrow \infty} P^n$ without having to go to the trouble of diagonalizing P , as we did in the introduction to discrete-time chains.

Answer

1. $f = \left(\frac{q}{p+q}, \frac{p}{p+q} \right)$
2. $\mu = \left(\frac{p+q}{q}, \frac{p+q}{p} \right)$
3. $P^n \rightarrow \frac{1}{p+q} \begin{bmatrix} q & p \\ q & p \end{bmatrix}$ as $n \rightarrow \infty$.

Consider a Markov chain with state space $S = \{a, b, c, d\}$ and transition matrix P given below:

$$P = \begin{bmatrix} \frac{1}{3} & \frac{2}{3} & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{bmatrix} \quad (16.6.40)$$

1. Draw the state diagram.
2. Determine the equivalent classes and classify each as transient or positive recurrent.
3. Find all invariant probability density functions.
4. Find the mean return time to each state.
5. Find $\lim_{n \rightarrow \infty} P^n$.

Answer

- 1.
2. $\{a, b\}$ recurrent; $\{c\}$ recurrent; $\{d\}$ transient.
3. $f = \left(\frac{3}{5}p, \frac{2}{5}p, 1-p, 0 \right), 0 \leq p \leq 1$
4. $\mu = \left(\frac{5}{3}, \frac{5}{2}, 1, \infty \right)$

State graph


$$5. P^n \rightarrow \begin{bmatrix} \frac{3}{5} & \frac{2}{5} & 0 & 0 \\ \frac{3}{5} & \frac{2}{5} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ \frac{2}{5} & \frac{4}{15} & \frac{1}{3} & 0 \end{bmatrix} \text{ as } n \rightarrow \infty$$

Consider a Markov chain with state space $S = \{1, 2, 3, 4, 5, 6\}$ and transition matrix P given below:

$$P = \begin{bmatrix} 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ \frac{1}{4} & 0 & \frac{1}{2} & 0 & \frac{1}{4} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 & \frac{2}{3} & 0 \\ 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} \end{bmatrix} \quad (16.6.41)$$

1. Sketch the state graph.
2. Find the equivalence classes and classify each as transient or positive recurrent.
3. Find all invariant probability density functions.
4. Find the mean return time to each state.
5. Find $\lim_{n \rightarrow \infty} P^n$.

Answer

- 1.
2. $\{1, 3, 5\}$ recurrent; $\{2, 6\}$ transient; $\{4\}$ recurrent.
3. $f = \left(\frac{2}{19}p, 0, \frac{8}{19}p, 1-p, \frac{9}{19}p, 0 \right), \quad 0 \leq p \leq 1$
4. $\mu = \left(\frac{19}{2}, \infty, \frac{19}{8}, 1, \frac{19}{8}, \infty \right)$

State graph


$$5. P^n \rightarrow \begin{bmatrix} \frac{2}{19} & 0 & \frac{8}{19} & 0 & \frac{9}{19} & 0 \\ \frac{1}{19} & 0 & \frac{4}{19} & \frac{1}{2} & \frac{9}{38} & 0 \\ \frac{2}{19} & 0 & \frac{8}{19} & 0 & \frac{9}{19} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ \frac{2}{19} & 0 & \frac{8}{19} & 0 & \frac{9}{19} & 0 \\ \frac{1}{19} & 0 & \frac{4}{19} & \frac{1}{2} & \frac{9}{38} & 0 \end{bmatrix} \text{ as } n \rightarrow \infty.$$

Consider a Markov chain with state space $S = \{1, 2, 3, 4, 5, 6\}$ and transition matrix P given below:

$$P = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\ \frac{1}{4} & \frac{3}{4} & 0 & 0 & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{4} & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix} \quad (16.6.42)$$

1. Sketch the state graph.
2. Find the equivalence classes and classify each as transient or positive recurrent.
3. Find all invariant probability density functions.
4. Find the mean return time to each state.
5. Find $\lim_{n \rightarrow \infty} P^n$.

Answer

1.

State graph



2. $\{1, 2\}$ recurrent; $\{3, 4\}$ transient; $\{5, 6\}$ recurrent.

3. $f = \left(\frac{1}{3}p, \frac{2}{3}p, 0, 0, \frac{1}{2}(1-p), \frac{1}{2}(1-p)\right), \quad 0 \leq p \leq 1$

4. $\mu = \left(3, \frac{3}{2}, \infty, \infty, 2, 2\right)$

$$5. P^n \rightarrow \begin{bmatrix} \frac{1}{3} & \frac{2}{3} & 0 & 0 & 0 & 0 \\ \frac{1}{3} & \frac{2}{3} & 0 & 0 & 0 & 0 \\ \frac{4}{15} & \frac{8}{15} & 0 & 0 & 0 & 0 \\ \frac{1}{5} & \frac{2}{5} & 0 & 0 & \frac{1}{5} & \frac{1}{5} \\ 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix} \text{ as } n \rightarrow \infty$$

Consider the Markov chain with state space $S = \{1, 2, 3, 4, 5, 6, 7\}$ and transition matrix P given below:

$$P = \begin{bmatrix} 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 & \frac{2}{3} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{2}{3} \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & \frac{3}{4} & \frac{1}{4} \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{4} & \frac{3}{4} & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (16.6.43)$$

1. Sketch the state digraph, and show that the chain is irreducible with period 3.

2. Identify the cyclic classes.

3. Find the invariant probability density function.

4. Find the mean return time to each state.

5. Find $\lim_{n \rightarrow \infty} P^{3n}$.

6. Find $\lim_{n \rightarrow \infty} P^{3n+1}$.

7. Find $\lim_{n \rightarrow \infty} P^{3n+2}$.

Answer

1.

State graph



2. Cyclic classes: $\{1, 2\}, \{3, 4, 5\}, \{6, 7\}$

3. $f = \frac{1}{1785}(232, 363, 237, 58, 300, 333, 262)$

4. $\mu = 1785 \left(\frac{1}{232}, \frac{1}{363}, \frac{1}{237}, \frac{1}{58}, \frac{1}{300}, \frac{1}{333}, \frac{1}{262} \right)$

$$5. P^{3n} \rightarrow \frac{1}{585} \begin{bmatrix} 232 & 363 & 0 & 0 & 0 & 0 & 0 \\ 232 & 363 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 237 & 58 & 300 & 0 & 0 \\ 0 & 0 & 237 & 58 & 300 & 0 & 0 \\ 0 & 0 & 237 & 58 & 300 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 333 & 262 \\ 0 & 0 & 0 & 0 & 0 & 333 & 262 \end{bmatrix} \text{ as } n \rightarrow \infty$$

$$\begin{aligned}
 6. P^{3n+1} &\rightarrow \frac{1}{585} \begin{bmatrix} 0 & 0 & 237 & 58 & 300 & 0 & 0 \\ 0 & 0 & 237 & 58 & 300 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 333 & 262 \\ 0 & 0 & 0 & 0 & 0 & 333 & 262 \\ 0 & 0 & 0 & 0 & 0 & 333 & 262 \\ 232 & 363 & 0 & 0 & 0 & 0 & 0 \\ 232 & 363 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \text{ as } n \rightarrow \infty \\
 7. P^{3n+2} &\rightarrow \frac{1}{585} \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 333 & 262 \\ 0 & 0 & 0 & 0 & 0 & 333 & 262 \\ 232 & 363 & 0 & 0 & 0 & 0 & 0 \\ 232 & 363 & 0 & 0 & 0 & 0 & 0 \\ 232 & 363 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 237 & 58 & 300 & 0 & 0 \\ 0 & 0 & 237 & 58 & 300 & 0 & 0 \end{bmatrix} \text{ as } n \rightarrow \infty
 \end{aligned}$$

Special Models

Read the discussion of invariant distributions and limiting distributions in the Ehrenfest chains.

Read the discussion of invariant distributions and limiting distributions in the Bernoulli-Laplace chain.

Read the discussion of positive recurrence and invariant distributions for the reliability chains.

Read the discussion of positive recurrence and limiting distributions for the birth-death chain.

Read the discussion of positive recurrence and for the queuing chains.

Read the discussion of positive recurrence and limiting distributions for the random walks on graphs.

This page titled [16.6: Stationary and Limiting Distributions of Discrete-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.7: Time Reversal in Discrete-Time Chains

The Markov property, stated in the form that the past and future are independent given the present, essentially treats the past and future symmetrically. However, there is a lack of symmetry in the fact that in the usual formulation, we have an *initial* time 0, but not a *terminal* time. If we introduce a terminal time, then we can run the process backwards in time. In this section, we are interested in the following questions:

- Is the new process still Markov?
- If so, how does the new transition probability matrix relate to the original one?
- Under what conditions are the forward and backward processes stochastically the same?

Consideration of these questions leads to *reversed chains*, an important and interesting part of the theory of Markov chains.

Basic Theory

Reversed Chains

Our starting point is a (homogeneous) discrete-time Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ with (countable) state space S and transition probability matrix P . Let m be a positive integer, which we will think of as the *terminal time* or *finite time horizon*. We won't bother to indicate the dependence on m notationally, since ultimately the terminal time will not matter. Define $\hat{X}_n = X_{m-n}$ for $n \in \{0, 1, \dots, m\}$. Thus, the process *forward in time* is $\mathbf{X} = (X_0, X_1, \dots, X_m)$ while the process *backwards in time* is

$$\hat{\mathbf{X}} = (\hat{X}_0, \hat{X}_1, \dots, \hat{X}_m) = (X_m, X_{m-1}, \dots, X_0) \quad (16.7.1)$$

For $n \in \{0, 1, \dots, m\}$, let

$$\hat{\mathcal{F}}_n = \sigma\{\hat{X}_0, \hat{X}_1, \dots, \hat{X}_n\} = \sigma\{X_{m-n}, X_{m-n+1}, \dots, X_m\} \quad (16.7.2)$$

denote the σ algebra of the events of the process $\hat{\mathbf{X}}$ up to time n . So of course, an event for $\hat{\mathbf{X}}$ up to time n is an event for \mathbf{X} from time $m - n$ forward. Our first result is that the reversed process is still a Markov chain, but not time homogeneous in general.

The process $\hat{\mathbf{X}} = (\hat{X}_0, \hat{X}_1, \dots, \hat{X}_m)$ is a Markov chain, but is not time homogenous in general. The one-step transition matrix at time $n \in \{0, 1, \dots, m-1\}$ is given by

$$\mathbb{P}(\hat{X}_{n+1} = y \mid \hat{X}_n = x) = \frac{\mathbb{P}(X_{m-n-1} = y)}{\mathbb{P}(X_{m-n} = x)} P(y, x), \quad (x, y) \in S^2 \quad (16.7.3)$$

Proof

Let $A \in \hat{\mathcal{F}}_n$ and $x, y \in S$. Then

$$\begin{aligned} \mathbb{P}(\hat{X}_{n+1} = y \mid \hat{X}_n = x, A) &= \frac{\mathbb{P}(\hat{X}_{n+1} = y, \hat{X}_n = x, A)}{\mathbb{P}(\hat{X}_n = x, A)} = \frac{\mathbb{P}(X_{m-n-1} = y, X_{m-n} = x, A)}{\mathbb{P}(X_{m-n} = x, A)} \\ &= \frac{\mathbb{P}(A \mid X_{m-n-1} = y, X_{m-n} = x) \mathbb{P}(X_{m-n} = x \mid X_{m-n-1} = y) \mathbb{P}(X_{m-n-1} = y)}{\mathbb{P}(A \mid X_{m-n} = x) \mathbb{P}(X_{m-n} = x)} \end{aligned}$$

But $A \in \sigma\{X_{m-n}, \dots, X_m\}$ and so by the Markov property for \mathbf{X} ,

$$\mathbb{P}(A \mid X_{m-n-1} = y, X_{m-n} = x) = \mathbb{P}(A \mid X_{m-n} = x) \quad (16.7.4)$$

By the time homogeneity of \mathbf{X} , $\mathbb{P}(X_{m-n} = x \mid X_{m-n-1} = y) = P(y, x)$. Substituting and simplifying gives

$$\mathbb{P}(\hat{X}_{n+1} = y \mid \hat{X}_n = x, A) = \frac{\mathbb{P}(X_{m-n-1} = y)}{\mathbb{P}(X_{m-n} = x)} P(y, x) \quad (16.7.5)$$

However, the backwards chain will be time homogeneous if X_0 has an invariant distribution.

Suppose that \mathbf{X} is irreducible and positive recurrent, with (unique) invariant probability density function f . If X_0 has the invariant probability distribution, then $\hat{\mathbf{X}}$ is a time-homogeneous Markov chain with transition matrix \hat{P} given by

$$\hat{P}(x, y) = \frac{f(y)}{f(x)} P(y, x), \quad (x, y) \in S^2 \quad (16.7.6)$$

Proof

This follows from the result [above](#). Recall that if X_0 has PDF f , then X_k has PDF f for each $k \in \mathbb{N}$.

Recall that a discrete-time Markov chain is *ergodic* if it is irreducible, positive recurrent, and aperiodic. For an ergodic chain, the previous result holds in the limit of the terminal time.

Suppose that \mathbf{X} is ergodic, with (unique) invariant probability density function f . Regardless of the distribution of X_0 ,

$$\mathbb{P}(\hat{X}_{n+1} = y \mid \hat{X}_n = x) \rightarrow \frac{f(y)}{f(x)} P(y, x) \text{ as } n \rightarrow \infty \quad (16.7.7)$$

Proof

This follows from the conditional probability [above](#) and our study of the limiting behavior of Markov chains. Since \mathbf{X} is ergodic, $\mathbb{P}(X_k = x) \rightarrow f(x)$ as $k \rightarrow \infty$ for every $x \in S$.

These three results are motivation for the definition that follows. We can generalize by defining the reversal of an irreducible Markov chain, as long as there is a positive, invariant function. Recall that a positive invariant function defines a positive measure on S , but of course not in general a probability distribution.

Suppose that \mathbf{X} is an irreducible Markov chain with transition matrix P , and that $g: S \rightarrow (0, \infty)$ is invariant for \mathbf{X} . The reversal of \mathbf{X} with respect to g is the Markov chain $\hat{\mathbf{X}} = (\hat{X}_0, \hat{X}_1, \dots)$ with transition probability matrix \hat{P} defined by

$$\hat{P}(x, y) = \frac{g(y)}{g(x)} P(y, x), \quad (x, y) \in S^2 \quad (16.7.8)$$

Proof

We need to show that \hat{P} is a valid transition probability matrix, so that the definition makes sense. Since g is invariant for \mathbf{X} ,

$$\sum_{y \in S} \hat{P}(x, y) = \frac{1}{g(x)} \sum_{y \in S} g(y) P(y, x) = \frac{g(x)}{g(x)} = 1, \quad x \in S \quad (16.7.9)$$

Recall that if g is a positive invariant function for \mathbf{X} then so is cg for every positive constant c . Note that g and cg generate the same reversed chain. So let's consider the cases:

Suppose that \mathbf{X} is an irreducible Markov chain on S .

1. If \mathbf{X} is recurrent, then \mathbf{X} always has a positive invariant function that is unique up to multiplication by positive constants. Hence the reversal of a recurrent chain \mathbf{X} always exists and is unique, and so we can refer to the reversal of \mathbf{X} without reference to the invariant function.
2. Even better, if \mathbf{X} is positive recurrent, then there exists a unique invariant probability density function, and the reversal of \mathbf{X} can be interpreted as the *time* reversal (with respect to a terminal time) when \mathbf{X} has the invariant distribution, as in the motivating exercises above.
3. If \mathbf{X} is transient, then there may or may not exist a positive invariant function, and if one does exist, it may not be unique (up to multiplication by positive constants). So a transient chain may have no reversals or more than one.

Nonetheless, the general definition is natural, because most of the important properties of the reversed chain follow from the *balance equation* between the transition matrices P and \hat{P} , and the invariant function g :

$$g(x) \hat{P}(x, y) = g(y) P(y, x), \quad (x, y) \in S^2 \quad (16.7.10)$$

We will see this balance equation repeated with other objects related to the Markov chains.

Suppose that \mathbf{X} is an irreducible Markov chain with invariant function $g : S \rightarrow (0, \infty)$, and that $\hat{\mathbf{X}}$ is the reversal of \mathbf{X} with respect to g . For $x, y \in S$,

1. $\hat{P}(x, x) = P(x, x)$
2. $\hat{P}(x, y) > 0$ if and only if $P(y, x) > 0$

Proof

These results follow immediately from the balance equation $g(x)\hat{P}(x, y) = g(y)P(y, x)$ for $(x, y) \in S^2$.

From part (b) it follows that the state graphs of \mathbf{X} and $\hat{\mathbf{X}}$ are reverses of each other. That is, to go from the state graph of one chain to the state graph of the other, simply reverse the direction of each edge. Here is a more complicated (but equivalent) version of the balance equation for chains of states:

Suppose again that \mathbf{X} is an irreducible Markov chain with invariant function $g : S \rightarrow (0, \infty)$, and that $\hat{\mathbf{X}}$ is the reversal of \mathbf{X} with respect to g . For every $n \in \mathbb{N}_+$ and every sequence of states $(x_1, x_2, \dots, x_n, x_{n+1}) \in S^{n+1}$,

$$g(x_1)\hat{P}(x_1, x_2)\hat{P}(x_2, x_3) \cdots \hat{P}(x_n, x_{n+1}) = g(x_{n+1})P(x_{n+1}, x_n) \cdots P(x_3, x_2)P(x_2, x_1) \quad (16.7.11)$$

Proof

This follows from repeated applications of the basic equation. When $n = 1$, we have the balance equation itself:

$$g(x_1)\hat{P}(x_1, x_2) = g(x_2)P(x_2, x_1) \quad (16.7.12)$$

For $n = 2$,

$$g(x_1)\hat{P}(x_1, x_2)\hat{P}(x_2, x_3) = g(x_2)P(x_2, x_1)\hat{P}(x_2, x_3) = g(x_3)P(x_3, x_2)P(x_2, x_1) \quad (16.7.13)$$

Continuing in this manner (or using induction) gives the general result.

The balance equation holds for the powers of the transition matrix:

Suppose again that \mathbf{X} is an irreducible Markov chain with invariant function $g : S \rightarrow (0, \infty)$, and that $\hat{\mathbf{X}}$ is the reversal of \mathbf{X} with respect to g . For every $(x, y) \in S^2$ and $n \in \mathbb{N}$,

$$g(x)\hat{P}^n(x, y) = g(y)P^n(y, x) \quad (16.7.14)$$

Proof

When $n = 0$, the left and right sides are $g(x)$ if $x = y$ and are 0 otherwise. When $n = 1$, we have the basic balance equation: $g(x)\hat{P}(x, y) = g(y)P(y, x)$. In general, for $n \in \mathbb{N}_+$, by the [previous result](#) we have

$$\begin{aligned} g(x)\hat{P}^n(x, y) &= \sum_{(x_1, \dots, x_{n-1}) \in S^{n-1}} g(x)\hat{P}(x, x_1)\hat{P}(x_1, x_2) \cdots \hat{P}(x_{n-1}, y) \\ &= \sum_{(x_1, \dots, x_{n-1}) \in S^{n-1}} g(y)P(y, x_{n-1})P(x_{n-1}, x_{n-2}) \cdots P(x_1, x) = g(y)P^n(y, x) \end{aligned}$$

We can now generalize the simple result [above](#).

Suppose again that \mathbf{X} is an irreducible Markov chain with invariant function $g : S \rightarrow (0, \infty)$, and that $\hat{\mathbf{X}}$ is the reversal of \mathbf{X} with respect to g . For $n \in \mathbb{N}$ and $(x, y) \in S^2$,

1. $P^n(x, x) = \hat{P}^n(x, x)$
2. $\hat{P}^n(x, y) > 0$ if and only if $P^n(y, x) > 0$

In terms of the state graphs, part (b) has an obvious meaning: If there exists a path of length n from y to x in the original state graph, then there exists a path of length n from x to y in the reversed state graph. The time reversal definition is symmetric with respect to the two Markov chains.

Suppose again that \mathbf{X} is an irreducible Markov chain with invariant function $g : S \rightarrow (0, \infty)$, and that $\hat{\mathbf{X}}$ is the reversal of \mathbf{X} with respect to g . Then

1. g is also invariant for $\hat{\mathbf{X}}$.
2. $\hat{\mathbf{X}}$ is also irreducible.
3. \mathbf{X} is the reversal of $\hat{\mathbf{X}}$ with respect to g .

Proof

1. For $y \in S$, using the balance equation,

$$\sum_{x \in S} g(x) \hat{P}(x, y) = \sum_{x \in S} g(y) P(y, x) = g(y) \quad (16.7.15)$$

2. Suppose $(x, y) \in S^2$. Since \mathbf{X} is irreducible, there exist $n \in \mathbb{N}$ with $P^n(y, x) > 0$. But then from the previous result, $\hat{P}^n(x, y) > 0$. Hence $\hat{\mathbf{X}}$ is also irreducible.
3. This is clear from the symmetric relationship in the fundamental result.

The balance equation also holds for the potential matrices.

Suppose that \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals with respect to the invariant function $g : S \rightarrow (0, \infty)$. For $\alpha \in (0, 1]$, the α potential matrices are related by

$$g(x) \hat{R}_\alpha(x, y) = g(y) R_\alpha(y, x), \quad (x, y) \in S^2 \quad (16.7.16)$$

Proof

This follows easily from the result [above](#) and the definition of the potential matrices:

$$\begin{aligned} g(x) \hat{R}_\alpha(x, y) &= g(x) \sum_{n=0}^{\infty} \alpha^n \hat{P}^n(x, y) = \sum_{n=0}^{\infty} \alpha^n g(x) \hat{P}^n(x, y) \\ &= \sum_{n=0}^{\infty} \alpha^n g(y) P^n(y, x) = g(y) \sum_{n=0}^{\infty} \alpha^n P^n(y, x) = g(y) R_\alpha(y, x) \end{aligned}$$

Markov chains that are time reversals share many important properties:

Suppose that \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals. Then

1. \mathbf{X} and $\hat{\mathbf{X}}$ are of the same type (transient, null recurrent, or positive recurrent).
2. \mathbf{X} and $\hat{\mathbf{X}}$ have the same period.
3. \mathbf{X} and $\hat{\mathbf{X}}$ have the same mean return time $\mu(x)$ for every $x \in S$.

Proof

Suppose that \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals with respect to the invariant function $g : S \rightarrow (0, \infty)$.

1. The expected number of visits to a state $x \in S$, starting in x , is the same for both chains: $\hat{R}(x, x) = R(x, x)$. Hence either both chains are transient (if the common potential is finite) or both chains are recurrent (if the common potential is infinite). If both chains are recurrent then the invariant function g is unique up to multiplication by positive constants, and both are null recurrent if $\sum_{x \in S} g(x) = \infty$ and both are positive recurrent if $\sum_{x \in S} g(x) < \infty$.
2. This follows since $P^n(x, x) = \hat{P}^n(x, x)$ for all $n \in \mathbb{N}$ and $x \in S$.
3. If both chains are transient or both are null recurrent, then $\mu(x) = \hat{\mu}(x) = \infty$ for all $x \in S$. If both chains are positive recurrent, then for all $n \in \mathbb{N}$ and $x \in S$, we have

$$\frac{1}{n} \sum_{k=1}^n P^k(x, x) = \frac{1}{n} \sum_{k=1}^n \hat{P}^k(x, x) \quad (16.7.17)$$

The left side converges to $1/\mu(x)$ as $n \rightarrow \infty$ while the right side converges to $1/\hat{\mu}(x)$ as $n \rightarrow \infty$.

The main point of the next result is that we don't need to know a-priori that g is invariant for \mathbf{X} , if we can guess g and \hat{P} .

Suppose again that \mathbf{X} is irreducible with transition probability matrix P . If there exists a function $g: S \rightarrow (0, \infty)$ and a transition probability matrix \hat{P} such that $g(x)\hat{P}(x, y) = g(y)P(y, x)$ for all $(x, y) \in S^2$, then

1. g is invariant for \mathbf{X} .
2. \hat{P} is the transition matrix of the reversal of \mathbf{X} with respect to g .

Proof

1. Since \hat{P} is a transition probability matrix, we have the same computation we have seen before:

$$gP(x) = \sum_{y \in S} g(y)P(y, x) = \sum_{y \in S} g(x)\hat{P}(x, y) = g(x), \quad x \in S \quad (16.7.18)$$

2. This follows from (a) and the definition.

As a corollary, if there exists a probability density function f on S and a transition probability matrix \hat{P} such that $f(x)\hat{P}(x, y) = f(y)P(y, x)$ for all $(x, y) \in S^2$ then in addition to the conclusions above, we know that the chains \mathbf{X} and $\hat{\mathbf{X}}$ are positive recurrent.

Reversible Chains

Clearly, an interesting special case occurs when the transition matrix of the reversed chain turns out to be the same as the original transition matrix. A chain of this type could be used to model a physical process that is stochastically the same, forward or backward in time.

Suppose again that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is an irreducible Markov chain with transition matrix P and invariant function $g: S \rightarrow (0, \infty)$. If the reversal of \mathbf{X} with respect to g also has transition matrix P , then \mathbf{X} is said to be *reversible* with respect to g . That is, \mathbf{X} is reversible with respect to g if and only if

$$g(x)P(x, y) = g(y)P(y, x), \quad (x, y) \in S^2 \quad (16.7.19)$$

Clearly if \mathbf{X} is reversible with respect to the invariant function $g: S \rightarrow (0, \infty)$ then \mathbf{X} is reversible with respect to the invariant function cg for every $c \in (0, \infty)$. So again, let's review the cases.

Suppose that \mathbf{X} is an irreducible Markov chain on S .

1. If \mathbf{X} is recurrent, there exists a positive invariant function that is unique up to multiplication by positive constants. So \mathbf{X} is either reversible or not, and we don't have to reference the invariant function g .
2. If \mathbf{X} is positive recurrent then there exists a unique invariant probability density function $f: S \rightarrow (0, 1)$, and again, either \mathbf{X} is reversible or not. If \mathbf{X} is reversible, then P is the transition matrix of \mathbf{X} forward or backward in time, when the chain has the invariant distribution.
3. If \mathbf{X} is transient, there may or may not exist positive invariant functions. If there are two or more positive invariant functions that are not multiplies of one another, \mathbf{X} might be reversible with respect to one function but not the others.

The non-symmetric simple random walk on \mathbb{Z} falls into the last case. Using the [last result](#) in the previous subsection, we can tell whether \mathbf{X} is reversible with respect to g without knowing a-priori that g is invariant.

Suppose again that \mathbf{X} is irreducible with transition matrix P . If there exists a function $g: S \rightarrow (0, \infty)$ such that $g(x)P(x, y) = g(y)P(y, x)$ for all $(x, y) \in S^2$, then

1. g is invariant for \mathbf{X} .
2. \mathbf{X} is reversible with respect to g

If we have reason to believe that a Markov chain is reversible (based on modeling considerations, for example), then the condition in the previous theorem can be used to find the invariant functions. This procedure is often easier than using the definition of invariance directly. The next two results are minor generalizations:

Suppose again that \mathbf{X} is irreducible and that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if for every $n \in \mathbb{N}_+$ and every sequence of states $(x_1, x_2, \dots, x_n, x_{n+1}) \in S^{n+1}$,

$$g(x_1)P(x_1, x_2)P(x_2, x_3) \cdots P(x_n, x_{n+1}) = g(x_{n+1})P(x_{n+1}, x_n) \cdots P(x_3, x_2)P(x_2, x_1) \quad (16.7.20)$$

Suppose again that \mathbf{X} is irreducible and that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if for every $(x, y) \in S^2$ and $n \in \mathbb{N}_+$,

$$g(x)P^n(x, y) = g(y)P^n(y, x) \quad (16.7.21)$$

Here is the condition for reversibility in terms of the potential matrices.

Suppose again that \mathbf{X} is irreducible and that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if

$$g(x)R_\alpha(x, y) = g(y)R_\alpha(y, x), \quad \alpha \in (0, 1], (x, y) \in S^2 \quad (16.7.22)$$

In the positive recurrent case (the most important case), the following theorem gives a condition for reversibility that does not directly reference the invariant distribution. The condition is known as the *Kolmogorov cycle condition*, and is named for Andrei Kolmogorov

Suppose that \mathbf{X} is irreducible and positive recurrent. Then \mathbf{X} is reversible if and only if for every sequence of states (x_1, x_2, \dots, x_n) ,

$$P(x_1, x_2)P(x_2, x_3) \cdots P(x_{n-1}, x_n)P(x_n, x_1) = P(x_1, x_n)P(x_n, x_{n-1}) \cdots P(x_3, x_2)P(x_2, x_1) \quad (16.7.23)$$

Proof

Suppose that \mathbf{X} is reversible. Applying the chain result [above](#) to the sequence $(x_1, x_2, \dots, x_n, x_1)$ gives the Kolmogorov cycle condition. Conversely, suppose that the Kolmogorov cycle condition holds, and let f denote the invariant probability density function of \mathbf{X} . From the cycle condition we have $P(x, y)P^k(y, x) = P(y, x)P^k(x, y)$ for every $(x, y) \in S$ and $k \in \mathbb{N}_+$. Averaging over k from 1 to n gives

$$P(x, y) \frac{1}{n} \sum_{k=1}^n P^k(y, x) = P(y, x) \frac{1}{n} \sum_{k=1}^n P^k(x, y), \quad (x, y) \in S^2, n \in \mathbb{N}_+ \quad (16.7.24)$$

Letting $n \rightarrow \infty$ gives $f(x)P(x, y) = f(y)P(y, x)$ for $(x, y) \in S^2$, so \mathbf{X} is reversible.

Note that the Kolmogorov cycle condition states that the probability of visiting states $(x_2, x_3, \dots, x_n, x_1)$ in sequence, starting in state x_1 is the same as the probability of visiting states $(x_n, x_{n-1}, \dots, x_2, x_1)$ in sequence, starting in state x_1 . The cycle condition is also known as the *balance equation for cycles*.




Figure 16.7.1: The Kolmogorov cycle condition

Examples and Applications

Finite Chains

Recall the general two-state chain \mathbf{X} on $S = \{0, 1\}$ with the transition probability matrix

$$P = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix} \quad (16.7.25)$$

where $p, q \in (0, 1)$ are parameters. The chain \mathbf{X} is reversible and the invariant probability density function is $f = \left(\frac{q}{p+q}, \frac{p}{p+q} \right)$.

Proof

All we have to do is note that

$$\begin{bmatrix} q & p \end{bmatrix} \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix} = \begin{bmatrix} q & p \end{bmatrix} \quad (16.7.26)$$

Suppose that \mathbf{X} is a Markov chain on a finite state space S with symmetric transition probability matrix P . Thus $P(x, y) = P(y, x)$ for all $(x, y) \in S^2$. The chain \mathbf{X} is reversible and that the uniform distribution on S is invariant.

Proof

All we have to do is note that $\mathbf{1}(x)P(x, y) = \mathbf{1}(y)P(y, x)$ where $\mathbf{1}$ is the constant function 1 on S .

Consider the Markov chain \mathbf{X} on $S = \{a, b, c\}$ with transition probability matrix P given below:

$$P = \begin{bmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{2} & 0 \end{bmatrix} \quad (16.7.27)$$

1. Draw the state graph of \mathbf{X} and note that the chain is irreducible.
2. Find the invariant probability density function f .
3. Find the mean return time to each state.
4. Find the transition probability matrix \hat{P} of the time-reversed chain $\hat{\mathbf{X}}$.
5. Draw the state graph of $\hat{\mathbf{X}}$.

Answer

1.

State graph of \mathbf{X}



2. $f = \left(\frac{6}{17}, \frac{6}{17}, \frac{5}{17} \right)$

3. $\mu = \left(\frac{17}{6}, \frac{17}{6}, \frac{17}{5} \right)$

4. $\hat{P} = \begin{bmatrix} \frac{1}{4} & \frac{1}{3} & \frac{5}{12} \\ \frac{1}{4} & \frac{1}{3} & \frac{5}{12} \\ \frac{3}{5} & \frac{2}{5} & 0 \end{bmatrix}$

5.

State graph of $\hat{\mathbf{X}}$



Special Models

Read the discussion of reversibility for the Ehrenfest chains.

Read the discussion of reversibility for the Bernoulli-Laplace chain.

Read the discussion of reversibility for the random walks on graphs.

Read the discussion of time reversal for the reliability chains.

Read the discussion of reversibility for the birth-death chains.

This page titled [16.7: Time Reversal in Discrete-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.8: The Ehrenfest Chains

Basic Theory

The *Ehrenfest chains*, named for Paul Ehrenfest, are simple, discrete models for the exchange of gas molecules between two containers. However, they can be formulated as simple ball and urn models; the balls correspond to the molecules and the urns to the two containers. Thus, suppose that we have two urns, labeled 0 and 1, that contain a total of m balls. The state of the system at time $n \in \mathbb{N}$ is the number of balls in urn 1, which we will denote by X_n . Our stochastic process is $\mathbf{X} = (X_0, X_1, X_2, \dots)$ with state space $S = \{0, 1, \dots, m\}$. Of course, the number of balls in urn 0 at time n is $m - X_n$.

The Models

In the basic Ehrenfest model, at each discrete time unit, independently of the past, a ball is selected at random and moved to the other urn.




Figure 16.8.1: The Ehrenfest model

\mathbf{X} is a discrete-time Markov chain on S with transition probability matrix P given by

$$P(x, x-1) = \frac{x}{m}, \quad P(x, x+1) = \frac{m-x}{m}, \quad x \in S \quad (16.8.1)$$

Proof

We will give a construction of the chain from a more basic process. Let V_n be the ball selected at time $n \in \mathbb{N}_+$. Thus $\mathbf{V} = (V_1, V_2, \dots)$ is a sequence of independent random variables, each uniformly distributed on $\{1, 2, \dots, m\}$. Let $X_0 \in S$ be independent of \mathbf{V} . (We can start the process any way that we like.) Now define the state process recursively as follows:

$$X_{n+1} = \begin{cases} X_n - 1, & V_{n+1} \leq X_n \\ X_n + 1, & V_{n+1} > X_n \end{cases}, \quad n \in \mathbb{N} \quad (16.8.2)$$

In the Ehrenfest experiment, select the basic model. For selected values of m and selected values of the initial state, run the chain for 1000 time steps and note the limiting behavior of the proportion of time spent in each state.

Suppose now that we modify the basic Ehrenfest model as follows: at each discrete time, independently of the past, we select a ball at random and a urn at random. We then put the chosen ball in the chosen urn.

\mathbf{X} is a discrete-time Markov chain on S with the transition probability matrix Q given by

$$Q(x, x-1) = \frac{x}{2m}, \quad Q(x, x) = \frac{1}{2}, \quad Q(x, x+1) = \frac{m-x}{2m}, \quad x \in S \quad (16.8.3)$$

Proof

Again, we can construct the chain from a more basic process. Let X_0 and \mathbf{V} be as in Theorem 1. Let U_n be the urn selected at time $n \in \mathbb{N}_+$. Thus $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent random variables, each uniformly distributed on $\{0, 1\}$ (so that \mathbf{U} is a fair, Bernoulli trials sequence). Also, \mathbf{U} is independent of \mathbf{V} and X_0 . Now define the state process recursively as follows:

$$X_{n+1} = \begin{cases} X_n - 1, & V_{n+1} \leq X_n, U_{n+1} = 0 \\ X_n + 1, & V_{n+1} > X_n, U_{n+1} = 1 \\ X_n, & \text{otherwise} \end{cases}, \quad n \in \mathbb{N} \quad (16.8.4)$$

Note that $Q(x, y) = \frac{1}{2}P(x, y)$ for $y \in \{x-1, x+1\}$.

In the Ehrenfest experiment, select the modified model. For selected values of m and selected values of the initial state, run the chain for 1000 time steps and note the limiting behavior of the proportion of time spent in each state.

Classification

The basic and modified Ehrenfest chains are irreducible and positive recurrent.

Proof

The chains are clearly irreducible since every state leads to every other state. It follows that the chains are positive recurrent since the state space S is finite.

The basic Ehrenfest chain is periodic with period 2. The cyclic classes are the set of even states and the set of odd states. The two-step transition matrix is

$$P^2(x, x-2) = \frac{x(x-1)}{m^2}, \quad P^2(x, x) = \frac{x(m-x+1) + (m-x)(x+1)}{m^2}, \quad P^2(x, x+2) = \frac{(m-x)(m-x-1)}{m^2}, \quad x \in S \quad (16.8.5)$$

Proof

Note that returns to a state can only occur at even times, so the chain has period 2. The form of P^2 follows from the [formula for \$P\$](#) above.

The modified Ehrenfest chain is aperiodic.

Proof

Note that $P(x, x) > 0$ for each $x \in S$.

Invariant and Limiting Distributions

For the basic and modified Ehrenfest chains, the invariant distribution is the binomial distribution with trial parameter m and success parameter $\frac{1}{2}$. So the invariant probability density function f is given by

$$f(x) = \binom{m}{x} \left(\frac{1}{2}\right)^m, \quad x \in S \quad (16.8.6)$$

Proof

For the basic chain we have

$$\begin{aligned} (fP)(y) &= f(y-1)P(y-1, y) + f(y+1)P(y+1, y) \\ &= \binom{m}{y-1} \left(\frac{1}{2}\right)^m \frac{m-y+1}{m} + \binom{m}{y+1} \left(\frac{1}{2}\right)^m \frac{y+1}{m} \\ &= \left(\frac{1}{2}\right)^m \left[\binom{m-1}{y-1} + \binom{m-1}{y} \right] = \left(\frac{1}{2}\right)^m \binom{m}{y} = f(y), \quad y \in S \end{aligned}$$

The last step uses a fundamental identity for binomial coefficients. For the modified chain we can use the result for the basic chain:

$$\begin{aligned} (fQ)(y) &= f(y-1)Q(y-1, y) + f(y)Q(y, y) + f(y+1)Q(y+1, y) \\ &= \frac{1}{2}f(y-1)P(y-1, y) + \frac{1}{2}f(y+1)P(y+1, y) + \frac{1}{2}f(y) = f(y), \quad y \in S \end{aligned}$$

Thus, the invariant distribution corresponds to placing each ball randomly and independently either in urn 0 or in urn 1.

The mean return time to state $x \in S$ for the basic or modified Ehrenfest chain is $\mu(x) = 2^m / \binom{m}{x}$.

Proof

This follows from the general theory and the invariant distribution [above](#).

For the basic Ehrenfest chain, the limiting behavior of the chain is as follows:

1. $P^{2n}(x, y) \rightarrow \binom{m}{y} \left(\frac{1}{2}\right)^{m-1}$ as $n \rightarrow \infty$ if $x, y \in S$ have the same parity (both even or both odd). The limit is 0 otherwise.
2. $P^{2n+1}(x, y) \rightarrow \binom{m}{y} \left(\frac{1}{2}\right)^{m-1}$ as $n \rightarrow \infty$ if $x, y \in S$ have opposite parity (one even and one odd). The limit is 0 otherwise.

Proof

These results follow from the general theory and the invariant distribution [above](#), and the fact that the chain is periodic with period 2, with the odd and even integers in S as the cyclic classes.

For the modified Ehrenfest chain, $Q^n(x, y) \rightarrow \binom{m}{y} \left(\frac{1}{2}\right)^m$ as $n \rightarrow \infty$ for $x, y \in S$.

Proof

Again, this follows from the general theory and the invariant distribution [above](#), and the fact that the chain is aperiodic.

In the Ehrenfest experiment, the limiting binomial distribution is shown graphically and numerically. For each model and for selected values of m and selected values of the initial state, run the chain for 1000 time steps and note the limiting behavior of the proportion of time spent in each state. How do the choices of m , the initial state, and the model seem to affect the rate of convergence to the limiting distribution?

Reversibility

The basic and modified Ehrenfest chains are reversible.

Proof

Let $g(x) = \binom{m}{x}$ for $x \in S$. The crucial observations are $g(x)P(x, y) = g(y)P(y, x)$ and $g(x)Q(x, y) = g(y)Q(y, x)$ for all $x, y \in S$. For the basic chain, if $x \in S$ then

$$g(x)P(x, x-1) = g(x-1)P(x-1, x) = \binom{m-1}{x-1}$$

$$g(x)P(x, x+1) = g(x+1)P(x+1, x) = \binom{m-1}{x}$$

In all other cases, $g(x)P(x, y) = g(y)P(y, x) = 0$. The reversibility condition for the modified chain follows trivially from that of the basic chain since $Q(x, y) = \frac{1}{2}P(x, y)$ for $y = x \pm 1$ (and of course the reversibility condition is trivially satisfied when $x = y$). Note that the invariant PDF f is simply g normalized. The reversibility condition gives another (and better) proof that f is invariant.

Run the simulation of the Ehrenfest experiment 10,000 time steps for each model, for selected values of m , and with initial state 0. Note that at first, you can see the “arrow of time”. After a long period, however, the direction of time is no longer evident.

Computational Exercises

Consider the basic Ehrenfest chain with $m = 5$ balls, and suppose that X_0 has the uniform distribution on S .

1. Compute the probability density function, mean and variance of X_1 .
2. Compute the probability density function, mean and variance of X_2 .
3. Compute the probability density function, mean and variance of X_3 .
4. Sketch the initial probability density function and the probability density functions in parts (a), (b), and (c) on a common set of axes.

Answer

1. $f_1 = \left(\frac{1}{30}, \frac{7}{30}, \frac{7}{30}, \frac{7}{30}, \frac{1}{30}\right), \mu_1 = \frac{5}{2}, \sigma_1^2 = \frac{19}{12}$
2. $f_2 = \left(\frac{7}{150}, \frac{19}{150}, \frac{49}{150}, \frac{49}{150}, \frac{19}{150}, \frac{7}{150}\right), \mu_2 = \frac{5}{2}, \sigma_2^2 = \frac{79}{60}$
3. $f_3 = \left(\frac{19}{750}, \frac{133}{750}, \frac{223}{150}, \frac{223}{150}, \frac{133}{150}, \frac{19}{150}\right), \mu_3 = \frac{5}{2}, \sigma_3^2 = \frac{431}{300}$

Consider the modified Ehrenfest chain with $m = 5$ balls, and suppose that the chain starts in state 2 (with probability 1).

1. Compute the probability density function, mean and standard deviation of X_1 .
2. Compute the probability density function, mean and standard deviation of X_2 .
3. Compute the probability density function, mean and standard deviation of X_3 .
4. Sketch the initial probability density function and the probability density functions in parts (a), (b), and (c) on a common set of axes.

Answer

1. $f_1 = (0, 0.2, 0.5, 0.3, 0, 0), \mu_1 = 2.1, \sigma_1 = 0.7$
2. $f_2 = (0.02, 0.20, 0.42, 0.30, 0.06, 0), \mu_2 = 2.18, \sigma_2 = 0.887$
3. $f_3 = (0.030, 0.194, 0.380, 0.300, 0.090, 0.006), \mu_3 = 2.244, \sigma_3 = 0.984$

This page titled [16.8: The Ehrenfest Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.9: The Bernoulli-Laplace Chain

Basic Theory

Introduction

The *Bernoulli-Laplace chain*, named for Jacob Bernoulli and Pierre Simon Laplace, is a simple discrete model for the diffusion of two incompressible gases between two containers. Like the Ehrenfest chain, it can also be formulated as a simple ball and urn model. Thus, suppose that we have two urns, labeled 0 and 1. Urn 0 contains j balls and urn 1 contains k balls, where $j, k \in \mathbb{N}_+$. Of the $j+k$ balls, r are red and the remaining $j+k-r$ are green. Thus $r \in \mathbb{N}_+$ and $0 < r < j+k$. At each discrete time, independently of the past, a ball is selected at random from each urn and then the two balls are switched. The balls of different colors correspond to molecules of different types, and the urns are the containers. The incompressible property is reflected in the fact that the number of balls in each urn remains constant over time.




Figure 16.9.1: The Bernoulli-Laplace model

Let X_n denote the number of red balls in urn 1 at time $n \in \mathbb{N}$. Then

1. $k - X_n$ is the number of green balls in urn 1 at time n .
2. $r - X_n$ is the number of red balls in urn 0 at time n .
3. $j - r + X_n$ is the number of green balls in urn 0 at time n .

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a discrete-time Markov chain with state space $S = \{\max\{0, r-j\}, \dots, \min\{k, r\}\}$ and with transition matrix P given by

$$P(x, x-1) = \frac{(j-r+x)x}{jk}, \quad P(x, x) = \frac{(r-x)x + (j-r+x)(k-x)}{jk}, \quad P(x, x+1) = \frac{(r-x)(k-x)}{jk}; \quad x \in S \quad (16.9.1)$$

Proof

For the state space, note from the previous result that the number of red balls x in urn 1 must satisfy the inequalities $x \geq 0$, $x \leq k$, $x \leq r$, and $x \geq r-j$. The Markov property is clear from the model. For the transition probabilities, note that to go from state x to state $x-1$ we must select a green ball from urn 0 and a red ball from urn 1. The probabilities of these events are $(j-r+x)/j$ and x/k for x and $x-1$ in S , and the events are independent. Similarly, to go from state x to state $x+1$ we must select a red ball from urn 0 and a green ball from urn 1. The probabilities of these events are $(r-x)/j$ and $(k-x)/k$ for x and $x+1$ in S , and the events are independent. Finally, to go from state x back to state x , we must select a red ball from both urns or a green ball from both urns. Of course also,

$$P(x, x) = 1 - P(x, x-1) - P(x, x+1).$$

This is a fairly complicated model, simply because of the number of parameters. Interesting special cases occur when some of the parameters are the same.

Consider the special case $j = k$, so that each urn has the same number of balls. The state space is $S = \{\max\{0, r-k\}, \dots, \min\{k, r\}\}$ and the transition probability matrix is

$$P(x, x-1) = \frac{(k-r+x)x}{k^2}, \quad P(x, x) = \frac{(r-x)x + (k-r+x)(k-x)}{k^2}, \quad P(x, x+1) = \frac{(r-x)(k-x)}{k^2}; \quad x \in S \quad (16.9.2)$$

Consider the special case $r = j$, so that the number of red balls is the same as the number of balls in urn 0. The state space is $S = \{0, \dots, \min\{j, k\}\}$ and the transition probability matrix is

$$P(x, x-1) = \frac{x^2}{jk}, \quad P(x, x) = \frac{x(j+k-2x)}{jk}, \quad P(x, x+1) = \frac{(j-x)(k-x)}{jk}; \quad x \in S \quad (16.9.3)$$

Consider the special case $r = k$, so that the number of red balls is the same as the number of balls in urn 1. The state space is $S = \{\max\{0, k-j\}, \dots, k\}$ and the transition probability matrix is

$$P(x, x-1) = \frac{(j-k+x)x}{jk}, \quad P(x, x) = \frac{(k-x)(j-k+2x)}{jk}, \quad P(x, x+1) = \frac{(k-x)^2}{jk}; \quad x \in S \quad (16.9.4)$$

Consider the special case $j = k = r$, so that each urn has the same number of balls, and this is also the number of red balls. The state space is $S = \{0, 1, \dots, k\}$ and the transition probability matrix is

$$P(x, x-1) = \frac{x^2}{k^2}, \quad P(x, x) = \frac{2x(k-x)}{k^2}, \quad P(x, x+1) = \frac{(k-x)^2}{k^2}; \quad x \in S \quad (16.9.5)$$

Run the simulation of the Bernoulli-Laplace experiment for 10000 steps and for various values of the parameters. Note the limiting behavior of the proportion of time spent in each state.

Invariant and Limiting Distributions

The Bernoulli-Laplace chain is irreducible.

Proof

Note that $P(x, x-1) > 0$ whenever $x, x-1 \in S$, and $P(x, x+1) > 0$ whenever $x, x+1 \in S$. Hence every state leads to every other state so the chain is irreducible.

Except in the trivial case $j = k = r = 1$, the Bernoulli-Laplace chain aperiodic.

Proof

Consideration of the state probabilities shows that except when $j = k = r = 1$, the chain has a state x with $P(x, x) > 0$, so state x is aperiodic. Since the chain is irreducible by the previous result, all states are aperiodic.

The invariant distribution is the hypergeometric distribution with population parameter $j+k$, sample parameter k , and type parameter r . The probability density function is

$$f(x) = \frac{\binom{r}{x} \binom{j+k-r}{k-x}}{\binom{j+k}{k}}, \quad x \in S \quad (16.9.6)$$

Proof

A direct proof that $(fP)(x) = f(x)$ for all $x \in S$ is straightforward but tedious. A better proof follows from the reversibility condition [below](#).

Thus, the invariant distribution corresponds to selecting a sample of k balls at random and without replacement from the $j+k$ balls and placing them in urn 1. The mean and variance of the invariant distribution are

$$\mu = k \frac{r}{j+k}, \quad \sigma^2 = k \frac{r}{j+k} \frac{j+k-r}{j+k} \frac{j}{j+k-1} \quad (16.9.7)$$

The mean return time to each state $x \in S$ is

$$\mu(x) = \frac{\binom{j+k}{k}}{\binom{r}{x} \binom{j+k-r}{k-x}} \quad (16.9.8)$$

Proof

This follows from the general theory and the invariant distribution [above](#).

$P^n(x, y) \rightarrow f(y) = \frac{\binom{r}{y} \binom{j+k-r}{k-y}}{\binom{j+k}{k}}$ as $n \rightarrow \infty$ for $(x, y) \in S^2$.

Proof

This follows from the general theory and the invariant distribution [above](#).

In the simulation of the Bernoulli-Laplace experiment, vary the parameters and note the shape and location of the limiting hypergeometric distribution. For selected values of the parameters, run the simulation for 10000 steps and note the limiting behavior of the proportion of time spent in each state.

Reversibility

The Bernoulli-Laplace chain is reversible.

Proof

Let

$$g(x) = \binom{r}{x} \binom{j+k-r}{k-x}, \quad x \in S \quad (16.9.9)$$

It suffices to show the reversibility condition $g(x)P(x, y) = g(y)P(y, x)$ for all $x, y \in S$. It then follows that \mathbf{X} is reversible and that g is invariant for \mathbf{X} . For $x \in S$ and $y = x - 1 \in S$, the left and right sides of the reversibility condition reduce to

$$\frac{1}{jk} \frac{r!}{(x-1)!(r-x)!} \frac{(j+k-r)!}{(k-x)!(j-r+x-1)!} \quad (16.9.10)$$

For $x \in S$ and $y = x + 1 \in S$, the left and right sides of the reversibility condition reduce to

$$\frac{1}{jk} \frac{r!}{x!(r-x-1)!} \frac{(j+k-r)!}{(k-x-1)!(j-r+x)!} \quad (16.9.11)$$

For all other values of $x, y \in S$, the reversibility condition is trivially satisfied. The hypergeometric PDF f above is simply g normalized, so this proves that f is also invariant.

Run the simulation of the Bernoulli-Laplace experiment 10,000 time steps for selected values of the parameters, and with initial state 0. Note that at first, you can see the “arrow of time”. After a long period, however, the direction of time is no longer evident.

Computational Exercises

Consider the Bernoulli-Laplace chain with $j = 10$, $k = 5$, and $r = 4$. Suppose that X_0 has the uniform distribution on S . Explicitly give each of the following:

1. The state space S
2. The transition matrix P .
3. The probability density function, mean and variance of X_1 .
4. The probability density function, mean and variance of X_2 .
5. The probability density function, mean and variance of X_3 .

Answer

1. $S = \{0, 1, 2, 3, 4\}$

$$2. P = \frac{1}{50} \begin{bmatrix} 30 & 20 & 0 & 0 & 0 \\ 7 & 31 & 12 & 0 & 0 \\ 0 & 16 & 28 & 6 & 0 \\ 0 & 0 & 27 & 21 & 2 \\ 0 & 0 & 0 & 40 & 10 \end{bmatrix}$$

$$3. f_1 = \frac{1}{250} (37, 67, 67, 67, 12), \mu_1 = \frac{9}{5}, \sigma_1^2 = \frac{32}{25}$$

$$4. f_2 = \frac{1}{12500} (1579, 3889, 4489, 2289, 254), \mu_2 = \frac{83}{50}, \sigma_2^2 = \frac{2413}{2500}$$

$$5. f_3 = \frac{1}{625000} (74593, 223963, 234163, 85163, 7118), \mu_3 = \frac{781}{500}, \sigma_3^2 = \frac{206427}{250000}$$

Consider the Bernoulli-Laplace chain with $j = k = 10$ and $r = 6$. Give each of the following explicitly:

1. The state space S
2. The transition matrix P
3. The invariant probability density function.

Answer

1. $S = \{0, 1, 2, 3, 4, 5, 6\}$

$$2. P = \frac{1}{100} \begin{bmatrix} 40 & 60 & 0 & 0 & 0 & 0 & 0 \\ 5 & 50 & 45 & 0 & 0 & 0 & 0 \\ 0 & 12 & 56 & 32 & 0 & 0 & 0 \\ 0 & 0 & 21 & 58 & 21 & 0 & 0 \\ 0 & 0 & 0 & 32 & 56 & 12 & 0 \\ 0 & 0 & 0 & 0 & 45 & 50 & 5 \\ 0 & 0 & 0 & 0 & 0 & 60 & 40 \end{bmatrix}$$

$$3. f = \frac{1}{1292} (7, 84, 315, 480, 315, 84, 7)$$

This page titled [16.9: The Bernoulli-Laplace Chain](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.10: Discrete-Time Reliability Chains

The Success-Runs Chain

Suppose that we have a sequence of *trials*, each of which results in either *success* or *failure*. Our basic assumption is that if there have been $x \in \mathbb{N}$ consecutive successes, then the probability of success on the next trial is $p(x)$, independently of the past, where $p : \mathbb{N} \rightarrow (0, 1)$. Whenever there is a failure, we start over, independently, with a new sequence of trials. Appropriately enough, p is called the *success function*. Let X_n denote the length of the run of successes after n trials.

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a discrete-time Markov chain with state space \mathbb{N} and transition probability matrix P given by

$$P(x, x+1) = p(x), \quad P(x, 0) = 1 - p(x); \quad x \in \mathbb{N} \quad (16.10.1)$$

The Markov chain \mathbf{X} is called the *success-runs chain*.




Figure 16.10.1: State graph of the success-runs chain

Now let T denote the trial number of the first failure, starting with a fresh sequence of trials. Note that in the context of the success-runs chain \mathbf{X} , $T = \tau_0$, the first return time to state 0, starting in 0. Note that T takes values in $\mathbb{N}_+ \cup \{\infty\}$, since presumably, it is possible that no failure occurs. Let $r(n) = \mathbb{P}(T > n)$ for $n \in \mathbb{N}$, the probability of at least n consecutive successes, starting with a fresh set of trials. Let $f(n) = \mathbb{P}(T = n+1)$ for $n \in \mathbb{N}$, the probability of exactly n consecutive successes, starting with a fresh set of trials.

The functions p , r , and f are related as follows:

1. $p(x) = r(x+1)/r(x)$ for $x \in \mathbb{N}$
2. $r(n) = \prod_{x=0}^{n-1} p(x)$ for $n \in \mathbb{N}$
3. $f(n) = [1 - p(n)] \prod_{x=0}^{n-1} p(x)$ for $n \in \mathbb{N}$
4. $r(n) = 1 - \sum_{x=0}^{n-1} f(x)$ for $n \in \mathbb{N}$
5. $f(n) = r(n) - r(n+1)$ for $n \in \mathbb{N}$

Thus, the functions p , r , and f give equivalent information. If we know one of the functions, we can construct the other two, and hence any of the functions can be used to define the success-runs chain. The function r is the reliability function associated with T .

The function r is characterized by the following properties:

1. r is positive.
2. $r(0) = 1$
3. r is strictly decreasing.

The function f is characterized by the following properties:

1. f is positive.
2. $\sum_{x=0}^{\infty} f(x) \leq 1$

Essentially, f is the probability density function of $T - 1$, except that it may be *defective* in the sense that the sum of its values may be less than 1. The leftover probability, of course, is the probability that $T = \infty$. This is the critical consideration in the classification of the success-runs chain, which we will consider shortly.

Verify that each of the following functions has the appropriate properties, and then find the other two functions:

1. p is a constant in $(0, 1)$.
2. $r(n) = 1/(n+1)$ for $n \in \mathbb{N}$.
3. $r(n) = (n+1)/(2n+1)$ for $n \in \mathbb{N}$.

4. $p(x) = 1/(x+2)$ for $x \in \mathbb{N}$.

Answer

1. $p(x) = p$ for $x \in \mathbb{N}$. $r(n) = p^n$ for $n \in \mathbb{N}$. $f(n) = (1-p)p^n$ for $n \in \mathbb{N}$.
2. $p(x) = \frac{x+1}{x+2}$ for $x \in \mathbb{N}$. $r(n) = \frac{1}{n+1}$ for $n \in \mathbb{N}$. $f(n) = \frac{1}{n+1} - \frac{1}{n}$ for $n \in \mathbb{N}$.
3. $p(x) = \frac{(x+2)(2x+1)}{(x+1)(2x+3)}$ for $x \in \mathbb{N}$. $r(n) = \frac{n+1}{2n+1}$ for $n \in \mathbb{N}$. $f(n) = \frac{n+1}{2n+1} - \frac{n+2}{2n+3}$ for $n \in \mathbb{N}$.
4. $p(x) = \frac{1}{x+2}$ for $x \in \mathbb{N}$. $r(n) = \frac{1}{(n+1)!}$ for $n \in \mathbb{N}$. $f(n) = \frac{1}{(n+1)!} - \frac{1}{(n+2)!}$ for $n \in \mathbb{N}$.

In part (a), note that the trials are Bernoulli trials. We have an app for this case.

The success-runs app is a simulation of the success-runs chain based on Bernoulli trials. Run the simulation 1000 times for various values of p and various initial states, and note the general behavior of the chain.

The success-runs chain is irreducible and aperiodic.

Proof

The chain is irreducible, since 0 leads to every other state, and every state leads back to 0. The chain is aperiodic since $P(0, 0) > 0$.

Recall that T has the same distribution as τ_0 , the first return time to 0 starting at state 0. Thus, the classification of the chain as recurrent or transient depends on $\alpha = \mathbb{P}(T = \infty)$. Specifically, the success-runs chain is transient if $\alpha > 0$ and recurrent if $\alpha = 0$. Thus, we see that the chain is recurrent if and only if a failure is sure to occur. We can compute the parameter α in terms of each of the three functions that define the chain.

In terms of p , r , and f ,

$$\alpha = \prod_{x=0}^{\infty} p(x) = \lim_{n \rightarrow \infty} r(n) = 1 - \sum_{x=0}^{\infty} f(x) \quad (16.10.2)$$

Compute α and determine whether the success-runs chain \mathbf{X} is transient or recurrent for each of the [examples above](#).

Answer

1. $\alpha = 0$, recurrent.
2. $\alpha = 0$, recurrent.
3. $\alpha = \frac{1}{2}$, transient.
4. $\alpha = 0$, recurrent.

Run the simulation of the success-runs chain 1000 times for various values of p , starting in state 0. Note the return times to state 0.

Let $\mu = \mathbb{E}(T)$, the expected trial number of the first failure, starting with a fresh sequence of trials.

μ is related to α , f , and r as follows:

1. If $\alpha > 0$ then $\mu = \infty$
2. If $\alpha = 0$ then $\mu = 1 + \sum_{n=0}^{\infty} n f(n)$
3. $\mu = \sum_{n=0}^{\infty} r(n)$

Proof

1. If $\alpha = \mathbb{P}(T = \infty) > 0$ then $\mu = \mathbb{E}(T) = \infty$.
2. If $\alpha = 0$, so that T takes values in \mathbb{N}_+ , then f is the PDF of $T-1$, so $\mu = 1 + \mathbb{E}(T-1)$.
3. This is a basic result from the general theory of expected value: $\mathbb{E}(T) = \sum_{n=0}^{\infty} \mathbb{P}(T > n)$.

The success-runs chain \mathbf{X} is positive recurrent if and only if $\mu < \infty$.

Proof

Since T is the return time to 0, starting at 0, and since the chain is irreducible, it follows from the general theory that the chain is positive recurrent if and only if $\mu = \mathbb{E}(T) < \infty$.

If \mathbf{X} is recurrent, then r is invariant for \mathbf{X} . In the positive recurrent case, when $\mu < \infty$, the invariant distribution has probability density function g given by

$$g(x) = \frac{r(x)}{\mu}, \quad x \in \mathbb{N} \quad (16.10.3)$$

Proof

If $y \in \mathbb{N}_+$ then from the result [above](#),

$$(rP)(y) = \sum_{x=0}^{\infty} r(x)P(x, y) = r(y-1)p(y-1) = r(y) \quad (16.10.4)$$

For $y = 0$, using the result [above](#) again,

$$(rP)(0) = \sum_{x=0}^{\infty} r(x)P(x, 0) = \sum_{x=0}^{\infty} r(x)[1 - p(x)] = \sum_{x=0}^{\infty} [r(x) - r(x)p(x)] = \sum_{x=0}^{\infty} [r(x) - r(x+1)] \quad (16.10.5)$$

If the chain is recurrent, $r(n) \rightarrow 0$ as $n \rightarrow \infty$ so the last sum collapses to $r(0) = 1$. Recall that $\mu = \sum_{n=0}^{\infty} r(n)$. Hence if $\mu < \infty$, so that the chain is positive recurrent, the function g (which is just r normalized) is the invariant PDF.

When \mathbf{X} is recurrent, we know from the general theory that every other nonnegative left invariant function is a nonnegative multiple of r

Determine whether the success-runs chain \mathbf{X} is transient, null recurrent, or positive recurrent for each of the [examples above](#). If the chain is positive recurrent, find the invariant probability density function.

Answer

1. $\mu = \frac{1}{1-p}$, positive recurrent. $g(x) = (1-p)p^x$ for $x \in \mathbb{N}$.
2. $\alpha = 0$, $\mu = \infty$, null recurrent.
3. $\alpha = \frac{1}{2}$, transient.
4. $\mu = e - 1$, positive recurrent. $g(x) = \frac{1}{(e-1)(x+1)!}$ for $x \in \mathbb{N}$.

From (a), the success-runs chain corresponding to Bernoulli trials with success probability $p \in (0, 1)$ has the geometric distribution on \mathbb{N} , with parameter $1 - p$, as the invariant distribution.

Run the simulation of the success-runs chain 1000 times for various values of p and various initial states. Compare the empirical distribution to the invariant distribution.

The Remaining Life Chain

Consider a device whose (discrete) time to failure U takes values in \mathbb{N} , with probability density function f . We assume that $f(n) > 0$ for $n \in \mathbb{N}$. When the device fails, it is immediately (and independently) replaced by an identical device. For $n \in \mathbb{N}$, let Y_n denote the time to failure of the device that is in service at time n .

$\mathbf{Y} = (Y_0, Y_1, Y_2, \dots)$ is a discrete-time Markov chain with state space \mathbb{N} and transition probability matrix Q given by

$$Q(0, x) = f(x), \quad Q(x+1, x) = 1; \quad x \in \mathbb{N} \quad (16.10.6)$$

The Markov chain \mathbf{Y} is called the *remaining life chain* with lifetime probability density function f , and has the state graph below.




Figure 16.10.2: State graph of the remaining life chain

We have an app for the remaining life chain whose lifetime distribution is the geometric distribution on \mathbb{N} , with parameter $1 - p \in (0, 1)$.

Run the simulation of the remaining-life chain 1000 times for various values of p and various initial states. Note the general behavior of the chain.

If U denotes the lifetime of a device, as before, note that $T = 1 + U$ is the return time to 0 for the chain \mathbf{Y} , starting at 0.

\mathbf{Y} is irreducible, aperiodic, and recurrent.

Proof

From the assumptions on f , state 0 leads to every other state (including 0), and every positive state leads (deterministically) to 0. Thus the chain is irreducible and aperiodic. By assumption, $\mathbb{P}(U \in \mathbb{N}) = 1$ so $\mathbb{P}(T < \infty) = 1$ and hence the chain is recurrent.

Now let $r(n) = \mathbb{P}(U \geq n) = \mathbb{P}(T > n)$ for $n \in \mathbb{N}$ and let $\mu = \mathbb{E}(T) = 1 + \mathbb{E}(U)$. Note that $r(n) = \sum_{x=n}^{\infty} f(x)$ and $\mu = 1 + \sum_{x=0}^{\infty} f(x) = \sum_{n=0}^{\infty} r(n)$.

The success-runs chain \mathbf{X} is positive recurrent if and only if $\mu < \infty$, in which case the invariant distribution has probability density function g given by

$$g(x) = \frac{r(x)}{\mu}, \quad x \in \mathbb{N} \quad (16.10.7)$$

Proof

Since the chain is irreducible, it is positive recurrent if and only if $\mu = E(T) < \infty$. The function r is invariant for Q : for $y \in \mathbb{N}$

$$\begin{aligned} (rQ)(y) &= \sum_{x \in \mathbb{N}} r(x)Q(x, y) = r(0)Q(0, y) + r(y+1)Q(y+1, y) \\ &= f(y) + r(y+1) = r(y) \end{aligned}$$

In the positive recurrent case, μ is the normalizing constant for r , so g is the invariant PDF.

Suppose that \mathbf{Y} is the remaining life chain whose lifetime distribution is the geometric distribution on \mathbb{N} with parameter $1 - p \in (0, 1)$. Then this distribution is also the invariant distribution.

Proof

By assumption, $f(x) = (1-p)p^x$ for $x \in \mathbb{N}$, and the mean of this distribution is $p/(1-p)$. Hence $\mu = 1 + p/(1-p) = 1/(1-p)$, and $r(x) = \sum_{y=x}^{\infty} f(y) = p^x$ for $x \in \mathbb{N}$. Hence $g = f$.

Run the simulation of the success-runs chain 1000 times for various values of p and various initial states. Compare the empirical distribution to the invariant distribution.

Time Reversal

You probably have already noticed similarities, in notation and results, between the success-runs chain and the remaining-life chain. There are deeper connections.

Suppose that f is a probability density function on \mathbb{N} with $f(n) > 0$ for $n \in \mathbb{N}$. Let \mathbf{X} be the success-runs chain associated with f and \mathbf{Y} the remaining life chain associated with f . Then \mathbf{X} and \mathbf{Y} are time reversals of each other.

Proof

Under the assumptions on f , both chains are recurrent and irreducible. Hence it suffices to show that

$$r(x)P(x, y) = r(y)Q(y, x), \quad x, y \in \mathbb{N} \quad (16.10.8)$$

It will then follow that the chains are time reversals of each other, and that r is a common invariant function (unique up to multiplication by positive constants). In the case that $\mu = \sum_{n=0}^{\infty} r(n) < \infty$, the function $g = r/\mu$ is the common invariant PDF. There are only two cases to consider. With $y = 0$, we have $r(x)P(x, 0) = r(x)[1 - p(x)]$ and $r(0)Q(y, 0) = f(x)$. But $r(x)[1 - p(x)] = f(x)$ by the result [above](#). When $x \in \mathbb{N}$ and $y = x + 1$, we have $r(x)P(x, x + 1) = r(x)p(x)$ and $r(x + 1)Q(x + 1, x) = r(x + 1)$. But $r(x)p(x) = r(x + 1)$ by the result [above](#).

In the context of reliability, it is also easy to see that the chains are time reversals of each other. Consider again a device whose random lifetime takes values in \mathbb{N} , with the device immediately replaced by an identical device upon failure. For $n \in \mathbb{N}$, we can think of X_n as the age of the device in service at time n and Y_n as the time remaining until failure for that device.

Run the simulation of the success-runs chain 1000 times for various values of p , starting in state 0. This is the time reversal of the simulation in the next exercise

Run the simulation of the remaining-life chain 1000 times for various values of p , starting in state 0. This is the time reversal of the simulation in the previous exercise.

This page titled [16.10: Discrete-Time Reliability Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.11: Discrete-Time Branching Chain

Basic Theory

Introduction

Generically, suppose that we have a system of *particles* that can generate or split into other particles of the same type. Here are some typical examples:

- The particles are biological organisms that reproduce.
- The particles are neutrons in a chain reaction.
- The particles are electrons in an electron multiplier.

We assume that each particle, at the end of its life, is replaced by a random number of new particles that we will refer to as *children* of the original particle. Our basic assumption is that the particles act independently, each with the same offspring distribution on \mathbb{N} . Let f denote the common probability density function of the number of offspring of a particle. We will also let $f^{*n} = f * f * \dots * f$ denote the convolution power of degree n of f ; this is the probability density function of the total number of children of n particles.

We will consider the evolution of the system in *real* time in our study of continuous-time branching chains. In this section, we will study the evolution of the system in *generational* time. Specifically, the particles that we start with are in generation 0, and recursively, the children of a particle in generation n are in generation $n + 1$.




Figure 16.11.1: Generations 0, 1, 2, and 3 of a branching chain.

Let X_n denote the number of particles in generation $n \in \mathbb{N}$. One way to construct the process mathematically is to start with an array of independent random variables $(U_{n,i} : n \in \mathbb{N}, i \in \mathbb{N}_+)$, each with probability density function f . We interpret $U_{n,i}$ as the number of children of the i th particle in generation n (if this particle exists). Note that we have more random variables than we need, but this causes no harm, and we know that we can construct a probability space that supports such an array of random variables. We can now define our state variables recursively by

$$X_{n+1} = \sum_{i=1}^{X_n} U_{n,i} \quad (16.11.1)$$

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a discrete-time Markov chain on \mathbb{N} with transition probability matrix P given by

$$P(x, y) = f^{*x}(y), \quad (x, y) \in \mathbb{N}^2 \quad (16.11.2)$$

The chain \mathbf{X} is the *branching chain* with offspring distribution defined by f .

Proof

The Markov property and the form of the transition matrix follow directly from the construction of the state variables given above. Since the variables $(U_{n,i} : n \in \mathbb{N}, i \in \mathbb{N}_+)$ are independent, each with PDF f , we have

$$\mathbb{P}(X_{n+1} = y \mid X_0 = x_0, \dots, X_{n-1} = x_{n-1}, X_n = x) = \mathbb{P}\left(\sum_{i=1}^x U_{n,i} = y\right) = f^{*x}(y) \quad (16.11.3)$$

The branching chain is also known as the *Galton-Watson process* in honor of Francis Galton and Henry William Watson who studied such processes in the context of the survival of (aristocratic) family names. Note that the descendants of each initial particle form a branching chain, and these chains are independent. Thus, the branching chain starting with x particles is equivalent to x independent copies of the branching chain starting with 1 particle. This features turns out to be very important in the analysis of the chain. Note also that 0 is an absorbing state that corresponds to *extinction*. On the other hand, the population may grow to infinity, sometimes called *explosion*. Computing the probability of extinction is one of the fundamental problems in branching chains; we will essentially solve this problem in the next subsection.

Extinction and Explosion

The behavior of the branching chain in expected value is easy to analyze. Let m denote the mean of the offspring distribution, so that

$$m = \sum_{x=0}^{\infty} x f(x) \quad (16.11.4)$$

Note that $m \in [0, \infty]$. The parameter m will turn out to be of fundamental importance.

Expected value properties

1. $\mathbb{E}(X_{n+1}) = m\mathbb{E}(X_n)$ for $n \in \mathbb{N}$
2. $\mathbb{E}(X_n) = m^n \mathbb{E}(X_0)$ for $n \in \mathbb{N}$
3. $\mathbb{E}(X_n) \rightarrow 0$ as $n \rightarrow \infty$ if $m < 1$.
4. $\mathbb{E}(X_n) = \mathbb{E}(X_0)$ for each $n \in \mathbb{N}$ if $m = 1$.
5. $\mathbb{E}(X_n) \rightarrow \infty$ as $n \rightarrow \infty$ if $m > 1$ and $\mathbb{E}(X_0) > 0$.

Proof

For part (a) we use a conditioning argument and the construction above. For $x \in \mathbb{N}$,

$$\mathbb{E}(X_{n+1} \mid X_n = x) = \mathbb{E}\left(\sum_{i=1}^x U_{n,i} \mid X_n = x\right) = \mathbb{E}\left(\sum_{i=1}^x U_{n,i}\right) = mx \quad (16.11.5)$$

That is, $\mathbb{E}(X_{n+1} \mid X_n) = mX_n$ so $\mathbb{E}(X_{n+1}) = \mathbb{E}[\mathbb{E}(X_{n+1} \mid X_n)] = m\mathbb{E}(X_n)$ Part (b) follows from (a) and then parts (c), (d), and (e) follow from (b).

Part (c) is *extinction in the mean*; part (d) is *stability in the mean*; and part (e) is *explosion in the mean*.

Recall that state 0 is absorbing (there are no particles), and hence $\{X_n = 0 \text{ for some } n \in \mathbb{N}\} = \{\tau_0 < \infty\}$ is the *extinction event* (where as usual, τ_0 is the time of the first return to 0). We are primarily concerned with the probability of extinction, as a function of the initial state. First, however, we will make some simple observations and eliminate some trivial cases.

Suppose that $f(1) = 1$, so that each particle is replaced by a single new particle. Then

1. Every state is absorbing.
2. The equivalence classes are the singleton sets.
3. With probability 1, $X_n = X_0$ for every $n \in \mathbb{N}$.

Proof

These properties are obvious since $P(x, x) = 1$ for every $x \in \mathbb{N}$.

Suppose that $f(0) > 0$ so that with positive probability, a particle will die without offspring. Then

1. Every state leads to 0.
2. Every positive state is transient.
3. With probability 1 either $X_n = 0$ for some $n \in \mathbb{N}$ (extinction) or $X_n \rightarrow \infty$ as $n \rightarrow \infty$ (explosion).

Proof

1. Note that $P(x, 0) = [f(0)]^x > 0$ for $x \in \mathbb{N}$, so every state leads to 0 in one step.

2. This follows from (a). If $x \in \mathbb{N}_+$, then x leads to the absorbing state 0 with positive probability. Hence a return to x , starting in x , cannot have probability 1.
3. This follows from (a) and (b). With probability 1, every positive state is visited only finitely many times. Hence the only possibilities are $X_n = 0$ for some $n \in \mathbb{N}$ or $X_n \rightarrow \infty$ as $n \rightarrow \infty$.

Suppose that $f(0) = 0$ and $f(1) < 1$, so that every particle is replaced by at least one particle, and with positive probability, more than one. Then

1. Every positive state is transient.
2. $\mathbb{P}(X_n \rightarrow \infty \text{ as } n \rightarrow \infty \mid X_0 = x) = 1$ for every $x \in \mathbb{N}_+$, so that explosion is certain, starting with at least one particle.

Proof

1. Let $x \in \mathbb{N}_+$. Under the assumptions on f , state x leads to some state $y > x$ but y does not lead back to x . Hence with positive probability, the chain starting in x will not return to x .
2. This follows from (a) and that the fact that positive states do not lead to 0.

Suppose that $f(0) > 0$ and $f(0) + f(1) = 1$, so that with positive probability, a particle will die without offspring, and with probability 1, a particle is not replaced by more than one particle. Then

1. Every state leads to 0.
2. Every positive state is transient.
3. With probability 1, $X_n = 0$ for some $n \in \mathbb{N}$, so extinction is certain.

Proof

1. As before, $P(x, 0) = [f(0)]^x > 0$ for $x \in \mathbb{N}$, so x leads to 0 in one step.
2. This follows from (a) and the fact that 0 is absorbing.
3. Under the assumptions on f , state x leads to state y only if $y \leq x$. So this follows from (a) and (b).

Thus, the interesting case is when $f(0) > 0$ and $f(0) + f(1) < 1$, so that with positive probability, a particle will die without offspring, and also with positive probability, the particle will be replaced by more than one new particles. We will assume these conditions for the remainder of our discussion. By the [state classification above](#) all states lead to 0 (extinction). We will denote the probability of extinction, starting with one particle, by

$$q = \mathbb{P}(\tau_0 < \infty \mid X_0 = 1) = \mathbb{P}(X_n = 0 \text{ for some } n \in \mathbb{N} \mid X_0 = 1) \quad (16.11.6)$$

The set of positive states \mathbb{N}_+ is a transient equivalence class, and the probability of extinction starting with $x \in \mathbb{N}$ particles is

$$q^x = \mathbb{P}(\tau_0 < \infty \mid X_0 = x) = \mathbb{P}(X_n = 0 \text{ for some } n \in \mathbb{N} \mid X_0 = x) \quad (16.11.7)$$

Proof

Under the assumptions on f , from any positive state the chain can move 2 or more units to the right and one unit to the left in one step. It follows that every positive state leads to every other positive state. On the other hand, every positive state leads to 0, which is absorbing. Thus, \mathbb{N}_+ is a transient equivalence class.

Recall that the branching chain starting with $x \in \mathbb{N}_+$ particles acts like x independent branching chains starting with one particle. Thus, the extinction probability starting with x particles is q^x .

The parameter q satisfies the equation

$$q = \sum_{x=0}^{\infty} f(x)q^x \quad (16.11.8)$$

Proof

This result follows from conditioning on the first state.

$$q = \mathbb{P}(\tau_0 < \infty \mid X_0 = 1) = \sum_{x=0}^{\infty} \mathbb{P}(\tau_0 < \infty \mid X_0 = 1, X_1 = x) \mathbb{P}(X_1 = x \mid X_0 = 1) \quad (16.11.9)$$

But by the Markov property and the previous result,

$$\mathbb{P}(\tau_0 < \infty \mid X_0 = 1, X_1 = x) = \mathbb{P}(\tau_0 < \infty \mid X_1 = x) = q^x \quad (16.11.10)$$

and of course $\mathbb{P}(X_1 = x \mid X_0 = 1) = P(1, x) = f(x)$.

Thus the extinction probability q starting with 1 particle is a *fixed point* of the probability generating function Φ of the offspring distribution:

$$\Phi(t) = \sum_{x=0}^{\infty} f(x)t^x, \quad t \in [0, 1] \quad (16.11.11)$$

Moreover, from the general discussion of hitting probabilities in the section on recurrence and transience, q is the smallest such number in the interval $(0, 1]$. If the probability generating function Φ can be computed in closed form, then q can sometimes be computed by solving the equation $\Phi(t) = t$.

Φ satisfies the following properties:

1. $\Phi(0) = f(0)$.
2. $\Phi(1) = 1$.
3. $\Phi'(t) > 0$ for $t \in (0, 1)$ so Φ is increasing on $(0, 1)$.
4. $\Phi''(t) > 0$ for $t \in (0, 1)$ so Φ is concave upward on $(0, 1)$.
5. $m = \lim_{t \uparrow 1} \Phi'(t)$.

Proof

These are basic properties of the probability generating function. Recall that the series that defines Φ is a power series about 0 with radius of convergence $r \geq 1$. A function defined by a power series is infinitely differentiable within the open interval of convergence, and the derivatives can be computed term by term. So

$$\begin{aligned} \Phi'(t) &= \sum_{x=1}^{\infty} x f(x) t^{x-1} > 0, \quad t \in (0, 1) \\ \Phi''(t) &= \sum_{x=2}^{\infty} x(x-1) f(x) t^{x-2} > 0, \quad t \in (0, 1) \end{aligned}$$

If $r > 1$ then $m = \Phi'(1)$. If $r = 1$, the limit result is the best we can do.

Our main result is next, and relates the extinction probability q and the mean of the offspring distribution m .

The extinction probability q and the mean of the offspring distribution m are related as follows:

1. If $m \leq 1$ then $q = 1$, so extinction is certain.
2. If $m > 1$ then $0 < q < 1$, so there is a positive probability of extinction and a positive probability of explosion.

Proof




Figure 16.11.2 The case of certain extinction.




Figure 16.11.3 The case of possible extinction and possible explosion.

Computational Exercises

Consider the branching chain with offspring probability density function f given by $f(0) = 1 - p$, $f(2) = p$, where $p \in (0, 1)$ is a parameter. Thus, each particle either dies or splits into two new particles. Find each of the following.

1. The transition matrix P .
2. The mean m of the offspring distribution.
3. The generating function Φ of the offspring distribution.
4. The extinction probability q .

Answer

Note that an offspring variable has the form $2I$ where I is an indicator variable with parameter p .

1. For $x \in \mathbb{N}$, f^{*x} is the PDF of $2U$ where U has the binomial distribution with parameters x and p . Hence

$$P(x, y) = f^{*x}(y) = \binom{x}{y/2} p^{y/2} (1-p)^{x-y/2}, \quad y \in \{0, 2, \dots, 2x\} \quad (16.11.12)$$

2. $m = 2p$.
3. $\Phi(t) = pt^2 + (1-p)$ for $t \in \mathbb{R}$.
4. $q = 1$ if $0 < p \leq \frac{1}{2}$ and $q = \frac{1-p}{p}$ if $\frac{1}{2} < p < 1$.

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{1}{3}$

Graphs

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{2}{3}$

Graphs

Consider the branching chain whose offspring distribution is the geometric distribution on \mathbb{N} with parameter $1 - p$, where $p \in (0, 1)$. Thus $f(n) = (1-p)p^n$ for $n \in \mathbb{N}$. Find each of the following:

1. The transition matrix P .
2. The mean m of the offspring distribution.
3. The generating function Φ of the offspring distribution.

4. The extinction probability q .

Answer

1. For $x \in \mathbb{N}$, f^{*x} is the PDF of the negative binomial distribution on \mathbb{N} with parameter $1 - p$. So

$$P(x, y) = f^{*x}(y) = \binom{x+y-1}{x-1} p^y (1-p)^x, \quad y \in \mathbb{N} \quad (16.11.13)$$

2. $m = \frac{p}{1-p}$.

3. $\Phi(t) = \frac{1-p}{1-pt}$ for $|t| < \frac{1}{p}$.

4. $q = 1$ if $0 < p \leq \frac{1}{2}$ and $q = \frac{1-p}{p}$ if $\frac{1}{2} < p < 1$.

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{1}{3}$

 Graphs

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{2}{3}$

 Graphs

Curiously, the extinction probability is the same as for the previous problem.

Consider the branching chain whose offspring distribution is the Poisson distribution with parameter $m \in (0, \infty)$. Thus $f(n) = e^{-m} m^n / n!$ for $n \in \mathbb{N}$. Find each of the following:

1. The transition matrix P .
2. The mean m of the offspring distribution.
3. The generating function Φ of the offspring distribution.
4. The approximate extinction probability q when $m = 2$ and when $m = 3$.

Answer

1. For $x \in \mathbb{N}$, f^{*x} is the PDF of the Poisson distribution with parameter mx . So

$$P(x, y) = f^{*x}(y) = e^{-mx} \frac{(mx)^y}{y!}, \quad y \in \mathbb{N} \quad (16.11.14)$$

2. The parameter m is the mean of the Poisson distribution, so the notation is consistent.

3. $\Phi(t) = e^{m(t-1)}$ for $t \in \mathbb{R}$.

4. $q = 1$ if $0 < m \leq 1$. If $m > 1$ then q is the solution in $(0, 1)$ of the equation $e^{m(q-1)} = q$ which can be expressed in terms of a special function known as the *Lambert W function*:

$$q = -\frac{1}{m} W(-me^{-m}) \quad (16.11.15)$$

For $m = 2$, $q \approx 0.20319$. For $m = 3$, $q \approx 0.059520$

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $m = \frac{1}{2}$

 Graphs

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $m = 2$

 Graphs

This page titled [16.11: Discrete-Time Branching Chain](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.12: Discrete-Time Queuing Chains

Basic Theory

Introduction

In a *queuing model*, customers arrive at a station for service. As always, the terms are generic; here are some typical examples:

- The customers are persons and the service station is a store.
- The customers are file requests and the service station is a web server.
- The customers are packages and the service station is a processing facility.



Figure 16.12.1: Ten customers and a server

Queuing models can be quite complex, depending on such factors as the probability distribution that governs the arrival of customers, the probability distribution that governs the service of customers, the number of servers, and the behavior of the customers when all servers are busy. Indeed, queuing theory has its own lexicon to indicate some of these factors. In this section, we will study one of the simplest, discrete-time queuing models. However, as we will see, this discrete-time chain is embedded in a much more realistic continuous-time queuing process known as the *M/G/1 queue*. In a general sense, the main interest in any queuing model is the number of customers in the system as a function of time, and in particular, whether the servers can adequately handle the flow of customers.

Our main assumptions are as follows:

1. If the queue is empty at a given time, then a random number of new customers arrive at the next time.
2. If the queue is nonempty at a given time, then one customer is served and a random number of new customers arrive at the next time.
3. The number of customers who arrive at each time period form an independent, identically distributed sequence.

Thus, let X_n denote the number of customers in the system at time $n \in \mathbb{N}$, and let U_n denote the number of new customers who arrive at time $n \in \mathbb{N}_+$. Then $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent random variables, with common probability density function f on \mathbb{N} , and

$$X_{n+1} = \begin{cases} U_{n+1}, & X_n = 0 \\ (X_n - 1) + U_{n+1}, & X_n > 0 \end{cases}, \quad n \in \mathbb{N} \quad (16.12.1)$$

$\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a discrete-time Markov chain with state space \mathbb{N} and transition probability matrix P given by

$$P(0, y) = f(y), \quad y \in \mathbb{N} \quad (16.12.2)$$

$$P(x, y) = f(y - x + 1), \quad x \in \mathbb{N}_+, y \in \{x - 1, x, x + 1, \dots\} \quad (16.12.3)$$

The chain \mathbf{X} is the *queuing chain* with arrival distribution defined by f .

Proof

The Markov property and the form of the transition matrix follow from the construction of the state process \mathbf{X} in terms of the IID sequence \mathbf{U} . Starting in state 0 (an empty queue), a random number of new customers arrive at the next time unit, governed by the PDF f . Hence the probability of going from state 0 to state y in one step is $f(y)$. Starting in state $x \in \mathbb{N}_+$, one customer is served and a random number of new customers arrive by the next time unit, again governed by the PDF f . Hence the probability of going from state x to state $y \in \{x - 1, x, x + 1, \dots\}$ is $f[y - (x - 1)]$.

Recurrence and Transience

From now on we will assume that $f(0) > 0$ and $f(0) + f(1) < 1$. Thus, at each time unit, it's possible that no new customers arrive or that at least 2 new customers arrive. Also, we let m denote the mean of the arrival distribution, so that

$$m = \sum_{x=0}^{\infty} x f(x) \quad (16.12.4)$$

Thus m is the average number of new customers who arrive during a time period.

The chain \mathbf{X} is irreducible and aperiodic.

Proof

In a positive state, the chain can move at least one unit to the right and can move one unit to the left at the next step. From state 0, the chain can move two or more units to the right or can stay in 0 at the next step. Thus, every state leads to every other state so the chain is irreducible. Since 0 leads back to 0, the chain is aperiodic.

Our goal in this section is to compute the probability that the chain reaches 0, as a function of the initial state (so that the server is able to serve all of the customers). As we will see, there are some curious and unexpected parallels between this problem and the problem of computing the extinction probability in the branching chain. As a corollary, we will also be able to classify the queuing chain as transient or recurrent. Our basic parameter of interest is $q = H(1, 0) = \mathbb{P}(\tau_0 < \infty \mid X_0 = 1)$, where as usual, H is the hitting probability matrix and $\tau_0 = \min\{n \in \mathbb{N}_+ : X_n = 0\}$ is the first positive time that the chain is in state 0 (possibly infinite). Thus, q is the probability that the queue eventually empties, starting with a single customer.

The parameter q satisfies the following properties:

1. $q = H(x, x-1)$ for every $x \in \mathbb{N}_+$.
2. $q^x = H(x, 0)$ for every $x \in \mathbb{N}_+$.

Proof

1. The critical observation is that if $x \in \mathbb{N}_+$ then $P(x, y) = P(1, y-x+1) = f(y-x+1)$ for $y \in \{x-1, x, x+1, \dots\}$. Thus, the chain, starting in x , and up until the time that it reaches $x-1$ (if it does), behaves stochastically like the chain starting in state 1, and up until it reaches 0.
2. In order to reach 0, starting in state $x \in \mathbb{N}_+$, the chain must first reach $x-1$ and then from $x-1$ must reach $x-2$, until finally reaching 0 from state 1. Each of these intermediate trips has probability q by part (a) and are independent by the Markov property.

The parameter q satisfies the equation:

$$q = \sum_{x=0}^{\infty} f(x) q^x \quad (16.12.5)$$

Proof

This follows from the previous theorem by conditioning on the first state.

$$\mathbb{P}(\tau_0 < \infty \mid X_0 = 1) = \sum_{x=0}^{\infty} \mathbb{P}(\tau_0 < \infty \mid X_0 = 1, X_1 = x) \mathbb{P}(X_1 = x \mid X_0 = 1) \quad (16.12.6)$$

Note first that $\mathbb{P}(\tau_0 < \infty \mid X_0 = 1, X_1 = 0) = 1 = q^0$. On the other hand, by the Markov property and the previous result,

$$\mathbb{P}(\tau_0 < \infty \mid X_0 = 1, X_1 = x) = \mathbb{P}(\tau_0 < \infty \mid X_1 = x) = q^x, \quad x \in \mathbb{N}_+ \quad (16.12.7)$$

Of course $\mathbb{P}(X_1 = x \mid X_0 = 1) = P(1, x) = f(x)$ for $x \in \mathbb{N}$.

Note that this is exactly the same equation that we considered for the branching chain, namely $\Phi(q) = q$, where Φ is the probability generating function of the distribution that governs the number of new customers that arrive during each period.




Figure 16.12.2 The graph of ϕ in the recurrent case




Figure 16.12.3 The graph of ϕ in the transient case

q is the smallest solution in $(0, 1]$ of the equation $\Phi(t) = t$. Moreover

1. If $m \leq 1$ then $q = 1$ and the chain is recurrent.
2. If $m > 1$ then $0 < q < 1$ and the chain is transient..

Proof

This follows from our analysis of branching chains. The graphs above show the two cases. Note that the condition in (a) means that on average, one or fewer new customers arrive for each customer served. The condition in (b) means that on average, more than one new customer arrives for each customer served.

Positive Recurrence

Our next goal is to find conditions for the queuing chain to be positive recurrent. Recall that m is the mean of the probability density function f ; that is, the expected number of new customers who arrive during a time period. As before, let τ_0 denote the first positive time that the chain is in state 0. We assume that the chain is recurrent, so $m \leq 1$ and $\mathbb{P}(\tau_0 < \infty) = 1$.

Let Ψ denote the probability generating function of τ_0 , starting in state 1. Then

1. Ψ is also the probability generating function of τ_0 starting in state 0.
2. Ψ^x is the probability generating function of τ_0 starting in state $x \in \mathbb{N}_+$.

Proof

1. The transition probabilities starting in state 1 are the same as those starting in state 0: $P(0, x) = P(1, x) = f(x)$ for $x \in \mathbb{N}$.
2. Starting in state $x \in \mathbb{N}_+$, the random time to reach 0 is the sum of the time to reach $x - 1$, the additional time to reach $x - 2$ from $x - 1$, and so forth, ending with the time to reach 0 from state 1. These random times are independent by the Markov property, and each has the same distribution as the time to reach 0 from state 1 by our argument above. Finally, recall that the PGF of a sum of independent variables is the product of the corresponding PGFs.

$$\Psi(t) = t\Phi[\Psi(t)] \text{ for } t \in [-1, 1].$$

Proof

Once again, the trick is to condition on the first state:

$$\Psi(t) = \mathbb{E}(t^{\tau_0} \mid X_0 = 1) = \sum_{x=0}^{\infty} \mathbb{E}(t^{\tau_0} \mid X_0 = 1, X_1 = x) \mathbb{P}(X_1 = x \mid X_0 = 1) \quad (16.12.8)$$

First note that $\mathbb{E}(t^{\tau_0} \mid X_0 = 1, X_1 = 0) = t^1 = t\Psi^0(t)$. On the other hand, by the Markov property and the [previous theorem](#),

$$\mathbb{E}(t^{\tau_0} \mid X_0 = 1, X_1 = x) = \mathbb{E}(t^{1+\tau_0} \mid X_0 = x) = t\mathbb{E}(t^{\tau_0} \mid X_0 = x) = t\Psi^x(t), \quad x \in \mathbb{N}_+ \quad (16.12.9)$$

Of course $\mathbb{P}(X_1 = x \mid X_0 = 1) = P(1, x) = f(x)$. Hence we have

$$\Psi(t) = \sum_{x=0}^{\infty} t\Psi^x(t)f(x) = t\Phi[\Psi(t)] \quad (16.12.10)$$

The PGF of any variable that takes positive integer values is defined on $[-1, 1]$, and maps this interval back into itself. Hence the representation is valid at least for $t \in [-1, 1]$.

The derivative of Ψ is

$$\Psi'(t) = \frac{\Phi[\Psi(t)]}{1 - t\Phi'[\Psi(t)]}, \quad t \in (-1, 1) \quad (16.12.11)$$

Proof

Recall that a PGF is infinitely differentiable on the open interval of convergence. Hence using the result in the [previous theorem](#) and the product and chain rules,

$$\Psi'(t) = \Phi[\Psi(t)] + t\Phi'[\Psi(t)]\Psi'(t) \quad (16.12.12)$$

Solving for $\Psi'(t)$ gives the result.

As usual, let $\mu_0 = \mathbb{E}(\tau_0 \mid X_0 = 0)$, the mean return time to state 0 starting in state 0. Then

1. $\mu_0 = \frac{1}{1-m}$ if $m < 1$ and therefore the chain is positive recurrent.
2. $\mu_0 = \infty$ if $m = 1$ and therefore the chain is null recurrent.

Proof

Recall that Ψ is the probability generating function of τ_0 , starting at 0. From basic properties of PGFs we know that $\Phi(t) \uparrow 1$, $\Psi(t) \uparrow 1$, $\Phi'(t) \uparrow m$, and $\Psi'(t) \uparrow \mu_0$ as $t \uparrow 1$. So letting $t \uparrow 1$ in the result of the previous theorem, we have $\mu_0 = 1/(1-m)$ if $m < 1$ and $\mu_0 = \infty$ if $m = 1$.

So to summarize, the queuing chain is positive recurrent if $m < 1$, null recurrent if $m = 1$, and transient if $m > 1$. Since m is the expected number of new customers who arrive during a service period, the results are certainly reasonable.

Computational Exercises

Consider the queuing chain with arrival probability density function f given by $f(0) = 1 - p$, $f(2) = p$, where $p \in (0, 1)$ is a parameter. Thus, at each time period, either no new customers arrive or two arrive.

1. Find the transition matrix P .
2. Find the mean m of the arrival distribution.
3. Find the generating function Φ of the arrival distribution.
4. Find the probability q that the queue eventually empties, starting with one customer.
5. Classify the chain as transient, null recurrent, or positive recurrent.
6. In the positive recurrent case, find μ_0 , the mean return time to 0.

Answer

1. $P(0, 0) = 1 - p$, $P(0, 2) = p$. For $x \in \mathbb{N}_+$, $P(x, x-1) = 1 - p$, $P(x, x+1) = p$.
2. $m = 2p$.

3. $\Phi(t) = pt^2 + (1-p)$ for $t \in \mathbb{R}$.
4. $q = 1$ if $0 < p \leq \frac{1}{2}$ and $q = \frac{1-p}{p}$ if $\frac{1}{2} < p < 1$.
5. The chain is transient if $p > \frac{1}{2}$, null recurrent if $p = \frac{1}{2}$, and positive recurrent if $p < \frac{1}{2}$.
6. $\mu_0 = \frac{1}{1-2p}$ for $p < \frac{1}{2}$.

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{1}{3}$

 Graphs

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{2}{3}$

 Graphs

Consider the queuing chain whose arrival distribution is the geometric distribution on \mathbb{N} with parameter $1-p$, where $p \in (0, 1)$. Thus $f(n) = (1-p)p^n$ for $n \in \mathbb{N}$.

1. Find the transition matrix P .
2. Find the mean m of the arrival distribution.
3. Find the generating function Φ of the arrival distribution.
4. Find the probability q that the queue eventually empties, starting with one customer.
5. Classify the chain as transient, null recurrent, or positive recurrent.
6. In the positive recurrent case, find μ_0 , the mean return time to 0.

Answer

1. $P(0, y) = (1-p)p^y$ for $y \in \mathbb{N}$. For $x \in \mathbb{N}_+$, $P(x, y) = (1-p)p^{y-x+1}$ for $y \in \{x-1, x, x+1, \dots\}$.
2. $m = \frac{p}{1-p}$.
3. $\Phi(t) = \frac{1-p}{1-pt}$ for $|t| < \frac{1}{p}$.
4. $q = 1$ if $0 < p \leq \frac{1}{2}$ and $q = \frac{1-p}{p}$ if $\frac{1}{2} < p < 1$.
5. The chain is transient if $p > \frac{1}{2}$, null recurrent if $p = \frac{1}{2}$, and positive recurrent if $p < \frac{1}{2}$.
6. $\mu_0 = \frac{1-p}{1-2p}$ for $p < \frac{1}{2}$.

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{1}{3}$

 Graphs

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $p = \frac{2}{3}$

 Graphs

Curiously, the parameter q and the classification of the chain are the same in the last two models.

Consider the queuing chain whose arrival distribution is the Poisson distribution with parameter $m \in (0, \infty)$. Thus $f(n) = e^{-m} m^n / n!$ for $n \in \mathbb{N}$. Find each of the following:

1. The transition matrix P
2. The mean m of the arrival distribution.
3. The generating function Φ of the arrival distribution.
4. The approximate value of q when $m = 2$ and when $m = 3$.
5. Classify the chain as transient, null recurrent, or positive recurrent.
6. In the positive recurrent case, find μ_0 , the mean return time to 0.

Answer

1. $P(0, y) = e^{-m} m^y / y!$ for $y \in \mathbb{N}$. For $x \in \mathbb{N}_+$, $P(x, y) = e^{-m} m^{y-x+1} / (y-x+1)!$ for $y \in \{x-1, x, x+1, \dots\}$.
2. The parameter m is the mean of the Poisson distribution, so the notation is consistent.
3. $\Phi(t) = e^{m(t-1)}$ for $t \in \mathbb{R}$.
4. $q = 1$ if $0 < m \leq 1$. If $m > 1$ then q is the solution in $(0, 1)$ of the equation $e^{m(q-1)} = q$ which can be expressed in terms of a special function known as the *Lambert W function*:

$$q = -\frac{1}{m} W(-me^{-m}) \quad (16.12.13)$$

For $m = 2$, $q \approx 0.20319$ For $m = 3$, $q \approx 0.059520$

5. The chain is transient if $m > 1$, null recurrent if $m = 1$, and positive recurrent if $m < 1$.
6. $\mu_0 = \frac{1}{1-m}$ for $m < 1$.

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $m = \frac{1}{2}$

 Graphs

Graphs of $t \mapsto \Phi(t)$ and $t \mapsto t$ when $m = 2$

 Graphs

This page titled [16.12: Discrete-Time Queuing Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.13: Discrete-Time Birth-Death Chains

Basic Theory

Introduction

Suppose that S is an interval of integers (that is, a set of consecutive integers), either finite or infinite. A (discrete-time) *birth-death* chain on S is a discrete-time Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ on S with transition probability matrix P of the form

$$P(x, x-1) = q(x), \quad P(x, x) = r(x), \quad P(x, x+1) = p(x); \quad x \in S \quad (16.13.1)$$

where p , q , and r are nonnegative functions on S with $p(x) + q(x) + r(x) = 1$ for $x \in S$.

If the interval S has a minimum value $a \in \mathbb{Z}$ then of course we must have $q(a) = 0$. If $r(a) = 1$, the boundary point a is *absorbing* and if $p(a) = 1$, then a is *reflecting*. Similarly, if the interval S has a maximum value $b \in \mathbb{Z}$ then of course we must have $p(b) = 0$. If $r(b) = 1$, the boundary point b is *absorbing* and if $q(b) = 1$, then b is *reflecting*. Several other special models that we have studied are birth-death chains; these are explored in [below](#).

In this section, as you will see, we often have sums of products. Recall that a sum over an empty index set is 0, while a product over an empty index set is 1.

Recurrence and Transience

If S is finite, classification of the states of a birth-death chain as recurrent or transient is simple, and depends only on the state graph. In particular, if the chain is irreducible, then the chain is positive recurrent. So we will study the classification of birth-death chains when $S = \mathbb{N}$. We assume that $p(x) > 0$ for all $x \in \mathbb{N}$ and that $q(x) > 0$ for all $x \in \mathbb{N}_+$ (but of course we must have $q(0) = 0$). Thus, the chain is irreducible.

Under these assumptions, the birth-death chain on \mathbb{N} is

1. Aperiodic if $r(x) > 0$ for some $x \in \mathbb{N}$.
2. Periodic with period 2 if $r(x) = 0$ for all $x \in \mathbb{N}$.

Proof

1. If $r(x) > 0$ for some $x \in \mathbb{N}$ then $P(x, x) > 0$ and hence the chain is aperiodic.
2. If $r(x) = 0$ for every $x \in \mathbb{N}$ then clearly the chain starting in x can be in state x again only at even times.

We will use the test for recurrence derived earlier with $A = \mathbb{N}_+$, the set of positive states. That is, we will compute the probability that the chain never hits 0, starting in a positive state.

The chain \mathbf{X} is recurrent if and only if

$$\sum_{x=0}^{\infty} \frac{q(1) \cdots q(x)}{p(1) \cdots p(x)} = \infty \quad (16.13.2)$$

Proof

Let P_+ denote the restriction of P to $\mathbb{N}_+ \times \mathbb{N}_+$, and define $u_+ : \mathbb{N}_+ \rightarrow [0, 1]$ by

$$u_+(x) = \mathbb{P}(X_1 > 0, X_2 > 0, \dots \mid X_0 = x), \quad x \in \mathbb{N}_+ \quad (16.13.3)$$

So $u_+(x)$ is the probability that chain never reaches 0, starting in $x \in \mathbb{N}_+$. From our general theory, we know that u_+ satisfies $u_+ = P_+ u_+$ and is the largest such function with values in $[0, 1]$. Furthermore, we know that either $u_+(x) = 0$ for all $x \in \mathbb{N}_+$ or that $\sup\{u_+(x) : x \in [0, 1]\} = 1$. In the first case the chain is recurrent, and in the second case the chain is transient.

The functional equation $P_+ u = u$ for a function $u : \mathbb{N}_+ \rightarrow [0, 1]$ is equivalent to the following system of equations:

$$u(2) - u(1) = \frac{q(1)}{p(1)}u(1) \quad (16.13.4)$$

$$u(x+1) - u(x) = \frac{q(x)}{p(x)}[u(x) - u(x-1)], \quad x \in \{2, 3, \dots\} \quad (16.13.5)$$

Solving this system of equations for the differences gives

$$u(x+1) - u(x) = \frac{q(1) \cdots q(x)}{p(1) \cdots p(x)}u(1), \quad x \in \mathbb{N}_+ \quad (16.13.6)$$

Solving this new systems gives

$$u(x) = u(1) \sum_{i=0}^{x-1} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)}, \quad x \in \mathbb{N}_+ \quad (16.13.7)$$

Note that $u(x)$ is increasing in $x \in \mathbb{N}_+$ and so has a limit as $x \rightarrow \infty$. Let $A = \sum_{i=0}^{\infty} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)}$.

1. Suppose that $A = \infty$. Letting $x \rightarrow \infty$ in the displayed equation above for $u(x)$ shows that $u(1) = 0$ and so $u(x) = 0$ for all x . Hence the chain is recurrent.
2. Suppose that $A < \infty$. Define $u(1) = 1/A$ and then more generally,

$$u(x) = \frac{1}{A} \sum_{i=0}^{x-1} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)}, \quad x \in \mathbb{N}_+ \quad (16.13.8)$$

The function u takes values in $(0, 1)$ and satisfies the functional equation $u = P_+ u$. Hence the chain is transient. Note that $u(x) \rightarrow 1$ as $x \rightarrow \infty$ and so in fact, $u = u_+$, the function that we discussed above that gives the probability of staying in \mathbb{N}_+ for all time. We will return to this function below in our discussion of [absorption](#).

Note that r , the function that assigns to each state $x \in \mathbb{N}$ the probability of an immediate return to x , plays no direct role in whether the chain is transient or recurrent. Indeed all that matters are the ratios $q(x)/p(x)$ for $x \in \mathbb{N}_+$.

Positive Recurrence and Invariant Distributions

Suppose again that we have a birth-death chain \mathbf{X} on \mathbb{N} , with $p(x) > 0$ for all $x \in \mathbb{N}$ and $q(x) > 0$ for all $x \in \mathbb{N}_+$. Thus the chain is irreducible.

The function $g: \mathbb{N} \rightarrow (0, \infty)$ defined by

$$g(x) = \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x)}, \quad x \in \mathbb{N} \quad (16.13.9)$$

is invariant for \mathbf{X} , and is the only invariant function, up to multiplication by constants. Hence \mathbf{X} is positive recurrent if and only if $B = \sum_{x=0}^{\infty} g(x) < \infty$, in which case the (unique) invariant probability density function f is given by $f(x) = \frac{1}{B}g(x)$ for $x \in \mathbb{N}$.

Proof

Recall that by convention, a product over an empty index set is 1. So first,

$$\begin{aligned} (gP)(0) &= g(0)P(0,0) + g(1)P(1,0) = g(0)r(0) + g(1)q(1) \\ &= 1r(0) + \frac{p(0)}{q(1)}q(1) = [1 - p(0)] + p(0) = 1 = g(0) \end{aligned}$$

Next, for $y \in \mathbb{N}_+$,

$$\begin{aligned} (gP)(y) &= g(y-1)P(y-1,y) + g(y)P(y,y) + g(y+1)P(y+1,y) \\ &= g(y-1)p(y-1) + g(y)r(y) + g(y+1)q(y+1) \\ &= g(y-1)p(y-1) + g(y)[1 - p(y) - q(y)] + g(y+1)q(y+1) \end{aligned}$$

But

$$g(y-1)p(y-1) = g(y)q(y) = \frac{p(0) \cdots p(y-1)}{q(1) \cdots q(y-1)}$$

$$g(y+1)q(y+1) = g(y)p(y) = \frac{p(0) \cdots p(y)}{q(1) \cdots q(y)}$$

so $(gP)(y) = g(y)$.

Conversely, suppose that $h : \mathbb{N} \rightarrow \mathbb{R}$ is invariant for \mathbf{X} . We will show by induction that $h(x) = h(0)g(x)$ for all $x \in \mathbb{N}$. The result is trivially true for $x = 0$ since $g(0) = 1$. Next, $(hP)(0) = h(0)$ gives $h(0)P(0,0) + h(1)P(1,0) = h(0)$. But $P(0,0) = r(0) = [1 - p(0)]$ and $P(1,0) = q(1)$, so substituting and solving for $h(1)$ gives

$$h(1) = h(0) \frac{p(0)}{q(1)} = h(0)g(1) \quad (16.13.10)$$

so the result is true when $x = 1$. Assume now that $y \in \mathbb{N}_+$ and that the result is true for all $x \in \mathbb{N}$ with $x \leq y$. Then $(hP)(y) = h(y)$ gives

$$h(y-1)P(y-1, y) + h(y)P(y, y) + h(y+1)P(y+1, y) = h(y) \quad (16.13.11)$$

But $P(y-1, y) = p(y-1)$, $P(y, y) = r(y) = 1 - p(y) - q(y)$, and $P(y+1, y) = q(y+1)$. Also, by the induction hypothesis, $h(y) = h(0)g(y)$ and $h(y-1) = h(0)g(y-1)$ so substituting and using the definition of g gives

$$\begin{aligned} q(y+1)h(y+1) &= [p(y) + q(y)]h(0) \frac{p(0) \cdots p(y-1)}{q(1) \cdots q(y)} - p(y-1)h(0) \frac{p(0) \cdots p(y-2)}{q(1) \cdots q(y-1)} \\ &= h(0) \frac{p(0) \cdots p(y)}{q(1) \cdots q(y)} \end{aligned}$$

Finally, solving gives

$$h(y+1) = h(0) \frac{p(0) \cdots p(y)}{q(1) \cdots q(y+1)} = h(0)g(y+1) \quad (16.13.12)$$

Here is a summary of the classification:

For the birth-death chain \mathbf{X} , define

$$A = \sum_{x=0}^{\infty} \frac{q(1) \cdots q(x)}{p(1) \cdots p(x)}, \quad B = \sum_{x=0}^{\infty} \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x)} \quad (16.13.13)$$

1. \mathbf{X} is transient if $A < \infty$
2. \mathbf{X} is null recurrent if $A = \infty$ and $B = \infty$.
3. \mathbf{X} is positive recurrent if $B < \infty$.

Note again that r , the function that assigns to each state $x \in \mathbb{N}$ the probability of an immediate return to x , plays no direct role in whether the chain is transient, null recurrent, or positive recurrent. Also, we know that an irreducible, recurrent chain has a positive invariant function that is unique up to multiplication by positive constants, but the birth-death chain gives an example where this is also true in the transient case.

Suppose now that $n \in \mathbb{N}_+$ and that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a birth-death chain on the integer interval $\mathbb{N}_n = \{0, 1, \dots, n\}$. We assume that $p(x) > 0$ for $x \in \{0, 1, \dots, n-1\}$ while $q(x) > 0$ for $x \in \{1, 2, \dots, n\}$. Of course, we must have $q(0) = p(n) = 0$. With these assumptions, \mathbf{X} is irreducible, and since the state space is finite, positive recurrent. So all that remains is to find the invariant distribution. The result is essentially the same as when the state space is \mathbb{N} .

The invariant probability density function f_n is given by

$$f_n(x) = \frac{1}{B_n} \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x)} \text{ for } x \in \mathbb{N}_n \text{ where } B_n = \sum_{x=0}^n \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x)} \quad (16.13.14)$$

Proof

Define

$$g_n(x) = \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x)}, \quad x \in \mathbb{N}_n \quad (16.13.15)$$

The proof that g_n is invariant for \mathbf{X} is the same as [before](#). The constant B_n is the normalizing constant.

Note that $B_n \rightarrow B$ as $n \rightarrow \infty$, and if $B < \infty$, $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$ for $x \in \mathbb{N}$. We will see this type of behavior again. Results for the birth-death chain on \mathbb{N}_n often converge to the corresponding results for the birth-death chain on \mathbb{N} as $n \rightarrow \infty$.

Absorption

Often when the state space $S = \mathbb{N}$, the state of a birth-death chain represents a *population* of individuals of some sort (and so the terms *birth* and *death* have their usual meanings). In this case state 0 is absorbing and means that the population is extinct. Specifically, suppose that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a birth-death chain on \mathbb{N} with $r(0) = 1$ and with $p(x), q(x) > 0$ for $x \in \mathbb{N}_+$. Thus, state 0 is absorbing and all positive states lead to each other and to 0. Let $N = \min\{n \in \mathbb{N} : X_n = 0\}$ denote the time until absorption, where as usual, $\min \emptyset = \infty$.

One of the following events will occur:

1. *Population extinction*: $N < \infty$ or equivalently, $X_m = 0$ for some $m \in \mathbb{N}$ and hence $X_n = 0$ for all $n \geq m$.
2. *Population explosion*: $N = \infty$ or equivalently $X_n \rightarrow \infty$ as $n \rightarrow \infty$.

Proof

Part (b) follows from the general theory, since 0 is absorbing, and all positive states lead to each other and to 0. Thus the positive states are transient and we know that with probability 1, a Markov chain will visit a transient state only finitely often. Thus $N = \infty$ is equivalent to $X_n \rightarrow \infty$ as $n \rightarrow \infty$.

Naturally we would like to find the probability of these complementary events, and happily we have already done so in our study of recurrence [above](#). Let

$$u(x) = \mathbb{P}(N = \infty) = \mathbb{P}(X_n \rightarrow \infty \text{ as } n \rightarrow \infty \mid X_0 = x), \quad x \in \mathbb{N} \quad (16.13.16)$$

so the absorption probability is

$$v(x) = 1 - u(x) = \mathbb{P}(N < \infty) = \mathbb{P}(X_n = 0 \text{ for some } n \in \mathbb{N} \mid X_0 = x), \quad x \in \mathbb{N} \quad (16.13.17)$$

For the birth-death chain \mathbf{X} ,

$$u(x) = \frac{1}{A} \sum_{i=0}^{x-1} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)} \text{ for } x \in \mathbb{N}_+ \text{ where } A = \sum_{i=0}^{\infty} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)} \quad (16.13.18)$$

Proof

For $x \in \mathbb{N}_+$, note that $u(x) = \mathbb{P}(X_n \in \mathbb{N}_+ \text{ for all } n \in \mathbb{N} \mid X_0 = x)$, the function that gives the probability of staying in the positive states for all time. The proof of the [theorem on recurrence](#) above has nothing to do with the transition probabilities in state 0, so the proof applies in this setting as well. In that proof we showed that $u(x)$ as the form given above, where of course the value is 0 if $A = \infty$. Trivially, $u(0) = 0$.

So if $A = \infty$ then $u(x) = 0$ for all $x \in S$. If $A < \infty$ then $u(x) > 0$ for all $x \in \mathbb{N}_+$ and $u(x) \rightarrow 1$ as $x \rightarrow \infty$. For the absorption probability, $v(x) = 1$ for all $x \in \mathbb{N}$ if $A = \infty$ and so absorption is certain. If $A < \infty$ then

$$v(x) = \frac{1}{A} \sum_{i=x}^{\infty} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)}, \quad x \in \mathbb{N} \quad (16.13.19)$$

Next we consider the mean time to absorption, so let $m(x) = \mathbb{E}(N \mid X_0 = x)$ for $x \in \mathbb{N}_+$.

The mean absorption function is given by

$$m(x) = \sum_{j=1}^x \sum_{k=j-1}^{\infty} \frac{p(j) \cdots p(k)}{q(j) \cdots q(k+1)}, \quad x \in \mathbb{N} \quad (16.13.20)$$

Probabilistic Proof

The number of steps required to go from state $x \in \mathbb{N}_+$ to $x-1$ has the same distribution as the number of steps required to go from state 1 to 0, except with parameters $p(y), q(y)$ for $y \in \{x, x+1, \dots\}$ instead of parameters $p(y), q(y)$ for $y \in \{1, 2, \dots\}$. So by the additivity of expected value, we just need to compute $m(1)$ as a function of the parameters. Starting in state 1, the chain will be absorbed in state 0 after a random number of returns to state 1 without absorption. Whenever the chain is in state 1, absorption occurs at the next time with probability $q(1)$ so it follows that the number of times that the chain is in state 1 before absorption has the geometric distribution on \mathbb{N}_+ with success parameter $q(1)$. The mean of this distribution is $1/q(1)$. On the other hand, starting in state 1, the number of steps until the chain is in state 1 again (without absorption) has the same distribution as the return time to state 0, starting in state 0 for the irreducible birth-death chain \mathbf{X}' considered [above](#) but with birth and death functions p' and q' given by $p'(x) = p(x+1)$ for $x \in \mathbb{N}$ and $q'(x) = q(x+1)$ for $x \in \mathbb{N}_+$. Thus, let

$$\mu = \sum_{k=0}^{\infty} \frac{p(1) \cdots p(k)}{q(2) \cdots q(k+1)} \quad (16.13.21)$$

Then μ is the mean return time to state 0 for the chain \mathbf{X}' . Specifically, note that if $\mu = \infty$ then \mathbf{X}' is either transient or null recurrent. If $\mu < \infty$ then $1/\mu$ is the invariant PDF at 0. So, it follows that

$$m(1) = \frac{1}{q(1)} \mu = \sum_{k=0}^{\infty} \frac{p(1) \cdots p(k)}{q(1) \cdots q(k+1)} \quad (16.13.22)$$

By our argument above, the mean time to go from state x to $x-1$ is

$$\sum_{k=x-1}^{\infty} \frac{p(x) \cdots p(k)}{q(x) \cdots q(k+1)} \quad (16.13.23)$$

Analytic Proof

Conditioning and using the Markov property, we have

$$m(x) = 1 + p(x)m(x+1) + q(x)m(x-1) + r(x)m(x), \quad x \in \mathbb{N}_+ \quad (16.13.24)$$

with initial condition $m(0) = 0$. Equivalently,

$$m(x+1) - m(x) = \frac{q(x)}{p(x)} [m(x) - m(x-1)] - \frac{1}{p(x)}, \quad x \in \mathbb{N}_+ \quad (16.13.25)$$

Solving gives

$$m(x+1) - m(x) = \frac{q(1) \cdots q(x)}{p(1) \cdots p(x)} m(1) - \sum_{y=1}^x \frac{q(y+1) \cdots q(x)}{p(y) \cdots p(x)}, \quad x \in \mathbb{N}_+ \quad (16.13.26)$$

Next, $m(x) = \sum_{y=0}^{x-1} [m(y+1) - m(y)]$ for $x \in \mathbb{N}$ which gives

$$m(x) = m(1) \sum_{y=0}^{x-1} \frac{q(1) \cdots q(y)}{p(1) \cdots p(y)} - \sum_{y=0}^{x-1} \sum_{z=1}^y \frac{q(z+1) \cdots q(y)}{p(z) \cdots p(y)}, \quad x \in \mathbb{N} \quad (16.13.27)$$

Finally, $m(1)$ is given as in the first proof. The expression for $m(x)$ is different, but equivalent, of course.

Next we will consider a birth-death chain on a finite integer interval with both endpoints absorbing. Our interest is in the probability of absorption in one endpoint rather than the other, and in the mean time to absorption. Thus suppose that $n \in \mathbb{N}_+$ and that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a birth-death chain on $\mathbb{N}_n = \{0, 1, \dots, n\}$ with $r(0) = r(n) = 1$ and with $p(x) > 0$ and $q(x) > 0$ for $x \in \{1, 2, \dots, n-1\}$. So the endpoints 0 and n are absorbing, and all other states lead to each other and to the endpoints. Let $N = \min\{n \in \mathbb{N} : X_n \in \{0, n\}\}$, the time until absorption, and for $x \in S$ let $v_n(x) = \mathbb{P}(X_N = 0 \mid X_0 = x)$ and $m_n(x) = \mathbb{E}(N \mid X_0 = x)$. The definitions make sense since N is finite with probability 1.

The absorption probability function for state 0 is given by

$$v_n(x) = \frac{1}{A_n} \sum_{i=x}^{n-1} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)} \text{ for } x \in \mathbb{N}_n \text{ where } A_n = \sum_{i=0}^{n-1} \frac{q(1) \cdots q(i)}{p(1) \cdots p(i)} \quad (16.13.28)$$

Proof

Conditioning and using the Markov property, v_n satisfies the second-order linear difference equation

$$v_n(x) = p(x)v_n(x+1) + q(x)v_n(x-1) + r(x)v_n(x), \quad x \in \{1, 2, \dots, n-1\} \quad (16.13.29)$$

with boundary conditions $v_n(0) = 1$, $v_n(n) = 0$. As we have seen before, the difference equation can be rewritten as

$$v_n(x+1) - v_n(x) = \frac{p(x)}{q(x)} [v_n(x) - v_n(x-1)], \quad x \in \{1, 2, \dots, n-2\} \quad (16.13.30)$$

Solving and applying the boundary conditions gives the result.

Note that $A_n \rightarrow A$ as $n \rightarrow \infty$ where A is the constant above for the absorption probability at 0 with the infinite state space \mathbb{N} . If $A < \infty$ then $v_n(x) \rightarrow v(x)$ as $n \rightarrow \infty$ for $x \in \mathbb{N}$.

The mean absorption time is given by

$$m_n(x) = m_n(1) \sum_{y=0}^{x-1} \frac{q(1) \cdots q(y)}{p(1) \cdots p(y)} - \sum_{y=0}^{x-1} \sum_{z=1}^y \frac{q(z+1) \cdots q(y)}{p(z) \cdots p(y)}, \quad x \in \mathbb{N}_n \quad (16.13.31)$$

where, with A_n as in the previous theorem,

$$m_n(1) = \frac{1}{A_n} \sum_{y=1}^{n-1} \sum_{z=1}^y \frac{q(z+1) \cdots q(y)}{p(z) \cdots p(y)} \quad (16.13.32)$$

Proof

The [probabilistic proof](#) above with state space \mathbb{N} and 0 absorbing does not work here, but the first part of the analytic proof does. So,

$$m_n(x) = m_n(1) \sum_{y=0}^{x-1} \frac{q(1) \cdots q(y)}{p(1) \cdots p(y)} - \sum_{y=0}^{x-1} \sum_{z=1}^y \frac{q(z+1) \cdots q(y)}{p(z) \cdots p(y)}, \quad x \in \{1, 2, \dots, n\} \quad (16.13.33)$$

Substituting $x = n$ and applying the boundary condition $m_n(n) = 0$, gives the result for $m_n(1)$ in the theorem.

Time Reversal

Our next discussion is on the time reversal of a birth-death chain. Essentially, every recurrent birth-death chain is reversible.

Suppose that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is an irreducible, recurrent birth-death chain on an integer interval S . Then \mathbf{X} is reversible.

Proof

We need to show that the Kolmogorov cycle condition is satisfied. That is, for every sequence of states $(x_0, x_1, x_2, \dots, x_n)$ with $x_0 = x_n$,

$$P(x_0, x_1)P(x_1, x_2) \cdots P(x_{n-1}, x_n) = P(x_n, x_{n-1})P(x_{n-1}, x_{n-2}) \cdots P(x_1, x_0) \quad (16.13.34)$$

We can restrict our attention to sequences where $x_{i+1} \in \{x_i, x_i - 1, x_i + 1\}$ for each $i \in \{1, 2, \dots, n\}$. For such sequences, the cycle condition is trivially satisfied.

If S is finite and the chain \mathbf{X} is irreducible, then of course \mathbf{X} is recurrent (in fact positive recurrent), so by the previous result, \mathbf{X} is reversible. In the case $S = \mathbb{N}$, we can use the invariant function above to show directly that the chain is reversible.

Suppose that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is a birth-death chain on \mathbb{N} with $p(x) > 0$ for $x \in \mathbb{N}$ and $q(x) > 0$ for $x \in \mathbb{N}_+$. Then \mathbf{X} is reversible.

Proof

With the function g defined [above](#), it suffices to show the reversibility condition $g(x)P(x, y) = g(y)P(y, x)$ for all $x, y \in \mathbb{N}$. It then follows that g is invariant for \mathbf{X} and that \mathbf{X} is reversible with respect to g . But since g is the only positive invariant function for \mathbf{X} , up to multiplication by positive constants, we can omit the qualifying phrase “with respect to g ”. For $x \in \mathbb{N}$ and $y = x + 1$ we have

$$g(x)P(x, y) = g(y)P(y, x) = \frac{p(0) \cdots p(x)}{q(1) \cdots q(x)} \quad (16.13.35)$$

For $x \in \mathbb{N}_+$ and $y = x - 1$ we have

$$g(x)P(x, y) = g(y)P(y, x) = \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x-1)} \quad (16.13.36)$$

In all other cases, the reversibility condition is trivially satisfied.

Thus, in the positive recurrent case, when the variables are given the invariant distribution, the transition matrix P describes the chain forward in time and backwards in time.

Examples and Special Cases

As always, be sure to try the problems yourself before looking at the solutions.

Constant Birth and Death Probabilities

Our first examples consider birth-death chains on \mathbb{N} with constant birth and death probabilities, except at the boundary points. Such chains are often referred to as *random walks*, although that term is used in a variety of different settings. The results are special cases of the general results above, but sometimes direct proofs are illuminating.

Suppose that $\mathbf{X} = (X_0, X_1, X_2, \dots)$ is the birth-death chain on \mathbb{N} with constant birth probability $p \in (0, \infty)$ on \mathbb{N} and constant death probability $q \in (0, \infty)$ on \mathbb{N}_+ , with $p + q \leq 1$. Then

1. \mathbf{X} is transient if $q < p$
2. \mathbf{X} is null recurrent if $q = p$
3. \mathbf{X} is positive recurrent if $q > p$, and the invariant distribution is the geometric distribution on \mathbb{N} with parameter p/q

$$f(x) = \left(1 - \frac{p}{q}\right) \left(\frac{p}{q}\right)^x, \quad x \in \mathbb{N} \quad (16.13.37)$$

Next we consider the random walk on \mathbb{N} with 0 absorbing. As in the discussion of [absorption](#) above, $v(x)$ denotes the absorption probability and $m(x)$ the mean time to absorption, starting in state $x \in \mathbb{N}$.

Suppose that $\mathbf{X} = (X_0, X_1, \dots)$ is the birth-death chain on \mathbb{N} with constant birth probability $p \in (0, \infty)$ on \mathbb{N}_+ and constant death probability $q \in (0, \infty)$ on \mathbb{N}_+ , with $p + q \leq 1$. Assume also that $r(0) = 1$, so that 0 is absorbing.

1. If $q \geq p$ then $v(x) = 1$ for all $x \in \mathbb{N}$. If $q < p$ then $v(x) = (q/p)^x$ for $x \in \mathbb{N}$.
2. If $q \leq p$ then $m(x) = \infty$ for all $x \in \mathbb{N}_+$. If $q > p$ then $m(x) = x/(q - p)$ for $x \in \mathbb{N}$.

Proof

1. This follows from the [general result above](#) for the absorption probability.
2. This also follows from the [general result above](#) for the mean absorption time, but we will give a direct proof using the same ideas. If $q < p$ then $\mathbb{P}(N = \infty \mid X_0 = x) > 0$ and hence $m(x) = \infty$ for $x \in \mathbb{N}_+$. So suppose that $q \geq p$ so that $\mathbb{P}(N < \infty \mid X_0 = x) = 1$ for $x \in \mathbb{N}$. Because of the spatial homogeneity, the time required to reach state $x - 1$ starting in state $x \in \mathbb{N}_+$ has the same distribution as the time required to reach state 0 starting in state 1. By the additivity of expected value, it follows that $m(x) = x m(1)$ for $x \in \mathbb{N}$. So it remains for us to compute $m(1)$. Starting in state 1, the chain will be absorbed into state 0 after a random number of intermediate returns to state 1 with absorption. In state 1, the

probability of absorption at the next step is q , so the number of times that the chain is in state 1 before absorption has the geometric distribution on \mathbb{N}_+ with success parameter q . So the mean number of visits is $1/q$. In state 1, the number of steps before a return to step 1 without absorption has the same distribution as the return time to state 0, starting in 0, for the recurrent chain considered in the [previous exercise](#). The mean of this distribution is ∞ if $q = p$ and is $1/f(0)$ if $q > p$, where f is the invariant distribution. It follows that

$$m(1) = \frac{1}{q} \frac{1}{1 - p/q} = \frac{1}{q - p} \quad (16.13.38)$$

This chain is essentially the *gambler's ruin chain*. Consider a gambler who bets on a sequence of independent games, where p and q are the probabilities of winning and losing, respectively. The gambler receives one monetary unit when she wins a game and must pay one unit when she loses a game. So X_n is the gambler's fortune after playing n games.

Next we consider random walks on a finite interval.

Suppose that $\mathbf{X} = (X_0, X_1, \dots)$ is the birth-death chain on $\mathbb{N}_n = \{0, 1, \dots, n\}$ with constant birth probability $p \in (0, \infty)$ on $\{0, 1, \dots, n-1\}$ and constant death probability $q \in (0, \infty)$ on $\{1, 2, \dots, n\}$, with $p + q \leq 1$. Then \mathbf{X} is positive recurrent and the invariant probability density function f_n is given as follows:

1. If $p \neq q$ then

$$f_n(x) = \frac{(p/q)^x (1 - p/q)}{1 - (p/q)^{n+1}}, \quad x \in \mathbb{N}_n \quad (16.13.39)$$

2. If $p = q$ then $f_n(x) = 1/(n+1)$ for $x \in \mathbb{N}_n$.

Note that if $p < q$ then the invariant distribution is a truncated geometric distribution, and $f_n(x) \rightarrow f(x)$ for $x \in \mathbb{N}$ where f is the invariant probability density function of the birth-death chain on \mathbb{N} considered [above](#). If $p = q$, the invariant distribution is uniform on \mathbb{N}_n , certainly a reasonable result. Next we consider the chain with both endpoints absorbing. As before, v_n is the function that gives the probability of absorption in state 0, while m_n is the function that gives the mean time to absorption.

Suppose that $\mathbf{X} = (X_0, X_1, \dots)$ is the birth-death chain on $\mathbb{N}_n = \{0, 1, \dots, n\}$ with constant birth probability $p \in (0, 1)$ and death probability $q \in (0, \infty)$ on $\{1, 2, \dots, n-1\}$, where $p + q \leq 1$. Assume also that $r(0) = r(n) = 1$, so that 0 and n are absorbing.

1. If $p \neq q$ then

$$v_n(x) = \frac{(q/p)^x - (q/p)^n}{1 - (q/p)^n}, \quad x \in \mathbb{N}_n \quad (16.13.40)$$

2. If $p = q$ then $v_n(x) = 1 - x/n$ for $x \in \mathbb{N}_n$.

Note that if $q < p$ then $v_n(x) \rightarrow v(x)$ as $n \rightarrow \infty$ for $x \in \mathbb{N}$.

Suppose again that $\mathbf{X} = (X_0, X_1, \dots)$ is the birth-death chain on $\mathbb{N}_n = \{0, 1, \dots, n\}$ with constant birth probability $p \in (0, 1)$ and death probability $q \in (0, \infty)$ on $\{1, 2, \dots, n-1\}$, where $p + q \leq 1$. Assume also that $r(0) = r(n) = 1$, so that 0 and n are absorbing.

1. If $p \neq q$ then

$$m_n(x) = \frac{n}{p - q} \frac{1 - (q/p)^x}{1 - (q/p)^n} + \frac{x}{q - p}, \quad x \in \mathbb{N}_n \quad (16.13.41)$$

2. If $p = q$ then

$$m_n(x) = \frac{1}{2p} x(n - x), \quad x \in \mathbb{N}_n \quad (16.13.42)$$

Special Birth-Death Chains

Some of the random processes that we have studied previously are birth-death Markov chains.

Describe each of the following as a birth-death chain.

1. The Ehrenfest chain.
2. The modified Ehrenfest chain.
3. The Bernoulli-Laplace chain
4. The simple random walk on \mathbb{Z} .

Answer

1. The Ehrenfest chain with parameter $m \in \mathbb{N}_+$ is a birth death chain on $S = \{0, 1, \dots, m\}$ with $q(x) = \frac{x}{m}$ and $p(x) = \frac{m-x}{m}$ for $x \in S$.
2. The modified Ehrenfest chain with parameter $m \in \mathbb{N}_+$ is a birth death chain on $S = \{0, 1, \dots, m\}$ with $q(x) = \frac{x}{2m}$, $r(x) = \frac{1}{2}$, and $p(x) = \frac{m-x}{2m}$ for $x \in S$.
3. The Bernoulli-Laplace chain with parameters $j, k, r \in \mathbb{N}_+$ with $r < j+k$ is a birth-death chain on $S = \{\max\{0, r-j\}, \dots, \min\{k, r\}\}$ with $q(x) = \frac{(j-r+x)x}{jk}$, $r(x) = \frac{(r-x)x+(j-r+x)(k-x)}{jk}$, and $p(x) = \frac{(r-x)(k-x)}{jk}$ for $x \in S$.
4. The simple random walk on \mathbb{Z} with parameter $p \in (0, 1)$ is a birth-death chain on \mathbb{Z} with $p(x) = p$ and $q(x) = 1-p$ for $x \in \mathbb{Z}$.

Other Examples

Consider the birth-death process on \mathbb{N} with $p(x) = \frac{1}{x+1}$, $q(x) = 1 - p(x)$, and $r(x) = 0$ for $x \in S$.

1. Find the invariant function g .
2. Classify the chain.

Answer

1. Note that $p(0) \cdots p(x-1) = \frac{1}{x!}$ and $q(1) \cdots q(x) = \frac{1}{x+1} = p(x)$ for $x \in \mathbb{N}$. Hence $g(x) = \frac{x+1}{x!}$.
2. Note that

$$\sum_{x=0}^{\infty} g(x) = \sum_{x=1}^{\infty} \frac{1}{(x-1)!} + \sum_{x=0}^{\infty} \frac{1}{x!} = 2e \quad (16.13.43)$$

So the chain is positive recurrent, with invariant PDF f given by

$$f(x) = e^{-2} \frac{(x+1)}{x!}, \quad x \in \mathbb{N} \quad (16.13.44)$$

Also, the chain is periodic with period 2.

This page titled [16.13: Discrete-Time Birth-Death Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.14: Random Walks on Graphs

Basic Theory

Introduction

Suppose that $G = (S, E)$ is a graph with vertex set S and edge set $E \subseteq S^2$. We assume that the graph is *undirected* (perhaps a better term would be *bi-directed*) in the sense that $(x, y) \in E$ if and only if $(y, x) \in E$. The vertex set S is countable, but may be infinite. Let $N(x) = \{y \in S : (x, y) \in E\}$ denote the set of *neighbors* of a vertex $x \in S$, and let $d(x) = \#N(x)$ denote the *degree* of x . We assume that $N(x) \neq \emptyset$ for $x \in S$, so G has no isolated points.

Suppose now that there is a *conductance* $c(x, y) > 0$ associated with each edge $(x, y) \in E$. The conductance is symmetric in the sense that $c(x, y) = c(y, x)$ for $(x, y) \in E$. We extend c to a function on all of $S \times S$ by defining $c(x, y) = 0$ for $(x, y) \notin E$. Let

$$C(x) = \sum_{y \in S} c(x, y), \quad x \in S \quad (16.14.1)$$

so that $C(x)$ is the total conductance of the edges coming from x . Our main assumption is that $C(x) < \infty$ for $x \in S$. As the terminology suggests, we imagine a fluid of some sort flowing through the edges of the graph, so that the conductance of an edge measures the capacity of the edge in some sense. One of the best interpretation is that the graph is an *electrical network* and the edges are *resistors*. In this interpretation, the conductance of a resistor is the reciprocal of the resistance.

In some applications, specifically the resistor network just mentioned, it's appropriate to impose the additional assumption that G has no loops, so that $(x, x) \notin E$ for each $x \in S$. However, that assumption is not mathematically necessary for the Markov chains that we will consider in this section.

The discrete-time Markov chain $\mathbf{X} = (X_0, X_1, X_2, \dots)$ with state space S and transition probability matrix P given by

$$P(x, y) = \frac{c(x, y)}{C(x)}, \quad (x, y) \in S^2 \quad (16.14.2)$$

is called a *random walk* on the graph G .

Justification

First, $P(x, y) \geq 0$ for $x, y \in S$. Next, by definition of C ,

$$\sum_{y \in S} P(x, y) = \sum_{y \in S} \frac{c(x, y)}{C(x)} = \frac{C(x)}{C(x)} = 1, \quad x \in S \quad (16.14.3)$$

so P is a valid transition matrix on S . Also, $P(x, y) > 0$ if and only if $c(x, y) > 0$ if and only if $(x, y) \in E$, so the state graph of \mathbf{X} is G , the graph we started with.

This chain governs a particle moving along the vertices of G . If the particle is at vertex $x \in S$ at a given time, then the particle will be at a neighbor of x at the next time; the neighbor is chosen randomly, in proportion to the conductance. In the setting of an electrical network, it is natural to interpret the particle as an electron. Note that multiplying the conductance function c by a positive constant has no effect on the associated random walk.

Suppose that $d(x) < \infty$ for each $x \in S$ and that c is constant on the edges. Then

1. $C(x) = cd(x)$ for every $x \in S$.
2. The transition matrix P is given by $P(x, y) = \frac{1}{d(x)}$ for $x \in S$ and $y \in N(x)$, and $P(x, y) = 0$ otherwise.

The discrete-time Markov chain \mathbf{X} is the *symmetric random walk* on G .

Proof

1. $C(x) = \sum_{y \in N(x)} c(x, y) = c\#N(x) = cd(x)$ for $x \in S$.
2. $P(x, y) = c(x, y)/C(x) = c/cd(x) = 1/d(x)$ for $x \in S$ and $y \in N(x)$

Thus, for the symmetric random walk, if the state is $x \in S$ at a given time, then the next state is equally likely to be any of the neighbors of x . The assumption that each vertex has finite degree means that the graph G is *locally finite*.

Let \mathbf{X} be a random walk on a graph G .

1. If G is connected then \mathbf{X} is irreducible.
2. If G is not connected then the equivalence classes of \mathbf{X} are the *components* of G (the maximal connected subsets of S).

Proof

1. Recall that there is a path of length $n \in \mathbb{N}_+$ between distinct states $x, y \in S$ in the state graph of \mathbf{X} if and only if $P^n(x, y) > 0$. If G is connected, there is a path between each pair of distinct vertices and hence the chain \mathbf{X} is irreducible.
2. This follows from (a).

So as usual, we will usually assume that G is connected, for otherwise we could simply restrict our attention to a component of G . In the case that G has no loops (again, an important special case because of applications), it's easy to characterize the periodicity of the chain. For the theorem that follows, recall that G is *bipartite* if the vertex set S can be partitioned into nonempty, disjoint sets A and B (the parts) such that every edge in E has one endpoint in A and one endpoint in B .

Suppose that \mathbf{X} is a random walk on a connected graph G with no loops. Then \mathbf{X} is either aperiodic or has period 2. Moreover, \mathbf{X} has period 2 if and only if G is bipartite, in which case the parts are the cyclic classes of \mathbf{X} .

Proof

First note that since G is connected, the chain \mathbf{X} is irreducible, and so all states have the same period. If $(x, y) \in E$ then $(y, x) \in E$ also, so returns to $x \in S$, starting at x can always occur at even positive integers. If G is bipartite, then returns to x starting at x can clearly only occur at even positive integers, so the period is 2. Conversely, if G is not bipartite then G has a cycle of odd length k . If x is a vertex on the cycle, then returns to x , starting at x , can occur in 2 steps or in k steps, so the period of x is 1.

Positive Recurrence and Invariant Distributions

Suppose again that \mathbf{X} is a random walk on a graph G , and assume that G is connected so that \mathbf{X} is irreducible.

The function C is invariant for P . The random walk \mathbf{X} is positive recurrent if and only if

$$K = \sum_{x \in S} C(x) = \sum_{(x,y) \in S^2} c(x, y) < \infty \quad (16.14.4)$$

in which case the invariant probability density function f is given by $f(x) = C(x)/K$ for $x \in S$.

Proof

For $y \in S$,

$$(CP)(y) = \sum_{x \in S} C(x)P(x, y) = \sum_{x \in N(y)} C(x) \frac{c(x, y)}{C(x)} = \sum_{x \in N(y)} c(x, y) = C(y) \quad (16.14.5)$$

so C is invariant for P . The other results follow from the general theory.

Note that K is the total conductance over all edges in G . In particular, of course, if S is finite then \mathbf{X} is positive recurrent, with f as the invariant probability density function. For the symmetric random walk, this is the only way that positive recurrence can occur:

The symmetric random walk on G is positive recurrent if and only if the set of vertices S is finite, in which case the invariant probability density function f is given by

$$f(x) = \frac{d(x)}{2m}, \quad x \in S \quad (16.14.6)$$

where d is the degree function and where m is the number of undirected edges.

Proof

If we take the conductance function to be the constant 1 on the edges, then $C(x) = d(x)$ and $K = 2m$.

On the other hand, when S is infinite, the classification of \mathbf{X} as recurrent or transient is complicated. We will consider an interesting special case below, the [symmetric random walk on \$\mathbb{Z}^k\$](#) .

Reversibility

Essentially, all reversible Markov chains can be interpreted as random walks on graphs. This fact is one of the reasons for studying such walks.

If \mathbf{X} is a random walk on a connected graph G , then \mathbf{X} is reversible with respect to C .

Proof

Since the graph is connected, \mathbf{X} is irreducible. The crucial observation is that

$$C(x)P(x, y) = C(y)P(y, x), \quad (x, y) \in S^2 \quad (16.14.7)$$

If $(x, y) \in E$ the left side is $c(x, y)$ and the right side is $c(y, x)$. If $(x, y) \notin E$, both sides are 0. It then follows from the general theory that C is invariant for \mathbf{X} and that \mathbf{X} is reversible with respect to C .

Of course, if \mathbf{X} is recurrent, then C is the only positive invariant function, up to multiplication by positive constants, and so \mathbf{X} is simply reversible.

Conversely, suppose that \mathbf{X} is an irreducible Markov chain on S with transition matrix P and positive invariant function g . If \mathbf{X} is reversible with respect to g then \mathbf{X} is the random walk on the state graph with conductance function c given by $c(x, y) = g(x)P(x, y)$ for $(x, y) \in S^2$.

Proof

Since \mathbf{X} is reversible with respect to g , g and P satisfy $g(x)P(x, y) = g(y)P(y, x)$ for every $(x, y) \in S^2$. Note that the state graph G of \mathbf{X} is bi-directed since $P(x, y) > 0$ if and only if $P(y, x) > 0$, and that the function c given in the theorem is symmetric, so that $c(x, y) = c(y, x)$ for all $(x, y) \in S^2$. Finally, note that

$$C(x) = \sum_{y \in S} c(x, y) = \sum_{y \in S} g(x)P(x, y) = g(x), \quad x \in S \quad (16.14.8)$$

so that $P(x, y) = c(x, y)/C(x)$ for $(x, y) \in S^2$, as required.

Again, in the important special case that \mathbf{X} is recurrent, there exists a positive invariant function g that is unique up to multiplication by positive constants. In this case the theorem states that an irreducible, recurrent, reversible chain is a random walk on the state graph.

Examples and Applications

The Wheatstone Bridge Graph

The graph below is called the *Wheatstone bridge* in honor of Charles Wheatstone.




Figure 16.14.1: The Wheatstone bridge network, with conductance values in red

In this subsection, let \mathbf{X} be the random walk on the Wheatstone bridge above, with the given conductance values.

For the random walk \mathbf{X} ,

1. Explicitly give the transition probability matrix P .
2. Given $X_0 = a$, find the probability density function of X_2 .

Answer

For the matrix and vector below, we use the ordered state space $S = (a, b, c, d)$.

$$1. P = \begin{bmatrix} 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{2} \\ 0 & \frac{1}{3} & 0 & \frac{2}{3} \\ \frac{1}{5} & \frac{2}{5} & \frac{2}{5} & 0 \end{bmatrix}$$

$$2. f_2 = \left(\frac{9}{40}, \frac{1}{5}, \frac{13}{40}, \frac{1}{4} \right)$$

For the random walk \mathbf{X} ,

1. Show that \mathbf{X} is aperiodic.
2. Find the invariant probability density function.
3. Find the mean return time to each state.
4. Find $\lim_{n \rightarrow \infty} P^n$.

Answer

For the matrix and vectors below, we use the ordered state space (a, b, c, d) .

1. The chain is aperiodic since the graph is not bipartite. (Note that the graph has triangles.)

$$2. f = \left(\frac{1}{7}, \frac{2}{7}, \frac{3}{14}, \frac{5}{14} \right)$$

$$3. \mu = \left(7, \frac{7}{2}, \frac{14}{3}, \frac{14}{5} \right)$$

$$4. P^n \rightarrow \begin{bmatrix} \frac{1}{7} & \frac{2}{7} & \frac{3}{14} & \frac{5}{14} \\ \frac{1}{7} & \frac{2}{7} & \frac{3}{14} & \frac{5}{14} \\ \frac{1}{7} & \frac{2}{7} & \frac{3}{14} & \frac{5}{14} \\ \frac{1}{7} & \frac{2}{7} & \frac{3}{14} & \frac{5}{14} \end{bmatrix} \text{ as } n \rightarrow \infty$$

The Cube Graph

The graph below is the 3-dimensional *cube graph*. The vertices are bit strings of length 3, and two vertices are connected by an edge if and only if the bit strings differ by a single bit.




Figure 16.14.2: The cube graph with conductance values in red

In this subsection, let \mathbf{X} denote the random walk on the cube graph above, with the given conductance values.

For the random walk \mathbf{X} ,

1. Explicitly give the transition probability matrix P .
2. Suppose that the initial distribution is the uniform distribution on $\{000, 001, 101, 100\}$. Find the probability density function of X_2 .

Answer

For the matrix and vector below, we use the ordered state space $S = (000, 001, 101, 110, 010, 011, 111, 101)$

$$1. P = \begin{bmatrix} 0 & \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{2} & 0 & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} & 0 & 0 & \frac{1}{2} & 0 & 0 \\ 0 & \frac{1}{4} & 0 & \frac{1}{4} & 0 & 0 & \frac{1}{2} & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} & 0 & 0 & 0 & 0 & \frac{1}{2} \\ \frac{1}{4} & 0 & 0 & 0 & 0 & \frac{3}{9} & 0 & \frac{3}{8} \\ 0 & \frac{1}{4} & 0 & 0 & \frac{3}{8} & 0 & \frac{3}{8} & 0 \\ 0 & 0 & \frac{1}{4} & 0 & 0 & \frac{3}{8} & 0 & \frac{3}{8} \\ 0 & 0 & 0 & \frac{1}{4} & \frac{3}{8} & 0 & \frac{3}{8} & 0 \end{bmatrix}$$

$$2. f_2 = \left(\frac{3}{32}, \frac{3}{32}, \frac{3}{32}, \frac{3}{32}, \frac{5}{32}, \frac{5}{32}, \frac{5}{32}, \frac{5}{32} \right)$$

For the random walk X ,

1. Show that the chain has period 2 and find the cyclic classes.
2. Find the invariant probability density function.
3. Find the mean return time to each state.
4. Find $\lim_{n \rightarrow \infty} P^{2n}$.
5. Find $\lim_{n \rightarrow \infty} P^{2n+1}$.

Answer

For the matrix and vector below, we use the ordered state space $S = (000, 001, 101, 110, 010, 011, 111, 101)$

1. The chain has period 2 since the graph is bipartite. The cyclic classes are $\{000, 011, 110, 101\}$ (bit strings with an even number of 1's) and $\{010, 001, 100, 111\}$ (bit strings with an odd number of 1's).

$$2. f = \left(\frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6} \right)$$

$$3. \mu = (12, 12, 12, 12, 6, 6, 6, 6)$$

$$4. P^{2n} \rightarrow \begin{bmatrix} \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \end{bmatrix} \text{ as } n \rightarrow \infty$$

5.

$$6. P^{2n+1} \rightarrow \begin{bmatrix} 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{3} & 0 & \frac{1}{3} & 0 \end{bmatrix} \text{ as } n \rightarrow \infty$$

Special Models

Recall that the basic Ehrenfest chain with $m \in \mathbb{N}_+$ balls is reversible. Interpreting the chain as a random walk on a graph, sketch the graph and find a conductance function.

Answer

The state graph G of the basic Ehrenfest chain with m balls is the path from 0 to m with no loops. A conductance function c is $c(x, x+1) = \binom{m-1}{x}$ for $x \in \{0, 1, \dots, m-1\}$.

Recall that the modified Ehrenfest chain with $m \in \mathbb{N}_+$ balls is reversible. Interpreting the chain as a random walk on a graph, sketch the graph and find a conductance function.

Answer

The state graph G of the modified Ehrenfest chain with m balls is the path from 0 to m with loops. A conductance function c is $c(x, x+1) = \frac{1}{2} \binom{m-1}{x}$ for $x \in \{0, 1, \dots, m-1\}$ and $c(x, x) = \frac{1}{2} \binom{m}{x}$ for $x \in \{0, 1, \dots, m\}$.

Recall that the Bernoulli-Laplace chain with $j \in \mathbb{N}_+$ balls in urn 0, $k \in \mathbb{N}_+$ balls in urn 1, and with $r \in \{0, \dots, j+k\}$ of the balls red, is reversible. Interpreting the chain as a random walk on a graph, sketch the graph and find a conductance function. Simplify the conductance function in the special case that $j = k = r$.

Answer

The state graph G of the Bernoulli-Laplace chain with j balls in urn 0, k balls in urn 1, and with r of the balls red, is the path from $\max\{0, r-j\}$ to $\min\{k, r\}$ with loops. A conductance function c is given by

$$c(x, x+1) = \binom{r}{x} \binom{j+k-r}{k-x} (r-x)(k-x), \quad x \in \{\max\{0, r-j\}, \dots, \min\{k, r\} - 1\}$$

$$c(x, x) = \binom{r}{x} \binom{j+k-r}{k-x} [(r-x)x + (j-r+x)(k-x)], \quad x \in \{\max\{0, r-j\}, \dots, \min\{k, r\}\}$$

In the special case that $j = k = r$, a conductance function is

$$c(x, x+1) = \binom{k}{x} \binom{k}{k-x} (k-x)^2, \quad x \in \{0, \dots, k-1\}$$

$$c(x, x) = \binom{k}{x} \binom{k}{k-x} 2x(k-x), \quad x \in \{0, \dots, k\}$$

Random Walks on \mathbb{Z}

Random walks on integer lattices are particularly interesting because of their classification as transient or recurrent. We consider the one-dimensional case in this subsection, and the higher dimensional case in the next subsection.

Let $\mathbf{X} = (X_0, X_1, X_2, \dots)$ be the discrete-time Markov chain with state space \mathbb{Z} and transition probability matrix P given by

$$P(x, x+1) = p, \quad P(x, x-1) = 1-p, \quad x \in \mathbb{Z} \quad (16.14.9)$$

where $p \in (0, 1)$. The chain \mathbf{X} is called the *simple random walk* on \mathbb{Z} with parameter p .

The term *simple* is used because the transition probabilities starting in state $x \in \mathbb{Z}$ do not depend on x . Thus the chain is *spatially* as well as *temporally* homogeneous. In the special case $p = \frac{1}{2}$, the chain \mathbf{X} is the *simple symmetric random walk* on \mathbb{Z} . Basic properties of the simple random walk on \mathbb{Z} , and in particular, the simple symmetric random walk were studied in the chapter on Bernoulli Trials. Of course, the state graph G of \mathbf{X} has vertex set \mathbb{Z} , and the neighbors of $x \in \mathbb{Z}$ are $x+1$ and $x-1$. It's not immediately clear that \mathbf{X} is a random walk on G associated with a conductance function, which after all, is the topic of this section. But that fact and more follow from the next result.

Let g be the function on \mathbb{Z} defined by

$$g(x) = \left(\frac{p}{1-p} \right)^x, \quad x \in \mathbb{Z} \quad (16.14.10)$$

Then

1. $g(x)P(x, y) = g(y)P(y, x)$ for all $(x, y) \in \mathbb{Z}^2$
2. g is invariant for \mathbf{X}
3. \mathbf{X} is reversible with respect to g
4. \mathbf{X} is the random walk on \mathbb{Z} with conductance function c given by $c(x, x+1) = p^{x+1} / (1-p)^x$ for $x \in \mathbb{Z}$.

Proof

1. For $x \in \mathbb{Z}$, we only need to consider $y = x \pm 1$.

$$g(x)P(x, x-1) = \frac{p^x}{(1-p)^{x-1}} = g(x-1)P(x-1, x)$$

$$g(x)P(x, x+1) = \frac{p^{x+1}}{(1-p)^x} = g(x+1)P(x+1, x)$$

2. This follows from (a) and the general theory.
3. This follows from (a) and (b) and the general theory.
4. From the result [above](#), \mathbf{X} is the random walk on G associated with the conductance function c given by $c(x, y) = g(x)P(x, y)$. By symmetry, it suffices to consider the edge $(x, x+1)$, and in this case, c is given in the second displayed equation above.

In particular, the simple symmetric random walk is the symmetric random walk on G .

The chain \mathbf{X} is irreducible and periodic with period 2. Moreover

$$P^{2n}(0, 0) = \binom{2n}{n} p^n (1-p)^n, \quad n \in \mathbb{N} \quad (16.14.11)$$

Proof

The chain is irreducible since G is connected. The chain is periodic since G has no loops and is bipartite, with the parts being the odd and even integers. Finally, note that starting in state 0, the chain returns to 0 at time $2n$ if and only if there are n steps to the right and n steps to the left.

Classification of the simple random walk on \mathbb{Z} .

1. If $p \neq \frac{1}{2}$ then \mathbf{X} is transient.
2. If $p = \frac{1}{2}$ then \mathbf{X} is null recurrent.

Proof

From the [previous result](#) and Stirling's approximation,

$$P^{2n}(0, 0) \approx \frac{[4p(1-p)]^n}{\sqrt{\pi n}} \text{ as } n \rightarrow \infty \quad (16.14.12)$$

Let $R(x, y) = \sum_{n=0}^{\infty} P^n(x, y)$ for $(x, y) \in \mathbb{Z}^2$, so that R is the potential matrix. Recall that $R(x, y)$ is the expected number of visits to y starting in x for $(x, y) \in \mathbb{Z}^2$. If $p \neq \frac{1}{2}$ then $R(0, 0) < \infty$ and hence \mathbf{X} is transient. If $p = \frac{1}{2}$ then $R(0, 0) = \infty$ and hence \mathbf{X} is recurrent. In this case \mathbf{X} must be null recurrent from our general results [above](#), since the vertex set is infinite.

So for the one-dimensional lattice \mathbb{Z} , the random walk \mathbf{X} is transient in the non-symmetric case, and null recurrent in the symmetric case. Let's return to the invariant functions of \mathbf{X}

Consider again the random walk \mathbf{X} on \mathbb{Z} with parameter $p \in (0, 1)$. The constant function $\mathbf{1}$ on \mathbb{Z} and the function g given by

$$g(x) = \left(\frac{p}{1-p} \right)^x, \quad x \in \mathbb{Z} \quad (16.14.13)$$

are invariant for \mathbf{X} . All other invariant functions are linear combinations of these two functions.

Proof

The condition for h to be invariant, $hP = h$, leads to the following linear, second order difference equation:

$$(1-p)h(y+1) - h(y) + (1+p)h(y-1), \quad y \in \mathbb{Z} \quad (16.14.14)$$

The characteristic equation is $(1-p)r^2 - r + (1+p) = 0$ which has roots $r = 1$ and $r = p/(1-p)$. The solutions corresponding to the roots are $\mathbf{1}$ and g , respectively. Hence the result follows from the general theory of linear difference equations.

Note that when $p = \frac{1}{2}$, the constant function $\mathbf{1}$ is the only positive invariant function, up to multiplication by positive constants. But we know this has to be the case since the chain is recurrent when $p = \frac{1}{2}$. Moreover, the chain is reversible. In the non-symmetric case, when $p \neq \frac{1}{2}$, we have an example of a transient chain which nonetheless has non-trivial invariant functions—in fact a two dimensional space of such functions. Also, \mathbf{X} is reversible with respect to g , as shown [above](#), but the reversal of \mathbf{X} with respect to $\mathbf{1}$ is the chain with transition matrix Q given by $Q(x, y) = P(y, x)$ for $(x, y) \in \mathbb{Z}^2$. This chain is just the simple random walk on Z with parameter $1-p$. So the non-symmetric simple random walk is an example of a transient chain that is reversible with respect to one invariant measure but not with respect to another invariant measure.

Random walks on \mathbb{Z}^k

More generally, we now consider \mathbb{Z}^k , where $k \in \mathbb{N}_+$. For $i \in \{1, 2, \dots, k\}$, let $\mathbf{u}_i \in \mathbb{Z}^k$ denote the unit vector with 1 in position i and 0 elsewhere. The k -dimensional integer lattice G has vertex set Z^k , and the neighbors of $\mathbf{x} \in \mathbb{Z}^k$ are $\mathbf{x} \pm \mathbf{u}_i$ for $i \in \{1, 2, \dots, k\}$. So in particular, each vertex has $2k$ neighbors.

Let $\mathcal{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots)$ be the Markov chain on \mathbb{Z}^k with transition probability matrix P given by

$$P(\mathbf{x}, \mathbf{x} + \mathbf{u}_i) = p_i, \quad P(\mathbf{x}, \mathbf{x} - \mathbf{u}_i) = q_i; \quad \mathbf{x} \in \mathbb{Z}^k, \quad i \in \{1, 2, \dots, k\} \quad (16.14.15)$$

where $p_i > 0$, $q_i > 0$ for $i \in \{1, 2, \dots, k\}$ and $\sum_{i=1}^k (p_i + q_i) = 1$. The chain \mathcal{X} is the *simple random walk on \mathbb{Z}^k* with parameters $\mathbf{p} = (p_1, p_2, \dots, p_k)$ and $\mathbf{q} = (q_1, q_2, \dots, q_k)$.

Again, the term *simple* means that the transition probabilities starting at $\mathbf{x} \in \mathbb{Z}^k$ do not depend on \mathbf{x} , so that the chain is spatially homogeneous as well as temporally homogeneous. In the special case that $p_i = q_i = \frac{1}{2k}$ for $i \in \{1, 2, \dots, k\}$, \mathcal{X} is the *simple symmetric random walk on Z^k* . The following theorem is the natural generalization of the result [abpve](#) for the one-dimensional case.

Define the function $g: \mathbb{Z}^k \rightarrow (0, \infty)$ by

$$g(x_1, x_2, \dots, x_k) = \prod_{i=1}^k \left(\frac{p_i}{q_i} \right)^{x_i}, \quad (x_1, x_2, \dots, x_k) \in \mathbb{Z}^k \quad (16.14.16)$$

Then

1. $g(\mathbf{x})P(\mathbf{x}, \mathbf{y}) = g(\mathbf{y})P(\mathbf{y}, \mathbf{x})$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^k$
2. g is invariant for \mathcal{X} .
3. \mathcal{X} is reversible with respect to g .
4. \mathcal{X} is the random walk on G with conductance function c given by $c(\mathbf{x}, \mathbf{y}) = g(\mathbf{x})P(\mathbf{x}, \mathbf{y})$ for $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^k$.

Proof

1. For $\mathbf{x} = (x_1, x_2, \dots, x_k) \in \mathbb{Z}^k$, the only cases of interest are $\mathbf{y} = \mathbf{x} \pm \mathbf{u}_i$ for $i \in \{1, 2, \dots, k\}$, since in all other cases, the left and right sides are 0. But

$$\begin{aligned} g(\mathbf{x})P(\mathbf{x}, \mathbf{x} + \mathbf{u}_i) &= \prod_{j \neq i} \left(\frac{p_j}{q_j} \right)^{x_j} \cdot \frac{p_i^{x_i+1}}{q_i^{x_i}} = g(\mathbf{x} + \mathbf{u}_i)P(\mathbf{x} + \mathbf{u}_i, \mathbf{x}) \\ g(\mathbf{x})P(\mathbf{x}, \mathbf{x} - \mathbf{u}_i) &= \prod_{j \neq i} \left(\frac{p_j}{q_j} \right)^{x_j} \cdot \frac{p_i^{x_i}}{q_i^{x_i-1}} = g(\mathbf{x} - \mathbf{u}_i)P(\mathbf{x} - \mathbf{u}_i, \mathbf{x}) \end{aligned}$$

2. This follows from (a).
3. This follows from (a) and (b).

4. This also follows from the general result [above](#).

It terms of recurrence and transience, it would certainly seem that the larger the dimension k , the less likely the chain is to be recurrent. That's generally true:

Classification of the simple random walk on \mathbb{Z}^k .

1. For $k \in \{1, 2\}$, \mathcal{X} is null recurrent in the symmetric case and transient for all other values of the parameters.
2. For $k \in \{3, 4, \dots\}$, \mathcal{X} is transient for all values of the parameters.

Proof sketch

For certain of the non-symmetric cases, we can use the result for dimension 1. Suppose $i \in \{1, 2, \dots, k\}$ with $p_i \neq q_i$. If we consider the times when coordinate i of the random walk \mathcal{X} changes, we have an embedded one-dimensional random walk with parameter $p = p_i / (p_i + q_i)$ (the probability of a step in the positive direction). Since $p \neq \frac{1}{2}$, this embedded random walk is transient and so will fail to return to 0, starting at 0, with positive probability. But if this embedded random walk fails to return to 0, starting at 0, then the parent random walk \mathcal{X} fails to return to $\mathbf{0}$ starting at $\mathbf{0}$. Hence \mathcal{X} is transient.

For the symmetric case, the general proof is similar in to the proof for [dimension 1](#), but the details are considerably more complex. A return to 0 can occur only at even times and

$$P^{2n}(\mathbf{0}, \mathbf{0}) \approx \frac{C_k}{n^{k/2}} \text{ as } n \rightarrow \infty \text{ where } C_k = \frac{k^{k/2}}{\pi^{k/2} 2^{k-1}} \quad (16.14.17)$$

Thus for the potential matrix R we have $R(\mathbf{0}, \mathbf{0}) = \infty$ and the chain is recurrent if $k \in \{1, 2\}$ while $R(\mathbf{0}, \mathbf{0}) < \infty$ and the chain is transient if $k \in \{3, 4, \dots\}$.

So for the simple, symmetric random walk on the integer lattice \mathbb{Z}^k , we have the following interesting *dimensional phase shift*: the chain is null recurrent in dimensions 1 and 2 and transient in dimensions 3 or more.

Let's return to the positive invariant functions for \mathcal{X} . Again, the results generalize those for the one-dimensional case.

For $J \subseteq \{1, 2, \dots, k\}$, define g_J on \mathbb{Z}^k by

$$g_J(x_1, x_2, \dots, x_k) = \prod_{j \in J} \left(\frac{p_j}{q_j} \right)^{x_j}, \quad (x_1, x_2, \dots, x_k) \in \mathbb{Z}^k \quad (16.14.18)$$

Let \mathcal{X}_J denote the simple random walk on \mathbb{Z}^k with transition matrix P_J , corresponding to the parameter vectors \mathbf{p}^J and \mathbf{q}^J , where $p_j^J = p_j, q_j^J = q_j$ for $j \in J$, and $p_j^J = q_j, q_j^J = p_j$ for $j \notin J$. Then

1. $g_J(\mathbf{x})P(\mathbf{x}, \mathbf{y}) = g_J(\mathbf{y})P_J(\mathbf{y}, \mathbf{x})$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^k$
2. g_J is invariant for \mathcal{X} .
3. \mathcal{X}_J is reversal of \mathcal{X} with respect to g_J .

Proof

Part (a) follows from simple substitution. Parts (b) and (c) follow from (a) and the general theory.

Note that when $J = \emptyset$, $g_J = \mathbf{1}$ and when $J = \{1, 2, \dots, k\}$, $g_J = g$, the invariant function introduced [above](#). So in the completely non-symmetric case where $p_i \neq q_i$ for every $i \in \{1, 2, \dots, k\}$, the random walk \mathcal{X} has 2^k positive invariant functions that are linearly independent, and \mathcal{X} is reversible with respect to one of them.

This page titled [16.14: Random Walks on Graphs](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.15: Introduction to Continuous-Time Markov Chains

This section begins our study of Markov processes in continuous time and with discrete state spaces. Recall that a Markov process with a discrete state space is called a *Markov chain*, so we are studying *continuous-time Markov chains*. It will be helpful if you review the section on general Markov processes, at least briefly, to become familiar with the basic notation and concepts. Also, discrete-time chains plays a fundamental role, so you will need review this topic also.

We will study continuous-time Markov chains from different points of view. Our point of view in this section, involving holding times and the embedded discrete-time chain, is the most intuitive from a probabilistic point of view, and so is the best place to start. In the next section, we study the transition probability matrices in continuous time. This point of view is somewhat less intuitive, but is closest to how other types of Markov processes are treated. Finally, in the third introductory section we study the Markov chain from the view point of potential matrices. This is the least intuitive approach, but analytically one of the best. Naturally, the interconnections between the various approaches are particularly important.

Preliminaries

As usual, we start with a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, so that Ω is the set of outcomes, \mathcal{F} the σ -algebra of events, and \mathbb{P} the probability measure on the sample space (Ω, \mathcal{F}) . The time space is $([0, \infty), \mathcal{T})$ where as usual, \mathcal{T} is the Borel σ -algebra on $[0, \infty)$ corresponding to the standard Euclidean topology. The state space is (S, \mathcal{S}) where S is countable and \mathcal{S} is the power set of S . So every subset of S is measurable, as is every function from S to another measurable space. Recall that \mathcal{S} is also the Borel σ algebra corresponding to the discrete topology on S . With this topology, every function from S to another topological space is continuous. Counting measure $\#$ is the natural measure on (S, \mathcal{S}) , so in the context of the general introduction, integrals over S are simply sums. Also, kernels on S can be thought of as matrices, with rows and sums indexed by S . The left and right kernel operations are generalizations of matrix multiplication.

Suppose now that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is stochastic process with state space (S, \mathcal{S}) . For $t \in [0, \infty)$, let $\mathcal{F}_t^0 = \sigma\{X_s : s \in [0, t]\}$, so that \mathcal{F}_t^0 is the σ -algebra of events defined by the process up to time t . The collection of σ -algebras $\mathfrak{F}^0 = \{\mathcal{F}_t^0 : t \in [0, \infty)\}$ is the natural filtration associated with \mathbf{X} . For technical reasons, it's often necessary to have a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ that is slightly finer than the natural one, so that $\mathcal{F}_t^0 \subseteq \mathcal{F}_t$ for $t \in [0, \infty)$ (or in equivalent jargon, \mathbf{X} is adapted to \mathfrak{F}). See the general introduction for more details on the common ways that the natural filtration is refined. We will also let $\mathcal{G}_t = \sigma\{X_s : s \in [t, \infty)\}$, the σ -algebra of events defined by the process from time t onward. If t is thought of as the present time, then \mathcal{F}_t is the collection of events in the past and \mathcal{G}_t is the collection of events in the future.

It's often necessary to impose assumptions on the continuity of the process \mathbf{X} in time. Recall that \mathbf{X} is *right continuous* if $t \mapsto X_t(\omega)$ is right continuous on $[0, \infty)$ for every $\omega \in \Omega$, and similarly \mathbf{X} has left limits if $t \mapsto X_t(\omega)$ has left limits on $(0, \infty)$ for every $\omega \in \Omega$. Since S has the discrete topology, note that if \mathbf{X} is right continuous, then for every $t \in [0, \infty)$ and $\omega \in \Omega$, there exists ϵ (depending on t and ω) such that $X_{t+s}(\omega) = X_t(\omega)$ for $s \in [0, \epsilon)$. Similarly, if \mathbf{X} has left limits, then for every $t \in (0, \infty)$ and $\omega \in \Omega$ there exists δ (depending on t and ω) such that $X_{t-s}(\omega)$ is constant for $s \in (0, \delta)$.

The Markov Property

There are a number of equivalent ways to state the Markov property. At the most basic level, the property states that the *past and future are conditionally independent, given the present*.

The process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a *Markov chain* on S if for every $t \in [0, \infty)$, $A \in \mathcal{F}_t$, and $B \in \mathcal{G}_t$,

$$\mathbb{P}(A \cap B \mid X_t) = \mathbb{P}(A \mid X_t) \mathbb{P}(B \mid X_t) \quad (16.15.1)$$

Another version is that the conditional distribution of a state in the future, given the past, is the same as the conditional distribution just given the present state.

The process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S if for every $s, t \in [0, \infty)$, and $x \in S$,

$$\mathbb{P}(X_{s+t} = x \mid \mathcal{F}_s) = \mathbb{P}(X_{s+t} = x \mid X_s) \quad (16.15.2)$$

Technically, in the last two definitions, we should say that \mathbf{X} is a Markov process *relative to* the filtration \mathfrak{F} . But recall that if \mathbf{X} satisfies the Markov property relative to a filtration, then it satisfies the Markov property relative to any coarser filtration, and in particular, relative to the natural filtration. For the natural filtration, the Markov property can also be stated without explicit reference to σ -algebras, although at the cost of additional clutter:

The process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S if and only if for every $n \in \mathbb{N}_+$, time sequence $(t_1, t_2, \dots, t_n) \in [0, \infty)^n$ with $t_1 < t_2 < \dots < t_n$, and state sequence $(x_1, x_2, \dots, x_n) \in S^n$,

$$\mathbb{P}(X_{t_n} = x_n \mid X_{t_1} = x_1, X_{t_2} = x_2, \dots, X_{t_{n-1}} = x_{n-1}) = \mathbb{P}(X_{t_n} = x_n \mid X_{t_{n-1}} = x_{n-1}) \quad (16.15.3)$$

As usual, we also assume that our Markov chain \mathbf{X} is *time homogeneous*, so that $\mathbb{P}(X_{s+t} = y \mid X_s = x) = \mathbb{P}(X_t = y \mid X_0 = x)$ for $s, t \in [0, \infty)$ and $x, y \in S$. So, for a homogeneous Markov chain on S , the process $\{X_{s+t} : t \in [0, \infty)\}$ given $X_s = x$, is independent of \mathcal{F}_s and equivalent to the process $\{X_t : t \in [0, \infty)\}$ given $X_0 = x$, for every $s \in [0, \infty)$ and $x \in S$. That is, if the chain is in state $x \in S$ at a particular time $s \in [0, \infty)$, it does not matter how the chain got to x ; the chain essentially starts over in state x .

The Strong Markov Property

Random times play an important role in the study of continuous-time Markov chains. It's often necessary to allow random times to take the value ∞ , so formally, a random time τ is a random variable on the underlying sample space (Ω, \mathcal{F}) taking values in $[0, \infty]$. Recall also that a random time τ is a stopping time (also called a *Markov time* or an *optional time*) if $\{\tau \leq t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$. If τ is a stopping time, the σ -algebra associated with τ is

$$\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \in [0, \infty)\} \quad (16.15.4)$$

So \mathcal{F}_τ is the collection of events up to the random time τ in the same way that \mathcal{F}_t is the collection of events up to the deterministic time $t \in [0, \infty)$. We usually want the Markov property to extend from deterministic times to stopping times.

The process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a *strong Markov chain* on S if for every stopping time $\tau, t \in [0, \infty)$, and $x \in S$,

$$\mathbb{P}(X_{\tau+t} = x \mid \mathcal{F}_\tau) = \mathbb{P}(X_{\tau+t} = x \mid X_\tau) \quad (16.15.5)$$

So, for a homogeneous strong Markov chain on S , the process $\{X_{\tau+t} : t \in [0, \infty)\}$ given $X_\tau = x$, is independent of \mathcal{F}_τ and equivalent to the process $\{X_t : t \in [0, \infty)\}$ given $X_0 = x$, for every stopping time τ and $x \in S$. That is, if the chain is in state $x \in S$ at a stopping time τ , then the chain essentially starts over at x , independently of the past.

Holding Times and the Jump Chain

For our first point of view, we still study *when* and *how* our Markov chain \mathbf{X} changes state. The discussion depends heavily on properties of the exponential distribution, so we need a quick review.

The Exponential Distribution

A random variable τ has the *exponential distribution* with rate parameter $r \in (0, \infty)$ if τ has a continuous distribution on $[0, \infty)$ with probability density function f given by $f(t) = re^{-rt}$ for $t \in [0, \infty)$. Equivalently, the right distribution function F^c is given by

$$F^c(t) = \mathbb{P}(\tau > t) = e^{-rt}, \quad t \in [0, \infty) \quad (16.15.6)$$

The mean of the distribution is $1/r$ and the variance is $1/r^2$. The exponential distribution has an amazing number of characterizations. One of the most important is the *memoryless property* which states that a random variable τ with values in $[0, \infty)$ has an exponential distribution if and only if the conditional distribution of $\tau - s$ given $\tau > s$ is the same as the distribution of τ itself, for every $s \in [0, \infty)$. It's easy to see that the memoryless property is equivalent to the *law of exponents* for right distribution function F^c , namely $F^c(s+t) = F^c(s)F^c(t)$ for $s, t \in [0, \infty)$. Since F^c is right continuous, the only solutions are exponential functions.

For our study of continuous-time Markov chains, it's helpful to extend the exponential distribution to two degenerate cases, $\tau = 0$ with probability 1, and $\tau = \infty$ with probability 1. In terms of the parameter, the first case corresponds to $r = \infty$ so that $F(t) = \mathbb{P}(\tau > t) = 0$ for every $t \in [0, \infty)$, and the second case corresponds to $r = 0$ so that $F(t) = \mathbb{P}(\tau > t) = 1$ for every $t \in [0, \infty)$. Note that in both cases, the function F satisfies the law of exponents, and so corresponds to a memoryless distribution in a general sense. In all cases, the mean of the exponential distribution with parameter $r \in [0, \infty]$ is $1/r$, where we interpret $1/0 = \infty$ and $1/\infty = 0$.

Holding Times

The Markov property implies the memoryless property for the random time when a Markov process first leaves its initial state. It follows that this random time must have an exponential distribution.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S , and let $\tau = \inf\{t \in [0, \infty) : X_t \neq X_0\}$. For $x \in S$, the conditional distribution of τ given $X_0 = x$ is exponential with parameter $\lambda(x) \in [0, \infty]$.

Proof

Let $x \in S$ and $s \in [0, \infty)$. The events $X_0 = x$ and $\tau > s$ imply $X_s = x$. By the Markov property, given $X_s = x$, the chain starts over at time s in state x , independent of $\{X_0 = x\}$ and $\{\tau > s\}$, since both events are in \mathcal{F}_s . Hence for $t \in [0, \infty)$,

$$\mathbb{P}(\tau > t+s \mid X_0 = x, \tau > s) = \mathbb{P}(\tau > t+s \mid X_0 = x, X_s = x, \tau > s) = \mathbb{P}(\tau > t \mid X_0 = x) \quad (16.15.7)$$

It follows that τ has the memoryless property, and hence has an exponential distribution with parameter $\lambda(x) \in [0, \infty]$.

So, associated with the Markov chain \mathbf{X} on S is a function $\lambda : S \rightarrow [0, \infty]$ that gives the exponential parameters for the holding times in the states. Considering the ordinary exponential distribution, and the two degenerate versions, we are led to the following classification of states:

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S with exponential parameter function λ . Let $x \in S$.

1. If $\lambda(x) = 0$ then $\mathbb{P}(\tau = \infty \mid X_0 = x) = 1$, and x is said to be an *absorbing state*.
2. If $\lambda(x) \in (0, \infty)$ then $\mathbb{P}(0 < \tau < \infty \mid X_0 = x) = 1$ and x is said to be an *stable state*.
3. If $\lambda(x) = \infty$ then $\mathbb{P}(\tau = 0 \mid X_0 = x) = 1$, and x is said to be an *instantaneous state*.

As you can imagine, an instantaneous state corresponds to weird behavior, since the chain starting in the state leaves the state at times arbitrarily close to 0. While mathematically possible, instantaneous states make no sense in most applications, and so are to be avoided. Also, the proof of the last result has some technical holes. We did not really show that τ is a valid random time, let alone a stopping time. Fortunately, one of our standard assumptions resolves these problems.

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S . If the process \mathbf{X} and the filtration \mathfrak{F} are right continuous, then

1. τ is a stopping time.
2. \mathbf{X} has no instantaneous states.
3. $\mathbb{P}(X_\tau \neq x \mid X_0 = x) = 1$ if $x \in S$ is stable.
4. \mathbf{X} is a strong Markov process.

Proof

1. Let $t \in [0, \infty)$. By right continuity,

$$\{\tau < t\} = \{X_s \neq X_0 \text{ for some } s \in (0, t)\} = \{X_s \neq X_0 \text{ for some rational } s \in (0, t)\} \quad (16.15.8)$$

But for $s \in (0, t)$, $\{X_s \neq X_0\} \in \mathcal{F}_s \subseteq \mathcal{F}_t$. The last event in the displayed equation is a countable union, so $\{\tau < t\} \in \mathcal{F}_t$. Since \mathfrak{F} is right continuous, τ is a stopping time.

2. Suppose that $\omega \in \Omega$ and $X_0(\omega) = x$. Since \mathbf{X} is right continuous, there exists $\epsilon > 0$ such that $X_t(\omega) = x$ for $0 \leq t < \epsilon$ and hence $\tau(\omega) \geq \epsilon > 0$. So $\mathbb{P}(\tau > 0 \mid X_0 = x) = 1$.
3. Similarly, suppose that $\omega \in \Omega$ and that $X_0(\omega) = x$ and $X_{\tau(\omega)}(\omega) = y$. Since \mathbf{X} is right continuous, there exists $\epsilon > 0$ such that $X_t(\omega) = y$ for $\tau(\omega) \leq t < \tau(\omega) + \epsilon$. But by definition of $\tau(\omega)$, there exists $t \in (\tau(\omega), \tau(\omega) + \epsilon)$ with $X_t(\omega) \neq x$. Hence $\mathbb{P}(X_\tau \neq x \mid X_0 = x) = 1$.

There is actually a converse to part (b) that states that if \mathbf{X} has no instantaneous states, then there is a version of \mathbf{X} that is right continuous. From now on, we will assume that our Markov chains are right continuous with probability 1, and hence have no instantaneous states. On the other hand, absorbing states are perfectly reasonable and often do occur in applications. Finally, if the chain enters a stable state, it will stay there for a (proper) exponentially distributed time, and then leave.

The Jump Chain

Without instantaneous states, we can now construct a sequence of stopping times. Basically, we let τ_n denote the n th time that the chain changes state for $n \in \mathbb{N}_+$, unless the chain has previously been caught in an absorbing state. Here is the formal construction:

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S . Let $\tau_0 = 0$ and $\tau_1 = \inf\{t \in [0, \infty) : X_t \neq X_0\}$. Recursively, suppose that τ_n is defined for $n \in \mathbb{N}_+$. If $\tau_n = \infty$ let $\tau_{n+1} = \infty$. Otherwise, let

$$\tau_{n+1} = \inf\{t \in [\tau_n, \infty) : X_t \neq X_{\tau_n}\} \quad (16.15.9)$$

Let $M = \sup\{n \in \mathbb{N} : \tau_n < \infty\}$.

In the definition of M , of course, $\sup(\mathbb{N}) = \infty$, so M is the number of changes of state. If $M < \infty$, the chain was sucked into an absorbing state at time τ_M . Since we have ruled out instantaneous states, the sequence of random times is strictly increasing up until the (random) term M . That is, with probability 1, if $n \in \mathbb{N}$ and $\tau_n < \infty$ then $\tau_n < \tau_{n+1}$. Of course by construction, if $\tau_n = \infty$ then $\tau_{n+1} = \infty$. The increments $\tau_{n+1} - \tau_n$ for $n \in \mathbb{N}$ with $n < M$ are the times spent in the states visited by \mathbf{X} . The process at the random times when the state changes forms an embedded discrete-time Markov chain.

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S . Let $\{\tau_n : n \in \mathbb{N}\}$ denote the stopping times and M the random index, as defined above. For $n \in \mathbb{N}$, let $Y_n = X_{\tau_n}$ if $n \leq M$ and $Y_n = X_{\tau_M}$ if $n > M$. Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a (homogenous) discrete-time Markov chain on S , known as the *jump chain* of \mathbf{X} .

Proof

For $n \in \mathbb{N}$ let $\mathcal{G}_n = \sigma\{Y_0, Y_1, \dots, Y_n\}$, the σ -algebra of events for the process \mathbf{Y} , up to the discrete time n . Let $x \in S$. If x is stable, then given $Y_n = x$, the random times τ_n and τ_{n+1} are finite with probability 1. (Note that we cannot get to x from an absorbing state.) So

$$\mathbb{P}(Y_{n+1} = y \mid Y_n = x, \mathcal{G}_n) = \mathbb{P}(X_{\tau_{n+1}} = y \mid X_{\tau_n} = x, \mathcal{G}_n), \quad y \in S \quad (16.15.10)$$

But by the strong Markov property, given $X_{\tau_n} = x$, the chain starts over at time τ_n in state x , independent of $\mathcal{G}_n \subseteq \mathcal{F}_{\tau_n}$. Hence

$$\mathbb{P}(Y_{n+1} = y \mid Y_n = x, \mathcal{G}_n) = \mathbb{P}(X_\tau = y \mid X_0 = x), \quad y \in S \quad (16.15.11)$$

On the other hand, if x is an absorbing state, then by construction,

$$\mathbb{P}(Y_{n+1} = y \mid Y_n = x, \mathcal{G}_n) = I(x, y), \quad y \in S \quad (16.15.12)$$

where I is the identity matrix on S .

As noted in the proof, the one-step transition probability matrix Q for the jump chain \mathbf{Y} is given for $(x, y) \in S^2$ by

$$Q(x, y) = \begin{cases} \mathbb{P}(X_\tau = y \mid X_0 = x), & x \text{ stable} \\ I(x, y), & x \text{ absorbing} \end{cases} \quad (16.15.13)$$

where I is the identity matrix on S . Of course Q satisfies the usual properties of a probability matrix on S , namely $Q(x, y) \geq 0$ for $(x, y) \in S^2$ and $\sum_{y \in S} Q(x, y) = 1$ for $x \in S$. But Q satisfies another interesting property as well. Since the state actually *changes* at time τ starting in a stable state, we must have $Q(x, x) = 0$ if x is stable and $Q(x, x) = 1$ if x is absorbing.

Given the initial state, the holding time and the next state are independent.

If $x, y \in S$ and $t \in [0, \infty)$ then $\mathbb{P}(Y_1 = y, \tau_1 > t \mid Y_0 = x) = Q(x, y)e^{-\lambda(x)t}$

Proof

Suppose that x is a stable state, so that given $Y_0 = X_0 = x$, the stopping time $\tau_1 = \tau$ has a proper exponential distribution with parameter $\lambda(x) \in (0, \infty)$. Note that

$$\mathbb{P}(Y_1 = y, \tau_1 > t \mid Y_0 = x) = \mathbb{P}(X_\tau = y, \tau > t \mid X_0 = x) = \mathbb{P}(X_\tau = y \mid \tau > t, X_0 = x) \mathbb{P}(\tau > t \mid X_0 = x) \quad (16.15.14)$$

Note that if $X_0 = x$ and $\tau > t$ then $X_t = x$ also. By the Markov property, given $X_t = x$, the chain starts over at time t in state x , independent of $\{X_0 = x\}$ and $\{\tau > t\}$, both events in \mathcal{F}_t . Hence

$$\mathbb{P}(X_\tau = y \mid \tau > t, X_0 = x) = \mathbb{P}(X_\tau = y \mid X_t = x, \tau > t, X_0 = x) = \mathbb{P}(X_\tau = y \mid X_0 = x) = Q(x, y) \quad (16.15.15)$$

Of course $\mathbb{P}(\tau > t \mid X_0 = x) = e^{-\lambda(x)t}$.

If x is an absorbing state then $\mathbb{P}(\tau = \infty \mid X_0 = x) = 1$, $\mathbb{P}(Y_1 = x \mid Y_0 = x) = 1$, and $\lambda(x) = 0$. Hence

$$\mathbb{P}(Y_1 = y, \tau_1 > t \mid Y_0 = x) = I(x, y) = Q(x, y)e^{-\lambda(x)t} \quad (16.15.16)$$

The following theorem is a generalization. The changes in state and the holding times are independent, given the initial state.

Suppose that $n \in \mathbb{N}_+$ and that (x_0, x_1, \dots, x_n) is a sequence of stable states and (t_1, t_2, \dots, t_n) is a sequence in $[0, \infty)$. Then

$$\begin{aligned} & \mathbb{P}(Y_1 = x_1, \tau_1 > t_1, Y_2 = x_2, \tau_2 - \tau_1 > t_2, \dots, Y_n = x_n, \tau_n - \tau_{n-1} > t_n \mid Y_0 = x_0) \\ &= Q(x_0, x_1)e^{-\lambda(x_0)t_1} Q(x_1, x_2)e^{-\lambda(x_1)t_2} \dots Q(x_{n-1}, x_n)e^{-\lambda(x_{n-1})t_n} \end{aligned}$$

Proof

The proof is by induction, and the essence is captured in the case $n = 2$. So suppose that x_0, x_1, x_2 are stable states and $t_1, t_2 \in [0, \infty)$. Then

$$\begin{aligned} & \mathbb{P}(Y_1 = x_1, \tau_1 > t_1, Y_2 = x_2, \tau_2 - \tau_1 > t_2 \mid Y_0 = x_0) \\ &= \mathbb{P}(Y_2 = x_2, \tau_2 - \tau_1 > t_2 \mid X_0 = x, Y_1 = x_1, \tau_1 > t_1) \mathbb{P}(Y_1 = x_1, \tau_1 > t_1 \mid Y_0 = x_0) \end{aligned}$$

But $\mathbb{P}(Y_1 = x_1, \tau_1 > t_1 \mid Y_0 = x_0) = Q(x_0, x_1)e^{-\lambda(x_0)t_1}$ by the previous theorem. Next, by definition,

$$\mathbb{P}(Y_2 = x_2, \tau_2 - \tau_1 > t_2 \mid X_0 = x, Y_1 = x_1, \tau_1 > t_1) = \mathbb{P}(X_{\tau_2} = x_2, \tau_2 - \tau_1 > t_2 \mid X_0 = x_0, X_{\tau_1} = x_1, \tau_1 > t_1) \quad (16.15.17)$$

But by the strong Markov property, given $X_{\tau_1} = x_1$, the chain starts over at time τ_1 in state x , independent of the events $\{X_0 = x_0\}$ and $\{\tau_1 > t_1\}$ (both events in \mathcal{F}_{τ_1}). Hence using the previous theorem again,

$$\mathbb{P}(Y_2 = x_2, \tau_2 - \tau_1 > t_2 \mid X_0 = x, Y_1 = x_1, \tau_1 > t_1) = \mathbb{P}(X_\tau = x_2, \tau > t_2 \mid X_0 = x_1) = Q(x_1, x_2)e^{-\lambda(x_1)t_2} \quad (16.15.18)$$

Regularity

We now know quite a bit about the structure of a continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ (without instantaneous states). Once the chain enters a given state $x \in S$, the holding time in state x has an exponential distribution with parameter $\lambda(x) \in [0, \infty)$, after which the next state $y \in S$ is chosen, independently of the holding time, with probability $Q(x, y)$. However, we don't know everything about the chain. For the sequence $\{\tau_n : n \in \mathbb{N}\}$ defined above, let $\tau_\infty = \lim_{n \rightarrow \infty} \tau_n$, which exists in $(0, \infty]$ of course, since the sequence is increasing. Even though the holding time in a state is positive with probability 1, it's possible that $\tau_\infty < \infty$ with positive probability, in which case we know nothing about X_t

for $t \geq \tau_\infty$. The event $\{\tau_\infty < \infty\}$ is known as *explosion*, since it means that the \mathbf{X} makes infinitely many transitions before the finite time τ_∞ . While not as pathological as the existence of instantaneous states, explosion is still to be avoided in most applications.

A Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on S is *regular* if each of the following events has probability 1:

1. \mathbf{X} is right continuous.
2. $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$.

There is a simple condition on the exponential parameters and the embedded chain that is equivalent to condition (b).

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a right-continuous Markov chain on S with exponential parameter function λ and embedded chain $\mathbf{Y} = (Y_0, Y_1, \dots)$. Then $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$ with probability 1 if and only if $\sum_{n=0}^{\infty} 1/\lambda(Y_n) = \infty$ with probability 1.

Proof

Given $\mathbf{Y} = (y_0, y_1, \dots)$, the distribution of $\tau_\infty = \lim_{n \rightarrow \infty} \tau_n$ is the distribution of $T_\infty = \sum_{n=0}^{\infty} T_n$ where (T_0, T_1, \dots) are independent, and T_n has the exponential distribution with parameter $\lambda(y_n)$. Note that $\mathbb{E}(T_\infty) = \sum_{n=0}^{\infty} 1/\lambda(y_n)$. In the section on the exponential distribution, it's shown that $\mathbb{P}(T_\infty = \infty) = 1$ if and only if $\mathbb{E}(T_\infty) = \infty$.

If λ is bounded, then \mathbf{X} is regular.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S with exponential parameter function λ . If λ is bounded, then \mathbf{X} is regular.

Proof

Suppose that $\lambda(x) \leq r$ for $x \in S$, where $r \in (0, \infty)$. Then in particular, \mathbf{X} has no instantaneous states and so is right continuous. Moreover, $1/\lambda(x) \geq 1/r$ for $x \in S$ so $\sum_{n=0}^{\infty} 1/\lambda(Y_n) = \infty$ with probability 1, where as usual, $\mathbf{Y} = (Y_0, Y_1, \dots)$ is the jump chain of \mathbf{X} .

Here is another sufficient condition that is useful when the state space is infinite.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S with exponential parameter function $\lambda : S \rightarrow [0, \infty)$. Let $S_+ = \{x \in S : \lambda(x) > 0\}$. Then \mathbf{X} is regular if

$$\sum_{x \in S_+} \frac{1}{\lambda(x)} = \infty \quad (16.15.19)$$

Proof

By assumption, $\lambda(x) < \infty$ for $x \in S$, so there are no instantaneous states and so we can take \mathbf{X} to be right continuous. Next,

$$\sum_{n=0}^{\infty} \frac{1}{\lambda(Y_n)} = \sum_{n=0}^{\infty} \sum_{x \in S} \frac{1}{\lambda(x)} \mathbf{1}(Y_n = x) = \sum_{x \in S} \frac{1}{\lambda(x)} \sum_{n=0}^{\infty} \mathbf{1}(Y_n = x) = \sum_{x \in S} \frac{N_x}{\lambda(x)} \quad (16.15.20)$$

where $N_x = \sum_{n=0}^{\infty} \mathbf{1}(Y_n = x)$ is the number of times that the jump chain \mathbf{Y} is in state x . Suppose that $\sum_{x \in S_+} 1/\lambda(x) = \infty$. Note that it must be the case that S_+ , and hence S , is infinite. With probability 1, either \mathbf{Y} enters an absorbing state (a state $x \in S$ with $\lambda(x) = 0$), or $N_x = \infty$ for some $x \in S_+$, or $N_x \geq 1$ for infinitely many $x \in S_+$. In any case,

$$\sum_{n=0}^{\infty} \frac{1}{\lambda(Y_n)} = \sum_{x \in S} \frac{N_x}{\lambda(x)} = \infty \quad (16.15.21)$$

As a corollary, note that if S is finite then λ is bounded, so a continuous-time Markov chain on a finite state space is regular. So to review, if the exponential parameter function λ is finite, the chain \mathbf{X} has no instantaneous states. Even better, if λ is bounded or if the conditions in the last theorem are satisfied, then \mathbf{X} is regular. A continuous-time Markov chain with bounded exponential parameter function λ is called *uniform*, for reasons that will become clear in the next section on transition matrices. As we will see in later section, a uniform continuous-time Markov chain can be constructed from a discrete-time chain and an independent Poisson process. For the next result, recall that to say that \mathbf{X} has *left limits* with probability 1 means that the random function $t \mapsto X_t$ has limits from the left on $(0, \infty)$ with probability 1.

If $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is regular then \mathbf{X} has left limits with probability 1.

Proof

Suppose first that there are no absorbing states. Under the assumptions, with probability 1, $0 < \tau_n < \infty$ for each $n \in \mathbb{N}$ and $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$. Moreover, $X_t = Y_n$ for $t \in [\tau_n, \tau_{n+1})$ and $n \in \mathbb{N}$. So $t \mapsto X_t$ has left limits on $(0, \infty)$ with probability 1. The same basic argument works with absorbing states, except that possibly $\tau_{n+1} = \infty$.

Thus, our standard assumption will be that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S . For such a chain, the behavior of \mathbf{X} is completely determined by the exponential parameter function λ that governs the holding times, and the transition probability matrix Q of the

jump chain \mathbf{Y} . Conversely, when modeling real stochastic systems, we often start with λ and Q . It's then relatively straightforward to construct the continuous-time Markov chain that has these parameters. For simplicity, we will assume that there are no absorbing states. The inclusion of absorbing states is not difficult, but mucks up the otherwise elegant exposition.

Suppose that $\lambda : S \rightarrow (0, \infty)$ is bounded and that Q is a probability matrix on S with the property that $Q(x, x) = 0$ for every $x \in S$. The regular, continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with exponential parameter function λ and jump transition matrix Q can be constructed as follows:

1. First construct the jump chain $\mathbf{Y} = (Y_0, Y_1, \dots)$ having transition matrix Q .
2. Next, given $\mathbf{Y} = (x_0, x_1, \dots)$, the transition times (τ_1, τ_2, \dots) are constructed so that the holding times $(\tau_1, \tau_2 - \tau_1, \dots)$ are independent and exponentially distributed with parameters $(\lambda(x_0), \lambda(x_1), \dots)$
3. Again given $\mathbf{Y} = (x_0, x_1, \dots)$, define $X_t = x_0$ for $0 \leq t < \tau_1$ and for $n \in \mathbb{N}_+$, define $X_t = x_n$ for $\tau_n \leq t < \tau_{n+1}$.

Additional details

Using product sets and product measures, it's straightforward to construct a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with the following objects and properties:

1. $\mathbf{Y} = (Y_0, Y_1, \dots)$ is a Markov chain on S with transition matrix Q .
2. $\mathbf{T} = \{T_x : x \in S\}$ is a collection of independent random variables with values in $[0, \infty)$ such that T_x has the exponential distribution with parameter $\lambda(x)$ for each $x \in S$.
3. \mathbf{Y} and \mathbf{T} are independent.

Define $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ as follows: First, $\tau_1 = T_{Y_0}$ and $X_t = Y_0$ for $0 \leq t < \tau_1$. Recursively, if X_t is defined on $[0, \tau_n)$, let $\tau_{n+1} = \tau_n + T_{Y_n}$ and then let $X_t = Y_n$ for $\tau_n \leq t < \tau_{n+1}$. Since λ is bounded, $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$, so X_t is well defined for $t \in [0, \infty)$. By construction, $t \mapsto X_t$ is right continuous and has left limits. The Markov property holds by the memoryless property of the exponential distribution and the fact that \mathbf{Y} is a Markov chain. Finally, by construction, \mathbf{X} has exponential parameter function λ and jump chain \mathbf{Y} .

Often, particularly when S is finite, the essential structure of a standard, continuous-time Markov chain can be succinctly summarized with a graph.

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S , with exponential parameter function λ and embedded transition matrix Q . The *state graph* of \mathbf{X} is the graph with vertex set S and directed edge set $E = \{(x, y) \in S^2 : Q(x, y) > 0\}$. The graph is *labeled* as follows:

1. Each vertex $x \in S$ is labeled with the exponential parameter $\lambda(x)$.
2. Each edge $(x, y) \in E$ is labeled with the transition probability $Q(x, y)$.

So except for the labels on the vertices, the state graph of \mathbf{X} is the same as the state graph of the discrete-time jump chain \mathbf{Y} . That is, there is a directed edge from state x to state y if and only if the chain, when in x , can move to y after the random holding time in x . Note that the only loops in the state graph correspond to absorbing states, and for such a state there are no outward edges.

Let's return again to the [construction above](#) of a continuous-time Markov chain from the jump transition matrix Q and the exponential parameter function λ . Again for simplicity, assume there are no absorbing states. We assume that $Q(x, x) = 0$ for all $x \in S$, so that the state really does *change* at the transition times. However, if we drop this assumption, the construction still produces a continuous-time Markov chain, but with an altered jump transition matrix and exponential parameter function.

Suppose that Q is a transition matrix on $S \times S$ with $Q(x, x) < 1$ for $x \in S$, and that $\lambda : S \rightarrow (0, \infty)$ is bounded. The stochastic process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ [constructed above](#) from Q and λ is a regular, continuous-time Markov chain with exponential parameter function $\tilde{\lambda}$ and jump transition matrix \tilde{Q} given by

$$\begin{aligned}\tilde{\lambda}(x) &= \lambda(x)[1 - Q(x, x)], \quad x \in S \\ \tilde{Q}(x, y) &= \frac{Q(x, y)}{1 - Q(x, x)}, \quad (x, y) \in S^2, x \neq y\end{aligned}$$

Proof 1

As before, the fact that \mathbf{X} is a continuous-time Markov chain follows from the memoryless property of the exponential distribution and the Markov property of the jump chain \mathbf{Y} . By construction, $t \mapsto X_t$ is right continuous and has left limits. The main point, however, is that (τ_1, τ_2, \dots) is not necessarily the sequence of transition times, when the state actually changes. So we just need to determine the parameters. Suppose $X_0 = x \in S$ and let $\tau = \tau_1$ have the exponential distribution with parameter $\lambda(x)$, as in the construction. Let T denote the time when the state actually does change. For $t \in [0, \infty)$, the event $T > t$ can happen in two ways: either $\tau > t$ or $\tau = s$ for some $s \in [0, t]$, the chain jumps back into state x at time s , and the process then stays in x for a period of at least $t - s$. Thus let $F_x(t) = \mathbb{P}(T > t \mid X_0 = x)$. Taking the two cases, conditioning on τ , and using the Markov property gives

$$F_x(t) = e^{-\lambda(x)t} + \int_0^t \lambda(x) e^{-\lambda(x)s} Q(x, x) F_x(t-s) ds \quad (16.15.22)$$

Using the change of variables $u = t - s$ and simplifying gives

$$F_x(t) = e^{-\lambda(x)t} \left[1 + \lambda(x) Q(x, x) \int_0^t e^{\lambda(x)u} F_x(u) du \right] \quad (16.15.23)$$

Differentiating with respect to t then gives

$$F'_x(t) = -\lambda(x)[1 - Q(x, x)]F_x(t) \quad (16.15.24)$$

with the initial condition $F_x(0) = 1$. The solution of course is $F_x(t) = \exp\{-\lambda(x)[1 - Q(x, x)]t\}$ for $t \in [0, \infty)$. When the state does change, the new state $y \neq x$ is chosen with probability

$$\mathbb{P}(Y_1 = y \mid Y_0 = x, Y_1 \neq x) = \frac{Q(x, y)}{1 - Q(x, x)} \quad (16.15.25)$$

Proof 2

As in the first proof, we just need to determine the parameters. Given $X_0 = Y_0 = x$, the discrete time N when Y first changes state has the geometric distribution on \mathbb{N}_+ with success parameter $1 - Q(x, x)$. Hence the time until X actually changes state has the distribution of $T = \sum_{i=1}^N U_i$ where $U = (U_1, U_2, \dots)$ is a sequence of independent variables, each exponentially distributed with parameter $\lambda(x)$ and with U independent of N . In the section on the exponential distribution, it is shown that T also has the exponential distribution, but with parameter $\lambda(x)[1 - Q(x, x)]$. (The proof is simple using generating functions.) As in the first proof, when the state does change, the new state $y \neq x$ is chosen with probability

$$\mathbb{P}(Y_1 = y \mid Y_0 = x, Y_1 \neq x) = \frac{Q(x, y)}{1 - Q(x, x)} \quad (16.15.26)$$

This construction will be important in our study of chains subordinate to the Poisson process.

Transition Times

The structure of a regular Markov chain on S , as described above, can be explained purely in terms of a family of independent, exponentially distributed random variables. The main tools are some additional special properties of the exponential distribution, that we need to restate in the setting of our Markov chain. Our interest is in how the process evolves among the stable states until it enters an absorbing state (if it does). Once in an absorbing state, the chain stays there forever, so the behavior from that point on is trivial.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S , with exponential parameter function λ and transition probability matrix Q . Define $\mu(x, y) = \lambda(x)Q(x, y)$ for $(x, y) \in S^2$. Then

1. $\lambda(x) = \sum_{y \in S} \mu(x, y)$ for $x \in S$.
2. $Q(x, y) = \mu(x, y) / \lambda(x)$ if $(x, y) \in S^2$ and x is stable.

The main point is that the new parameters $\mu(x, y)$ for $(x, y) \in S^2$ determine the exponential parameters $\lambda(x)$ for $x \in S$, and the transition probabilities $Q(x, y)$ when $x \in S$ is stable and $y \in S$. Of course we know that if $\lambda(x) = 0$, so that x is absorbing, then $Q(x, x) = 1$. So in fact, the new parameters, as specified by the function μ , completely determine the old parameters, as specified by the functions λ and Q . But so what?

Consider the functions μ , λ , and Q as given in the previous result. Suppose that $T_{x,y}$ has the exponential distribution with parameter $\mu(x, y)$ for each $(x, y) \in S^2$ and that $\{T_{x,y} : (x, y) \in S^2\}$ is a set of independent random variables. Then

1. $T_x = \inf\{T_{x,y} : y \in S\}$ has the exponential distribution with parameter $\lambda(x)$ for $x \in S$.
2. $\mathbb{P}(T_x = T_{x,y}) = Q(x, y)$ for $(x, y) \in S^2$.

Proof

These are basic results proved in the section on the exponential distribution.

So here's how we can think of a regular, continuous-time Markov chain on S : There is a *timer* associated with each $(x, y) \in S^2$, set to the random time $T_{x,y}$. All of the timers function independently. When the chain enters state $x \in S$, the timers on (x, y) for $y \in S$ are started simultaneously. As soon as the first alarm goes off for a particular (x, y) , the chain immediately moves to state y , and the process repeats. Of course, if $\mu(x, y) = 0$ then $T_{x,y} = \infty$ with probability 1, so only the timers with $\lambda(x) > 0$ and $Q(x, y) > 0$ matter (these correspond to the non-loop edges in the state graph). In particular, if x is absorbing, then the timers on (x, y) are set to infinity for each y , and no alarm ever sounds.

The new collection of exponential parameters can be used to give an alternate version of the state graph. Again, the vertex set is S and the edge set is $E = \{(x, y) \in S^2 : Q(x, y) > 0\}$. But now each edge (x, y) is labeled with the exponential rate parameter $\mu(x, y)$. The exponential rate

parameters are closely related to the generator matrix, a matrix of fundamental importance that we will study in the next section.

Examples and Exercises

The Two-State Chain

The two-state chain is the simplest non-trivial, continuous-time Markov chain, but yet this chain illustrates many of the important properties of general continuous-time chains. So consider the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on the set of states $S = \{0, 1\}$, with transition rate $a \in [0, \infty)$ from 0 to 1 and transition rate $b \in [0, \infty)$ from 1 to 0.

The transition matrix Q for the embedded chain is given below. Draw the state graph in each case.

1. $Q = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ if $a > 0$ and $b > 0$, so that both states are stable.
2. $Q = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}$ if $a = 0$ and $b > 0$, so that a is absorbing and b is stable.
3. $Q = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$ if $a > 0$ and $b = 0$, so that a is stable and b is absorbing.
4. $Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ if $a = 0$ and $b = 0$, so that both states are absorbing.

We will return to the two-state chain in subsequent sections.

Computational Exercises

Consider the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1, 2\}$ with exponential parameter function $\lambda = (4, 1, 3)$ and embedded transition matrix

$$Q = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ \frac{1}{3} & \frac{2}{3} & 0 \end{bmatrix} \quad (16.15.27)$$

1. Draw the state graph and classify the states.
2. Find the matrix of transition rates.
3. Classify the jump chain in terms of recurrence and period.
4. Find the invariant distribution of the jump chain.

Answer

1. The edge set is $E = \{(0, 1), (0, 2), (1, 0), (2, 0), (2, 1)\}$ All states are stable.
2. The matrix of transition rates is

$$\begin{bmatrix} 0 & 2 & 2 \\ 1 & 0 & 0 \\ 1 & 2 & 0 \end{bmatrix} \quad (16.15.28)$$

3. The jump chain is irreducible, positive recurrent, and aperiodic.
4. The invariant distribution for the jump chain has PDF

$$f = \left[\frac{6}{14} \quad \frac{5}{14} \quad \frac{3}{14} \right] \quad (16.15.29)$$

Special Models

Read the introduction to chains subordinate to the Poisson process.

Read the introduction to birth-death chains.

Read the introduction to continuous-time queuing chains.

Read the introduction to continuous-time branching chains.

This page titled [16.15: Introduction to Continuous-Time Markov Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.16: Transition Matrices and Generators of Continuous-Time Chains

16. Transition Matrices and Generators of Continuous-Time Chains

Preliminaries

This is the second of the three introductory sections on continuous-time Markov chains. Thus, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and with state space (S, \mathcal{S}) . By the very meaning of *Markov chain*, the set of states S is countable and the σ -algebra \mathcal{S} is the collection of all subsets of S . So every subset of S is measurable, as is every function from S to another measurable space. Recall that \mathcal{S} is also the Borel σ algebra corresponding to the discrete topology on S . With this topology, every function from S to another topological space is continuous. Counting measure $\#$ is the natural measure on (S, \mathcal{S}) , so in the context of the general introduction, integrals over S are simply sums. Also, kernels on S can be thought of as matrices, with rows and sums indexed by S . The left and right kernel operations are generalizations of matrix multiplication.

A space of functions on S plays an important role. Let \mathcal{B} denote the collection of bounded functions $f : S \rightarrow \mathbb{R}$. With the usual pointwise definitions of addition and scalar multiplication, \mathcal{B} is a vector space. The *supremum norm* on \mathcal{B} is given by

$$\|f\| = \sup\{|f(x)| : x \in S\}, \quad f \in \mathcal{B} \quad (16.16.1)$$

Of course, if S is finite, \mathcal{B} is the set of all real-valued functions on S , and $\|f\| = \max\{|f(x)| : x \in S\}$ for $f \in \mathcal{B}$.

In the last section, we studied \mathbf{X} in terms of *when* and *how* the state changes. To review briefly, let $\tau = \inf\{t \in (0, \infty) : X_t \neq X_0\}$. Assuming that \mathbf{X} is right continuous, the Markov property of \mathbf{X} implies the memoryless property of τ , and hence the distribution of τ given $X_0 = x$ is exponential with parameter $\lambda(x) \in [0, \infty)$ for each $x \in S$. The assumption of right continuity rules out the pathological possibility that $\lambda(x) = \infty$, which would mean that x is an *instantaneous state* so that $\mathbb{P}(\tau = 0 | X_0 = x) = 1$. On the other hand, if $\lambda(x) \in (0, \infty)$ then x is a *stable state*, so that τ has a proper exponential distribution given $X_0 = x$ with $\mathbb{P}(0 < \tau < \infty | X_0 = x) = 1$. Finally, if $\lambda(x) = 0$ then x is an *absorbing state*, so that $\mathbb{P}(\tau = \infty | X_0 = x) = 1$. Next we define a sequence of stopping times: First $\tau_0 = 0$ and $\tau_1 = \tau$. Recursively, if $\tau_n < \infty$ then $\tau_n = \inf\{t > \tau_n : X_t \neq X_{\tau_n}\}$, while if $\tau_n = \infty$ then $\tau_{n+1} = \infty$. With $M = \sup\{n \in \mathbb{N} : \tau_n < \infty\}$ we define $Y_n = X_{\tau_n}$ if $n \in \mathbb{N}$ with $n \leq M$ and $Y_n = Y_M$ if $n \in \mathbb{N}$ with $n > M$. The sequence $\mathbf{Y} = (Y_0, Y_1, \dots)$ is a discrete-time Markov chain on S with one-step transition matrix Q given by $Q(x, y) = \mathbb{P}(X_\tau = y | X_0 = x)$ if $x, y \in S$ with x stable, and $Q(x, x) = 1$ if $x \in S$ is absorbing. Assuming that \mathbf{X} is *regular*, which means that $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$ with probability 1 (ruling out the *explosion* event of infinitely many transitions in finite time), the structure of \mathbf{X} is completely determined by the sequence of stopping times $\boldsymbol{\tau} = (\tau_0, \tau_1, \dots)$ and the discrete-time jump chain $\mathbf{Y} = (Y_0, Y_1, \dots)$. Analytically, the distribution \mathbf{X} is determined by the exponential parameter function λ and the one-step transition matrix Q of the jump chain.

In this section, we will study the Markov chain \mathbf{X} in terms of the transition matrices in continuous time and a fundamentally important matrix known as the *generator*. Naturally, the connections between the two points of view are particularly interesting.

The Transition Semigroup

Definition and basic Properties

The first part of our discussion is very similar to the treatment for a general Markov processes, except for simplifications caused by the discrete state space. We assume that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S .

The *transition probability matrix* P_t of \mathbf{X} corresponding to $t \in [0, \infty)$ is

$$P_t(x, y) = \mathbb{P}(X_t = y | X_0 = x), \quad (x, y) \in S^2 \quad (16.16.2)$$

In particular, $P_0 = I$, the identity matrix on S

Proof

The mapping $y \mapsto P_t(x, y)$ is the PDF of X_t given $X_0 = x$. Hence P_t is a probability matrix. That is, $P_t(x, y) \geq 0$ for $(x, y) \in S^2$ and $\sum_{y \in S} P_t(x, y) = 1$ for $x \in S$. Trivially, $P_0 = I$ by definition.

Note that since we are assuming that the Markov chain is homogeneous,

$$P_t(x, y) = \mathbb{P}(X_{s+t} = y | X_s = x), \quad (x, y) \in S^2 \quad (16.16.3)$$

for every $s, t \in [0, \infty)$. The *Chapman-Kolmogorov equation* given next is essentially yet another restatement of the Markov property. The equation is named for Andrei Kolmogorov and Sydney Chapman,

Suppose that $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ is the collection of transition matrices for the chain \mathbf{X} . Then $P_s P_t = P_{s+t}$ for $s, t \in [0, \infty)$. Explicitly,

$$P_{s+t}(x, z) = \sum_{y \in S} P_s(x, y) P_t(y, z), \quad x, z \in S \quad (16.16.4)$$

Proof

We condition on X_s .

$$P_{s+t}(x, z) = \mathbb{P}(X_{s+t} = z \mid X_0 = x) = \sum_{y \in S} \mathbb{P}(X_{s+t} = z \mid X_s = y, X_0 = x) \mathbb{P}(X_s = y \mid X_0 = x) \quad (16.16.5)$$

But by the Markov and time homogeneous properties,

$$\mathbb{P}(X_{s+t} = z \mid X_s = y, X_0 = x) = \mathbb{P}(X_{s+t} = z \mid X_s = y) = P_t(y, z) \quad (16.16.6)$$

Of course by definition, $\mathbb{P}(X_s = y \mid X_0 = x) = P_s(x, y)$. So the first displayed equation above becomes

$$P_{s+t}(x, y) = \sum_{y \in S} P_s(x, y) P_t(y, z) = P_s P_t(x, z) \quad (16.16.7)$$

Restated in another form of jargon, the collection $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ is a *semigroup* of probability matrices. The semigroup of transition matrices \mathbf{P} , along with the initial distribution, determine the finite-dimensional distributions of \mathbf{X} .

Suppose that X_0 has probability density function f . If $(t_1, t_2, \dots, t_n) \in [0, \infty)^n$ is a time sequence with $0 < t_1 < \dots < t_n$ and $(x_0, x_1, \dots, x_n) \in S^{n+1}$ is a state sequence, then

$$\mathbb{P}(X_0 = x_0, X_{t_1} = x_1, \dots, X_{t_n} = x_n) = f(x_0) P_{t_1}(x_0, x_1) P_{t_2 - t_1}(x_1, x_2) \cdots P_{t_n - t_{n-1}}(x_{n-1}, x_n) \quad (16.16.8)$$

Proof

To simplify the notation, we will just give the cases $n = 1$ and $n = 2$, which capture the essence of the proof. First suppose $x, y \in S$ and $t \in [0, \infty)$. Then

$$\mathbb{P}(X_0 = x, X_t = y) = \mathbb{P}(X_0 = x) \mathbb{P}(X_t = y \mid X_0 = x) = f(x) P_t(x, y) \quad (16.16.9)$$

Next suppose that $x, y, z \in S$ and $s, t \in [0, \infty)$ with $s < t$. Then

$$\mathbb{P}(X_0 = x, X_s = y, X_t = z) = \mathbb{P}(X_t = z \mid X_0 = x, X_s = y) \mathbb{P}(X_0 = x, X_s = y) \quad (16.16.10)$$

But by the Markov and time homogeneous properties, $\mathbb{P}(X_t = z \mid X_0 = x, X_s = y) = P_{t-s}(y, z)$. By the $n = 1$ case, $\mathbb{P}(X_0 = x, X_s = y) = f(x) P_s(x, y)$. Hence

$$\mathbb{P}(X_0 = x, X_s = y, X_t = z) = f(x) P_s(x, y) P_{t-s}(y, z) \quad (16.16.11)$$

As with any matrix on S , the transition matrices define left and right operations on functions which are generalizations of matrix multiplication. For a transition matrix, both have natural interpretations.

Suppose that $f : S \rightarrow \mathbb{R}$, and that either f is nonnegative or $f \in \mathcal{B}$. Then for $t \in [0, \infty)$,

$$P_t f(x) = \sum_{y \in S} P_t(x, y) f(y) = \mathbb{E}[f(X_t) \mid X_0 = x], \quad x \in S \quad (16.16.12)$$

The mapping $f \mapsto P_t f$ is a bounded, linear operator on \mathcal{B} and $\|P_t\| = 1$.

Proof

Since $P_t(x, \cdot)$ is the conditional probability density function of X_t given $X_0 = x$, it follows that $P_t f(x) = \mathbb{E}[f(X_t) \mid X_0 = x]$. The statement about $f \mapsto P_t f$ follows from general results on probability kernels.

If f is nonnegative and S is infinite, then it's possible that $P_t f(x) = \infty$. In general, the left operation of a positive kernel acts on positive measures on the state space. In the setting here, if μ is a positive (Borel) measure on (S, \mathcal{S}) , then the function $f : S \rightarrow [0, \infty)$ given by

$f(x) = \mu\{x\}$ for $x \in S$ is the density function of μ with respect to counting measure $\#$ on (S, \mathcal{S}) . This simply means that $\mu(A) = \sum_{x \in A} f(x)$ for $A \subseteq S$. Conversely, given $f : S \rightarrow [0, \infty)$, the set function $\mu(A) = \sum_{x \in A} f(x)$ for $A \subseteq S$ defines a positive measure on (S, \mathcal{S}) with f as its density function. So for the left operation of P_t , it's natural to consider only nonnegative functions.

If $f : S \rightarrow [0, \infty)$ then

$$fP_t(y) = \sum_{x \in S} f(x)P_t(x, y), \quad y \in S \quad (16.16.13)$$

If X_0 has probability density function f then X_t has probability density function fP_t .

Proof

If X_0 has PDF f , then conditioning gives

$$\mathbb{P}(X_t = y) = \sum_{x \in S} \mathbb{P}(X_t = y \mid X_0 = x) \mathbb{P}(X_0 = x) = \sum_{x \in S} P_t(x, y) f(x) = fP_t(y), \quad y \in S \quad (16.16.14)$$

More generally, if f is the density function of a positive measure μ on (S, \mathcal{S}) then fP_t is the density function of the measure μP_t , defined by

$$\mu P_t(A) = \sum_{x \in S} \mu\{x\} P_t(x, A) = \sum_{x \in S} f(x) P_t(x, A), \quad A \subseteq S \quad (16.16.15)$$

A function $f : S \rightarrow [0, \infty)$ is *invariant* for the Markov chain \mathbf{X} (or for the transition semigroup \mathbf{P}) if $fP_t = f$ for every $t \in [0, \infty)$.

It follows that if X_0 has an invariant probability density function f , then X_t has probability density function f for every $t \in [0, \infty)$, so \mathbf{X} is identically distributed. Invariant and limiting distributions are fundamentally important for continuous-time Markov chains.

Standard Semigroups

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain on S with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. Once again, continuity assumptions need to be imposed on \mathbf{X} in order to rule out strange behavior that would otherwise greatly complicate the theory. In terms of the transition semigroup \mathbf{P} , here is the basic assumption:

The transition semigroup \mathbf{P} is *standard* if $P_t(x, x) \rightarrow 1$ as $t \downarrow 0$ for each $x \in S$.

Since $P_0(x, x) = 1$ for $x \in S$, the standard assumption is clearly a continuity assumption. It actually implies much stronger smoothness properties that we will build up by stages.

If the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ is standard, then the function $t \mapsto P_t(x, y)$ is right continuous for each $(x, y) \in S^2$.

Proof

First note that if $(x, y) \in S^2$ with $x \neq y$ then $P_h(x, y) \leq 1 - P_h(x, x) \rightarrow 0$ as $h \downarrow 0$. Hence $P_h(x, y) \rightarrow I(x, y)$ as $h \downarrow 0$ for all $(x, y) \in S^2$. Suppose next that $t \in (0, \infty)$ and $(x, y) \in S^2$. By the semigroup property,

$$P_{t+h}(x, y) = P_t P_h(x, y) = \sum_{z \in S} P_t(x, z) P_h(z, y) \quad (16.16.16)$$

But $P_h(z, y) \rightarrow I(z, y)$ as $h \downarrow 0$ so by the bounded convergence theorem, $P_{t+h}(x, y) \rightarrow P_t(x, y)$ as $h \downarrow 0$.

Our next result connects one of the basic assumptions in the section on transition times and the embedded chain with the standard assumption here.

If the Markov chain \mathbf{X} has no instantaneous states then the transition semigroup \mathbf{P} is standard.

Proof

Given $X_0 = x \in S$ note that $\tau > t$ implies $X_t = x$. Hence

$$P_t(x, x) = \mathbb{P}(X_t = x \mid X_0 = x) \geq \mathbb{P}(\tau > t \mid X_0 = x) = e^{-\lambda(x)t} \quad (16.16.17)$$

Since \mathbf{X} has no instantaneous states, $0 \leq \lambda(x) < \infty$ so $e^{-\lambda(x)t} \rightarrow 1$ as $t \downarrow 0$.

Recall that the non-existence of instantaneous states is essentially equivalent to the right continuity of \mathbf{X} . So we have the nice result that if \mathbf{X} is right continuous, then so is \mathbf{P} . For the remainder of our discussion, we assume that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$, exponential function λ and one-step transition matrix Q for the jump chain. Our next result is the fundamental integral equations relating \mathbf{P} , λ , and Q .

For $t \in [0, \infty)$,

$$P_t(x, y) = I(x, y)e^{-\lambda(x)t} + \int_0^t \lambda(x)e^{-\lambda(x)s} Q P_{t-s}(x, y) ds, \quad (x, y) \in S^2 \quad (16.16.18)$$

Proof

If x is an absorbing state, then the equation trivially holds, since $\lambda(x) = 0$ and $P_t(x, y) = I(x, y)$. So suppose that x is a stable state, and as above, let $\tau = \inf\{t \in [0, \infty) : X_t \neq X_0\}$. Given $X_0 = x$, τ has a proper exponential distribution with parameter $\lambda(x) \in (0, \infty)$. Taking cases,

$$P_t(x, y) = \mathbb{P}(X_t = y \mid X_0 = x) = \mathbb{P}(X_t = y, \tau > t \mid X_0 = x) + \mathbb{P}(X_t = y, \tau \leq t \mid X_0 = x) \quad (16.16.19)$$

The first term on the right is 0 if $y \neq x$ and is $\mathbb{P}(\tau > t \mid X_0 = x) = e^{-\lambda(x)t}$ if $y = x$. In short,

$$\mathbb{P}(X_t = y, \tau > t \mid X_0 = x) = I(x, y)e^{-\lambda(x)t} \quad (16.16.20)$$

For the second term on the right in the displayed equation, we condition on τ and $Y_1 = X_\tau$. By a result in the last section on transition times and the embedded chain, the joint PDF of (τ, Y_1) at $s \in [0, \infty)$ and $z \in S$, given $X_0 = x$, is $\lambda(x)e^{-\lambda(x)s} Q(x, z)$ (continuous in time, discrete in space). Also, given $\tau = s \in [0, t]$ and $Y_1 = z \in S$, we can use the strong Markov property to “restart the clock” at s giving

$$\mathbb{P}(X_t = y \mid X_0 = x, \tau = s, Y_1 = z) = \mathbb{P}(X_{t-s} = y \mid X_0 = z) = P_{t-s}(z, y) \quad (16.16.21)$$

Putting the pieces together we have

$$\mathbb{P}(X_t = y, \tau \leq t \mid X_0 = x) = \int_0^t \lambda(x)e^{-\lambda(x)s} \sum_{z \in S} Q(x, z) P_{t-s}(z, y) ds = \int_0^t \lambda(x)e^{-\lambda(x)s} Q P_{t-s}(x, y) ds \quad (16.16.22)$$

We can now improve on the continuity result that we got earlier. First recall the leads to relation for the jump chain \mathbf{Y} : For $(x, y) \in S^2$, x leads to y if $Q^n(x, y) > 0$ for some $n \in \mathbb{N}$. So by definition, x leads to x for each $x \in S$, and for $(x, y) \in S^2$ with $x \neq y$, x leads to y if and only if the discrete-time chain starting in x eventually reaches y with positive probability.

For $(x, y) \in S^2$,

1. $t \mapsto P_t(x, y)$ is continuous.
2. If x leads to y then $P_t(x, y) > 0$ for every $t \in (0, \infty)$.
3. If x does not lead to y then $P_t(x, y) = 0$ for every $t \in (0, \infty)$.

Proof

For $t \in [0, \infty)$, we can use the change of variables $r = t - s$ in the fundamental integral equation to get

$$P_t(x, y) = I(x, y)e^{-\lambda(x)t} + \lambda(x)e^{-\lambda(x)t} \int_0^t e^{\lambda(x)r} Q P_r(x, y) dr, \quad (x, y) \in S^2 \quad (16.16.23)$$

1. In the displayed equation, $r \mapsto P_r(x, y)$ is right continuous for every $(x, y) \in S^2$, and hence by the bounded convergence theorem again, so is $r \mapsto Q P_r(x, y)$. Since the integrand in the displayed equation is bounded and right continuous, the integral is a continuous function of t . Hence $t \mapsto P_t(x, y)$ is continuous for $(x, y) \in S^2$.
2. For $x \in S$, note that $P_t(x, x) \geq e^{-\lambda(x)t} > 0$ for $t \in [0, \infty)$. If x leads to y and $x \neq y$ then there exists $n \in \mathbb{N}_+$ and $(x_1, x_2, \dots, x_{n-1}) \in S^{n-1}$ such that $Q(x, x_1) > 0, \dots, Q(x_{n-1}, y) > 0$. Then

$$P_t(x, y) = \mathbb{P}(X_t = y \mid X_0 = x) \geq \mathbb{P}(Y_1 = x_1, \dots, Y_{n-1} = x_{n-1}, Y_n = y, \tau_n \leq t < \tau_{n+1}) > 0 \quad (16.16.24)$$

3. This is clear from the definition of the embedded chain \mathbf{Y} .

Parts (b) and (c) are known as the *Lévy dichotomy*, named for Paul Lévy. It's possible to prove the Lévy dichotomy just from the semigroup property of \mathbf{P} , but this proof is considerably more complicated. In light of the dichotomy, the *leads to* relation clearly makes

sense for the continuous-time chain \mathbf{X} as well as the discrete-time embedded chain \mathbf{Y} .

The Generator Matrix

Definition and Basic Properties

In this discussion, we assume again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$, exponential parameter function λ and one-step transition matrix Q for the embedded jump chain. The fundamental integral equation above now implies that the transition probability matrix P_t is differentiable in t . The derivative at 0 is particularly important.

The matrix function $t \mapsto P_t$ has a (right) derivative at 0:

$$\frac{P_t - I}{t} \rightarrow G \text{ as } t \downarrow 0 \quad (16.16.25)$$

where the *infinitesimal generator matrix* G is given by $G(x, y) = -\lambda(x)I(x, y) + \lambda(x)Q(x, y)$ for $(x, y) \in S^2$.

Proof

As before the change of variables $r = t - s$ in the fundamental integral equation gives

$$P_t(x, y) = I(x, y)e^{-\lambda(x)t} + \lambda(x)e^{-\lambda(x)t} \int_0^t e^{\lambda(x)r} Q P_r(x, y) dr \quad (16.16.26)$$

The first term is clearly differentiable in t , and the second term is also differentiable in t since we now know that the integrand is a continuous function of r . The result then follows from standard calculus.

Note that $\lambda(x)Q(x, x) = 0$ for every $x \in S$, since $\lambda(x) = 0$ if x is absorbing, while $Q(x, x) = 0$ if x is stable. So $G(x, x) = -\lambda(x)$ for $x \in S$, and $G(x, y) = \lambda(x)Q(x, y)$ for $(x, y) \in S^2$ with $y \neq x$. Thus, the generator matrix G determines the exponential parameter function λ and the jump transition matrix Q , and thus determines the distribution of the Markov chain \mathbf{X} .

Given the generator matrix G of \mathbf{X} ,

1. $\lambda(x) = -G(x, x)$ for $x \in S$
2. $Q(x, y) = -G(x, y)/G(x, x)$ if $x \in S$ is stable and $y \in S - \{x\}$

The infinitesimal generator has a nice interpretation in terms of our discussion in the last section. Recall that when the chain first enters a stable state x , we set independent, exponentially distributed “timers” on (x, y) , for each $y \in S - \{x\}$. Note that $G(x, y)$ is the exponential parameter for the timer on (x, y) . As soon as an alarm sounds for a particular (x, y) , the chain moves to state y and the process continues.

The generator matrix G satisfies the following properties for every $x \in S$:

1. $G(x, x) \leq 0$
2. $\sum_{y \in S} G(x, y) = 0$

The matrix function $t \mapsto P_t$ is differentiable on $[0, \infty)$, and satisfies the *Kolmogorov backward equation*: $P'_t = GP_t$. Explicitly,

$$P'_t(x, y) = -\lambda(x)P_t(x, y) + \sum_{z \in S} \lambda(x)Q(x, z)P_t(z, y), \quad (x, y) \in S^2 \quad (16.16.27)$$

Proof

The proof is just like before, and follows from standard calculus and the integral equation

$$P_t(x, y) = I(x, y)e^{-\lambda(x)t} + \lambda(x)e^{-\lambda(x)t} \int_0^t e^{\lambda(x)r} Q P_r(x, y) dr \quad (16.16.28)$$

The backward equation is named for Andrei Kolmogorov. In continuous time, the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ can be obtained from the single, generator matrix G in a way that is reminiscent of the fact that in discrete time, the transition semigroup $\mathbf{P} = \{P^n : n \in \mathbb{N}\}$ can be obtained from the single, one-step matrix P . From a modeling point of view, we often start with the generator matrix G and then solve the backward equation, subject to the initial condition $P_0 = I$, to obtain the semigroup of transition matrices \mathbf{P} .

As with any matrix on S , the generator matrix G defines left and right operations on functions that are analogous to ordinary matrix multiplication. The right operation is defined for functions in \mathcal{B} .

If $f \in \mathcal{B}$ then Gf is given by

$$Gf(x) = -\lambda(x)f(x) + \sum_{y \in S} \lambda(x)Q(x, y)f(y), \quad x \in S \quad (16.16.29)$$

Proof

By definition,

$$Gf(x) = \sum_{y \in S} G(x, y)f(y) = -\lambda(x)f(x) + \sum_{y \in S - \{x\}} \lambda(x)Q(x, y)f(y) \quad (16.16.30)$$

In the second term, we can sum over all $y \in S$ since $\lambda(x) = 0$ if x is absorbing and $Q(x, x) = 0$ if x is stable. Note that Gf is well defined since

$$\sum_{y \in S - \{x\}} \lambda(x)Q(x, y) |f(y)| \leq \sum_{y \in S - \{x\}} \lambda(x)Q(x, y) \|f\| = \lambda(x) \|f\| \quad (16.16.31)$$

But note that Gf is not in \mathcal{B} unless $\lambda \in \mathcal{B}$. Without this additional assumption, G is a linear operator from the vector space \mathcal{B} of bounded functions from S to \mathbb{R} into the vector space of *all* functions from S to \mathbb{R} . We will return to this point in our next discussion.

Uniform Transition Semigroups

We can obtain stronger results for the generator matrix if we impose stronger continuity assumptions on \mathbf{P} .

The transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ is *uniform* if $P_t(x, x) \rightarrow 1$ as $t \downarrow 0$ uniformly in $x \in S$.

If \mathbf{P} is uniform, then the operator function $t \mapsto P_t$ is continuous on the vector space \mathcal{B} .

Proof

The statement means that for $f \in \mathcal{B}$, the function $t \mapsto P_t f$ is continuous with respect to the supremum norm on \mathcal{B} .

As usual, we want to look at this new assumption from different points of view.

The following are equivalent:

1. The transition semigroup \mathbf{P} is uniform.
2. The exponential parameter function λ is bounded.
3. The generator matrix G defines a bounded linear operator on \mathcal{B} .

Proof

From our remarks above we know that $\lambda \in \mathcal{B}$ if and only if the generator matrix G defines a bounded linear operator on \mathcal{B} . So we just need to show the equivalence of (a) and (b). If $\lambda \in \mathcal{B}$ then

$$P_t(x, x) = \mathbb{P}(X_t = x \mid X_0 = x) \geq \mathbb{P}(\tau > t \mid X_0 = x) = \exp[-\lambda(x)t] \geq \exp(-\|\lambda\|t) \quad (16.16.32)$$

The last term converges to 1 as $t \downarrow 0$ uniformly in x .

So when the equivalent conditions are satisfied, the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is also said to be uniform. As we will see in a later section, a uniform, continuous-time Markov chain can be constructed from a discrete-time Markov chain and an independent Poisson process. For a uniform transition semigroup, we have a companion to the backward equation.

Suppose that \mathbf{P} is a uniform transition semigroup. Then $t \mapsto P_t$ satisfies the *Kolmogorov forward equation* $P'_t = P_t G$. Explicitly,

$$P'_t(x, y) = -\lambda(y)P_t(x, y) + \sum_{z \in S} P_t(x, z)\lambda(z)Q(z, y), \quad (x, y) \in S^2 \quad (16.16.33)$$

The backward equation holds with more generality than the forward equation, since we only need the transition semigroup \mathbf{P} to be *standard* rather than *uniform*. It would seem that we need stronger conditions on λ for the forward equation to hold, for otherwise it's not even obvious that $\sum_{z \in S} P_t(x, z)\lambda(z)Q(z, y)$ is finite for $(x, y) \in S$. On the other hand, the forward equation is sometimes easier to

solve than the backward equation, and the assumption that λ is bounded is met in many applications (and of course holds automatically if S is finite).

As a simple corollary, the transition matrices and the generator matrix commute for a uniform semigroup: $P_t G = G P_t$ for $t \in [0, \infty)$. The forward and backward equations formally look like the differential equations for the exponential function. This actually holds with the operator exponential.

Suppose again that $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ is a uniform transition semigroup with generator G . Then

$$P_t = e^{tG} = \sum_{n=0}^{\infty} \frac{t^n}{n!} G^n, \quad t \in [0, \infty) \quad (16.16.34)$$

Proof

First e^{tG} is well defined as a bounded linear operator on \mathcal{B} for $t \in [0, \infty)$ (and hence also simply as a matrix), since G is a bounded linear operator on \mathcal{B} . Trivially $e^{0G} = I$, and by basic properties of the matrix exponential,

$$\frac{d}{dt} e^{tG} = G e^{tG}, \quad t \in (0, \infty) \quad (16.16.35)$$

It follows that $P_t = e^{tG}$ for $t \in [0, \infty)$.

We can characterize the generators of uniform transition semigroups. We just need the minimal conditions that the diagonal entries are nonpositive and the row sums are 0.

Suppose that G a matrix on S with $\|G\| < \infty$. Then G is the generator of a uniform transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ if and only if for every $x \in S$,

1. $G(x, x) \leq 0$
2. $\sum_{y \in S} G(x, y) = 0$

Proof

We know of course that if G is the generator of a transition semigroup, then conditions (a) and (b) hold. For the converse, we can use the previous result. Let

$$P_t = e^{tG} = \sum_{n=0}^{\infty} \frac{t^n}{n!} G^n, \quad t \in [0, \infty) \quad (16.16.36)$$

which makes sense since G is bounded in norm. Then $P_t(x, y) \geq 0$ for $(x, y) \in S^2$. By part (b), $\sum_{y \in S} G^n(x, y) = 0$ for every $x \in S$ and $n \in \mathbb{N}_+$, and hence $\sum_{y \in S} P_t(x, y) = \sum_{y \in S} I(x, y) = 1$ for $x \in S$. Finally, the semigroup property is a consequence of the law of exponents, which holds for the exponential of a matrix.

$$P_s P_t = e^{sG} e^{tG} = e^{(s+t)G} = P_{s+t} \quad (16.16.37)$$

Examples and Exercises

The Two-State Chain

Let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be the Markov chain on the set of states $S = \{0, 1\}$, with transition rate $a \in [0, \infty)$ from 0 to 1 and transition rate $b \in [0, \infty)$ from 1 to 0. This *two-state Markov chain* was studied in the previous section. To avoid the trivial case with both states absorbing, we will assume that $a + b > 0$.

The generator matrix is

$$G = \begin{bmatrix} -a & a \\ b & -b \end{bmatrix} \quad (16.16.38)$$

Show that for $t \in [0, \infty)$,

$$P_t = \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix} - \frac{1}{a+b} e^{-(a+b)t} \begin{bmatrix} -a & a \\ b & -b \end{bmatrix} \quad (16.16.39)$$

1. By solving the Kolmogorov backward equation.

2. By solving the Kolmogorov forward equation.
3. By computing $P_t = e^{tG}$.

You probably noticed that the forward equation is easier to solve because there is less coupling of terms than in the backward equation.

Define the probability density function f on S by $f(0) = \frac{b}{a+b}$, $f(1) = \frac{a}{a+b}$. Show that

1. $P_t \rightarrow \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix}$ as $t \rightarrow \infty$, the matrix with f in both rows.
2. $fP_t = f$ for all $t \in [0, \infty)$, so that f is invariant for P .
3. $fG = 0$.

Computational Exercises

Consider the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1, 2\}$ with exponential parameter function $\lambda = (4, 1, 3)$ and embedded transition matrix

$$Q = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ \frac{1}{3} & \frac{2}{3} & 0 \end{bmatrix} \quad (16.16.40)$$

1. Draw the state graph and classify the states.
2. Find the generator matrix G .
3. Find the transition matrix P_t for $t \in [0, \infty)$.
4. Find $\lim_{t \rightarrow \infty} P_t$.

Answer

1. The edge set is $E = \{(0, 1), (0, 2), (1, 0), (2, 0), (2, 1)\}$. All states are stable.
2. The generator matrix is

$$G = \begin{bmatrix} -4 & 2 & 2 \\ 1 & -1 & 0 \\ 1 & 2 & -3 \end{bmatrix} \quad (16.16.41)$$

3. For $t \in [0, \infty)$,

$$P_t = \frac{1}{15} \begin{bmatrix} 3 + 12e^{-5t} & 10 - 10e^{-3t} & 2 - 12e^{-5t} + 10e^{-3t} \\ 3 - 3e^{-5t} & 10 + 5e^{-3t} & 2 + 3e^{-5t} - 5e^{-3t} \\ 3 - 3e^{-5t} & 10 - 10e^{-3t} & 2 + 3e^{-5t} + 10e^{-3t} \end{bmatrix} \quad (16.16.42)$$

- 4.

$$P_t \rightarrow \frac{1}{15} \begin{bmatrix} 3 & 10 & 2 \\ 3 & 10 & 2 \\ 3 & 10 & 2 \end{bmatrix} \quad (16.16.43)$$

Special Models

Read the discussion of generator and transition matrices for chains subordinate to the Poisson process.

Read the discussion of the infinitesimal generator for continuous-time birth-death chains.

Read the discussion of the infinitesimal generator for continuous-time queuing chains.

Read the discussion of the infinitesimal generator for continuous-time branching chains.

This page titled [16.16: Transition Matrices and Generators of Continuous-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.17: Potential Matrices

Preliminaries

This is the third of the introductory sections on *continuous-time Markov chains*. So our starting point is a time-homogeneous Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and with discrete state space (S, \mathcal{S}) . Thus S is countable and \mathcal{S} is the power set of S , so every subset of S is measurable, as is every function from S into another measurable space. In addition, S is given the discrete topology so that \mathcal{S} can also be thought of as the Borel σ -algebra. Every function from S to another topological space is continuous. Counting measure $\#$ is the natural measure on (S, \mathcal{S}) , so in the context of the general introduction, integrals over S are simply sums. Also, kernels on S can be thought of as matrices, with rows and sums indexed by S , so the left and right kernel operations are generalizations of matrix multiplication. As before, let \mathcal{B} denote the collection of bounded functions $f : S \rightarrow \mathbb{R}$. With the usual pointwise definitions of addition and scalar multiplication, \mathcal{B} is a vector space. The *supremum norm* on \mathcal{B} is given by

$$\|f\| = \sup\{|f(x)| : x \in S\}, \quad f \in \mathcal{B} \quad (16.17.1)$$

Of course, if S is finite, \mathcal{B} is the set of all real-valued functions on S , and $\|f\| = \max\{|f(x)| : x \in S\}$ for $f \in \mathcal{B}$. The time space is $([0, \infty), \mathcal{T})$ where as usual, \mathcal{T} is the Borel σ -algebra on $[0, \infty)$ corresponding to the standard Euclidean topology. Lebesgue measure is the natural measure on $([0, \infty), \mathcal{T})$.

In our first point of view, we studied \mathbf{X} in terms of *when* and *how* the state changes. To review briefly, let $\tau = \inf\{t \in (0, \infty) : X_t \neq X_0\}$. Assuming that \mathbf{X} is right continuous, the Markov property of \mathbf{X} implies the memoryless property of τ , and hence the distribution of τ given $X_0 = x$ is exponential with parameter $\lambda(x) \in [0, \infty)$ for each $x \in S$. The assumption of right continuity rules out the pathological possibility that $\lambda(x) = \infty$, which would mean that x is an *instantaneous state* so that $\mathbb{P}(\tau = 0 | X_0 = x) = 1$. On the other hand, if $\lambda(x) \in (0, \infty)$ then x is a *stable state*, so that τ has a proper exponential distribution given $X_0 = x$ with $\mathbb{P}(0 < \tau < \infty | X_0 = x) = 1$. Finally, if $\lambda(x) = 0$ then x is an *absorbing state*, so that $\mathbb{P}(\tau = \infty | X_0 = x) = 1$. Next we define a sequence of stopping times: First $\tau_0 = 0$ and $\tau_1 = \tau$. Recursively, if $\tau_n < \infty$ then $\tau_n = \inf\{t > \tau_n : X_t \neq X_{\tau_n}\}$, while if $\tau_n = \infty$ then $\tau_{n+1} = \infty$. With $M = \sup\{n \in \mathbb{N} : \tau_n < \infty\}$ we define $Y_n = X_{\tau_n}$ if $n \in \mathbb{N}$ with $n \leq M$ and $Y_n = Y_M$ if $n \in \mathbb{N}$ with $n > M$. The sequence $\mathbf{Y} = (Y_0, Y_1, \dots)$ is a discrete-time Markov chain on S with one-step transition matrix Q given by $Q(x, y) = \mathbb{P}(X_\tau = y | X_0 = x)$ if $x, y \in S$ with x stable, and $Q(x, x) = 1$ if $x \in S$ is absorbing. Assuming that \mathbf{X} is *regular*, which means that $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$ with probability 1 (ruling out the *explosion* event of infinitely many transitions in finite time), the structure of \mathbf{X} is completely determined by the sequence of stopping times $\boldsymbol{\tau} = (\tau_0, \tau_1, \dots)$ and the embedded discrete-time *jump chain* $\mathbf{Y} = (Y_0, Y_1, \dots)$. Analytically, the distribution \mathbf{X} is determined by the exponential parameter function λ and the one-step transition matrix Q of the jump chain.

In our second point of view, we studied \mathbf{X} in terms of the collection of *transition matrices* $\mathbf{P} = \{P_t : t \in [0, \infty)\}$, where for $t \in [0, \infty)$,

$$P_t(x, y) = \mathbb{P}(X_t = y | X_0 = x), \quad (x, y) \in S^2 \quad (16.17.2)$$

The Markov and time-homogeneous properties imply the *Chapman-Kolmogorov* equations $P_s P_t = P_{s+t}$ for $s, t \in [0, \infty)$, so that \mathbf{P} is a *semigroup* of transition matrices. The semigroup \mathbf{P} , along with the initial distribution of X_0 , completely determines the distribution of \mathbf{X} . For a regular Markov chain \mathbf{X} , the fundamental integral equation connecting the two points of view is

$$P_t(x, y) = I(x, y)e^{-\lambda(x)t} + \int_0^t \lambda(x)e^{-\lambda(x)s} Q P_{t-s}(x, y) ds, \quad (x, y) \in S^2 \quad (16.17.3)$$

which is obtained by conditioning on τ and X_τ . It then follows that the matrix function $t \mapsto P_t$ is differentiable, with the derivative satisfying the *Kolmogorov backward equation* $P'_t = G P_t$ where the *generator matrix* G is given by

$$G(x, y) = -\lambda(x)I(x, y) + \lambda(x)Q(x, y), \quad (x, y) \in S^2 \quad (16.17.4)$$

If the exponential parameter function λ is bounded, then the transition semigroup \mathbf{P} is *uniform*, which leads to stronger results. The generator G is a bounded operator on \mathcal{B} , the backward equation holds as well as a companion forward equation $P'_t = P_t G$, as operators on \mathcal{B} (so with respect to the supremum norm rather than just pointwise). Finally, we can represent the transition matrix as an exponential: $P_t = e^{tG}$ for $t \in [0, \infty)$.

In this section, we study the Markov chain \mathbf{X} in terms of a family of matrices known as *potential matrices*. This is the least intuitive of the three points of view, but analytically one of the best approaches. Essentially, the potential matrices are transforms of the transition matrices.

Basic Theory

We assume again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. Our first discussion closely parallels the general theory, except for simplifications caused by the discrete state space.

Definitions and Properties

For $\alpha \in [0, \infty)$, the α -*potential matrix* U_α of \mathbf{X} is defined as follows:

$$U_\alpha(x, y) = \int_0^\infty e^{-\alpha t} P_t(x, y) dt, \quad (x, y) \in S^2 \quad (16.17.5)$$

1. The special case $U = U_0$ is simply the *potential matrix* of \mathbf{X} .
2. For $(x, y) \in S^2$, $U(x, y)$ is the expected amount of time that \mathbf{X} spends in y , starting at x .
3. The family of matrices $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$ is known as the *resolvent* of \mathbf{X} .

Proof

Since $t \mapsto P_t(x, y)$ is continuous, $U_\alpha(x, y)$ makes sense for $(x, y) \in S^2$. The interpretation of $U(x, y)$ involves an interchange of integrals:

$$U(x, y) = \int_0^\infty P_t(x, y) dt = \int_0^\infty \mathbb{E}[\mathbf{1}(X_t = y) \mid X_0 = x] dt = \mathbb{E}\left(\int_0^\infty \mathbf{1}(X_t = y) dt \mid X_0 = x\right) \quad (16.17.6)$$

The inside integral is the Lebesgue measure of $\{t \in [0, \infty) : X_t = y\}$.

It's quite possible that $U(x, y) = \infty$ for some $(x, y) \in S^2$, and knowing when this is the case is of considerable interest. If $f : S \rightarrow \mathbb{R}$ and $\alpha \geq 0$, then giving the right operation in its many forms,

$$\begin{aligned} U_\alpha f(x) &= \sum_{y \in S} U_\alpha(x, y) f(y) = \int_0^\infty e^{-\alpha t} P_t f(x) dt \\ &= \int_0^\infty e^{-\alpha t} \sum_{y \in S} P_t(x, y) f(y) dt = \int_0^\infty e^{-\alpha t} \mathbb{E}[f(X_t) \mid X_0 = x] dt, \quad x \in S \end{aligned}$$

assuming, as always, that the sums and integrals make sense. This will be the case in particular if f is nonnegative (although ∞ is a possible value), or as we will now see, if $f \in \mathcal{B}$ and $\alpha > 0$.

If $\alpha > 0$, then $U_\alpha(x, S) = \frac{1}{\alpha}$ for all $x \in S$.

Proof

For $x \in S$,

$$U_\alpha(x, S) = \int_0^\infty e^{-\alpha t} P_t(x, S) dt = \int_0^\infty e^{-\alpha t} dt = \frac{1}{\alpha} \quad (16.17.7)$$

It follows that for $\alpha \in (0, \infty)$, the right potential operator U_α is a bounded, linear operator on \mathcal{B} with $\|U_\alpha\| = \frac{1}{\alpha}$. It also follows that αU_α is a probability matrix. This matrix has a nice interpretation.

If $\alpha > 0$ then $\alpha U_\alpha(x, \cdot)$ is the conditional probability density function of X_T given $X_0 = x$, where T is independent of \mathbf{X} and has the exponential distribution on $[0, \infty)$ with parameter α .

Proof

Suppose that $(x, y) \in S^2$. The random time T has PDF $f(t) = \alpha e^{-\alpha t}$ for $t \in [0, \infty)$. Hence, conditioning on T gives

$$\mathbb{P}(X_T = y \mid X_0 = x) = \int_0^\infty \alpha e^{-\alpha t} \mathbb{P}(X_T = y \mid T = t, X_0 = x) dt \quad (16.17.8)$$

But by the substitution rule and the assumption of independence,

$$\mathbb{P}(X_T = y \mid T = t, X_0 = x) = \mathbb{P}(X_t = y \mid T = t, X_0 = x) = \mathbb{P}(X_t = y \mid X_0 = x) = P_t(x, y) \quad (16.17.9)$$

Substituting gives

$$\mathbb{P}(X_T = y \mid X_0 = x) = \int_0^\infty \alpha e^{-\alpha t} P_t(x, y) dt = \alpha U_\alpha(x, y) \quad (16.17.10)$$

So αU_α is a transition probability matrix, just as P_t is a transition probability matrix, but corresponding to the *random* time T (with $\alpha \in (0, \infty)$ as a parameter), rather than the deterministic time $t \in [0, \infty)$. The potential matrix can also be interpreted in economic terms. Suppose that we receive money at a rate of one unit per unit time whenever the process \mathbf{X} is in a particular state $y \in S$. Then $U(x, y)$ is the expected total amount of money that we receive, starting in state $x \in S$. But money that we receive later is of less value to us *now* than money that we will receive sooner. Specifically, suppose that one monetary unit at time $t \in [0, \infty)$ has a present value of $e^{-\alpha t}$ where $\alpha \in (0, \infty)$ is the *inflation factor* or *discount factor*. Then $U_\alpha(x, y)$ is the total, expected, discounted amount that we receive, starting in $x \in S$. A bit more generally, suppose that $f \in \mathcal{B}$ and that $f(y)$ is the reward (or cost, depending on the sign) per unit time that we receive when the process is in state $y \in S$. Then $U_\alpha f(x)$ is the expected, total, discounted reward, starting in state $x \in S$.

$\alpha U_\alpha \rightarrow I$ as $\alpha \rightarrow \infty$.

Proof

Note first that with a change of variables $s = \alpha t$,

$$\alpha U_\alpha = \int_0^\infty \alpha e^{-\alpha t} P_t dt = \int_0^\infty e^{-s} P_{s/\alpha} ds \quad (16.17.11)$$

But for $s \in [0, \infty)$, $s/\alpha \rightarrow 0$ and hence $P_{s/\alpha} \rightarrow I$ as $\alpha \rightarrow \infty$. The result then follows from the dominated convergence theorem.

If $f : S \rightarrow [0, \infty)$, then giving the left potential operation in its various forms,

$$\begin{aligned} fU_\alpha(y) &= \sum_{x \in S} f(x) U_\alpha(x, y) = \int_0^\infty e^{-\alpha t} f P_t(y) dt \\ &= \int_0^\infty e^{-\alpha t} \left[\sum_{x \in S} f(x) P_t(x, y) \right] dt = \int_0^\infty e^{-\alpha t} \left[\sum_{x \in S} f(x) \mathbb{P}(X_t = y) \right] dt, \quad y \in S \end{aligned}$$

In particular, suppose that $\alpha > 0$ and that f is the probability density function of X_0 . Then $f P_t$ is the probability density function of X_t for $t \in [0, \infty)$, and hence from the last result, $\alpha f U_\alpha$ is the probability density function of X_T , where again, T is independent of \mathbf{X} and has the exponential distribution on $[0, \infty)$ with parameter α . The family of potential kernels gives the same information as the family of transition kernels.

The resolvent $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$ completely determines the family of transition kernels $\mathbf{P} = \{P_t : t \in (0, \infty)\}$.

Proof

Note that for $(x, y) \in S^2$, the function $\alpha \mapsto U_\alpha(x, y)$ on $(0, \infty)$ is the Laplace transform of the function $t \mapsto P_t(x, y)$ on $[0, \infty)$. The Laplace transform of a continuous function determines the function uniquely.

Although not as intuitive from a probability view point, the potential matrices are in some ways nicer than the transition matrices because of additional smoothness. In particular, the resolvent $\{U_\alpha : \alpha \in [0, \infty)\}$, along with the initial distribution, completely determine the finite dimensional distributions of the Markov chain \mathbf{X} . The potential matrices commute with the transition matrices and with each other.

Suppose that $\alpha, \beta, t \in [0, \infty)$. Then

1. $P_t U_\alpha = U_\alpha P_t = \int_0^\infty e^{-\alpha s} P_{s+t} ds$
2. $U_\alpha U_\beta = U_\beta U_\alpha = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_{s+t} ds dt$

Proof

The interchanges of matrix multiplication and integrals below are interchanges of sums and integrals, and are justified since the underlying integrands are nonnegative. The other tool used is the semigroup property of $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. You may want to write out the proofs explicitly to convince yourself

1. First,

$$U_\alpha P_t = \left(\int_0^\infty e^{-\alpha s} P_s ds \right) P_t = \int_0^\infty e^{-\alpha s} P_s P_t ds = \int_0^\infty e^{-\alpha s} P_{s+t} ds \quad (16.17.12)$$

Similarly

$$P_t U_\alpha = P_t \int_0^\infty e^{-\alpha s} P_s ds = \int_0^\infty e^{-\alpha s} P_t P_s ds = \int_0^\infty e^{-\alpha s} P_{s+t} ds \quad (16.17.13)$$

2. First

$$U_\alpha U_\beta = \left(\int_0^\infty e^{-\alpha s} P_s ds \right) \left(\int_0^\infty e^{-\beta t} P_t dt \right) = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_s P_t ds dt = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_{s+t} ds dt \quad (16.17.14)$$

The other direction is similar.

The equations above are matrix equations, and so hold pointwise. The same identities hold for the right operators on the space \mathcal{B} under the additional restriction that $\alpha > 0$ and $\beta > 0$. The fundamental equation that relates the potential kernels, known as the *resolvent equation*, is given in the next theorem:

If $\alpha, \beta \in [0, \infty)$ with $\alpha \leq \beta$ then $U_\alpha = U_\beta + (\beta - \alpha)U_\alpha U_\beta$.

Proof

If $\alpha = \beta$ the equation is trivial, so assume $\alpha < \beta$. From the previous result,

$$U_\alpha U_\beta = \int_0^\infty \int_0^\infty e^{-\alpha s} e^{-\beta t} P_{s+t} dt ds \quad (16.17.15)$$

The transformation $u = s + t, v = s$ maps $[0, \infty)^2$ one-to-one onto $\{(u, v) \in [0, \infty)^2 : u \geq v\}$. The inverse transformation is $s = v, t = u - v$ with Jacobian -1 . Hence we have

$$\begin{aligned}
 U_\alpha U_\beta &= \int_0^\infty \int_0^u e^{-\alpha v} e^{-\beta(u-v)} P_u dv du = \int_0^\infty \left(\int_0^u e^{(\beta-\alpha)v} dv \right) e^{-\beta u} P_u du \\
 &= \frac{1}{\beta-\alpha} \int_0^\infty \left[e^{(\beta-\alpha)u} - 1 \right] e^{-\beta u} P_u du \\
 &= \frac{1}{\beta-\alpha} \left(\int_0^\infty e^{-\alpha u} P_u du - \int_0^\infty e^{-\beta u} P_u du \right) = \frac{1}{\beta-\alpha} (U_\alpha - U_\beta)
 \end{aligned}$$

Simplifying gives the result. Note that U_β is finite since $\beta > 0$, so we don't have to worry about the dreaded indeterminate form $\infty - \infty$.

The equation above is a matrix equation, and so holds pointwise. The same identity holds for the right potential operators on the space \mathcal{B} , under the additional restriction that $\alpha > 0$.

Connections with the Generator

Once again, assume that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular Markov chain on S with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$, infinitesimal generator G , resolvent $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$, exponential parameter function λ , and one-step transition matrix Q for the jump chain. There are fundamental connections between the potential U_α and the generator matrix G , and hence between U_α and the function λ and the matrix Q .

If $\alpha \in (0, \infty)$ then $I + GU_\alpha = \alpha U_\alpha$. In terms of λ and Q ,

$$U_\alpha(x, y) = \frac{1}{\alpha + \lambda(x)} I(x, y) + \frac{\lambda(x)}{\alpha + \lambda(x)} Q U_\alpha(x, y), \quad (x, y) \in S^2 \quad (16.17.16)$$

Proof 1

First,

$$GU_\alpha = G \int_0^\infty e^{-\alpha t} P_t dt = \int_0^\infty e^{-\alpha t} G P_t dt = \int_0^\infty e^{-\alpha t} P'_t dt \quad (16.17.17)$$

Passing G through the integrand is justified since $G P_t(x, y)$ is a sum with just one negative term for $(x, y) \in S^2$. The second identity in the displayed equation follows from the backward equation. Integrating by parts then gives

$$GU_\alpha = e^{-\alpha t} P_t \Big|_0^\infty + \int_0^\infty \alpha e^{-\alpha t} P_t dt = -I + \alpha U_\alpha \quad (16.17.18)$$

Proof 2

This proof use the fundamental integral equation relating \mathbf{P} , λ , and Q as well as the definition of U_α and interchanges of integrals. The interchange is justified since the integrand is nonnegative. So for $\alpha \in [0, \infty)$ and $(x, y) \in S^2$,

$$\begin{aligned}
 U_\alpha(x, y) &= \int_0^\infty e^{-\alpha t} P_t(x, y) dt \\
 &= \int_0^\infty e^{-\alpha t} \left[e^{-\lambda(x)t} I(x, y) + \lambda(x) e^{-\lambda(x)t} \int_0^t e^{\lambda(x)r} Q P_r(x, y) dr \right] dt \\
 &= I(x, y) \int_0^\infty e^{-[\alpha+\lambda(x)]t} dt + \lambda(x) \int_0^\infty \int_0^t e^{-[\alpha+\lambda(x)]t} e^{\lambda(x)r} Q P_r(x, y) dr dt \\
 &= \frac{1}{\alpha + \lambda(x)} I(x, y) + \lambda(x) \int_0^\infty \int_r^\infty e^{-[\alpha+\lambda(x)]t} e^{\lambda(x)r} Q P_r(x, y) dt dr \\
 &= \frac{1}{\alpha + \lambda(x)} I(x, y) + \frac{\lambda(x)}{\alpha + \lambda(x)} \int_0^\infty e^{-[\alpha+\lambda(x)]r} e^{\lambda(x)r} Q P_r(x, y) dr \\
 &= \frac{1}{\alpha + \lambda(x)} I(x, y) + \frac{\lambda(x)}{\alpha + \lambda(x)} \int_0^\infty e^{-\alpha r} Q P_r(x, y) dr = \frac{1}{\alpha + \lambda(x)} I(x, y) + \frac{\lambda(x)}{\alpha + \lambda(x)} Q U_\alpha(x, y)
 \end{aligned}$$

Proof 3

Recall that $\alpha U_\alpha(x, y) = \mathbb{P}(X_T = y \mid X_0 = x)$ where T is independent of \mathbf{X} and has the exponential distribution with parameter α . This proof works by conditioning on whether $T < \tau_1$ or $T \geq \tau_1$:

$$\alpha U_\alpha(x, y) = \mathbb{P}(X_T = y \mid X_0 = x, T < \tau_1) \mathbb{P}(T < \tau_1 \mid X_0 = x) + \mathbb{P}(X_T = y \mid X_0 = x, T \geq \tau_1) \mathbb{P}(T \geq \tau_1 \mid X_0 = x) \quad (16.17.19)$$

But $X_0 = x$ and $T < \tau_1$ imply $X_T = x$ so $\mathbb{P}(X_T = y \mid X_0 = x, T < \tau_1) = I(x, y)$. And by a basic property of independent exponential variables that we have seen many times before,

$$\mathbb{P}(T < \tau_1 \mid X_0 = x) = \frac{\alpha}{\alpha + \lambda(x)} \quad (16.17.20)$$

Next, for the first factor in the second term of the displayed equation, we condition on X_{τ_1} :

$$\mathbb{P}(X_T = y \mid X_0 = x, T \geq \tau_1) = \sum_{z \in S} \mathbb{P}(X_T = y \mid X_0 = x, X_{\tau_1} = z, T \geq \tau_1) \mathbb{P}(X_{\tau_1} = z \mid X_0 = x, T \geq \tau_1) \quad (16.17.21)$$

But by the strong Markov property, given $X_{\tau_1} = z$, we can restart the clock at time τ_1 in state z . Moreover, by the memoryless property and independence, the distribution of $T - \tau_1$ given $T \geq \tau_1$ is the same as the distribution of T , mainly exponential with parameter α . It follows that

$$\mathbb{P}(X_T = y \mid X_0 = x, X_{\tau_1} = z, T \geq \tau_1) = \mathbb{P}(X_T = y \mid X_0 = z) = \alpha U_\alpha(z, y) \quad (16.17.22)$$

Also, X_{τ_1} is independent of τ_1 and T so

$$\mathbb{P}(X_{\tau_1} = z \mid X_0 = x, T \geq \tau_1) = Q(x, z) \quad (16.17.23)$$

Finally using the basic property of exponential distributions again,

$$\mathbb{P}(T \geq \tau_1 \mid X_0 = x) = \frac{\lambda(x)}{\alpha + \lambda(x)} \quad (16.17.24)$$

Putting all the pieces together we have

$$\alpha U_\alpha(x, y) = \frac{\alpha}{\alpha + \lambda(x)} I(x, y) = \frac{\lambda(x)}{\alpha + \lambda(x)} \sum_{z \in S} Q(x, z) \alpha U_\alpha(z, y) = \frac{\alpha}{\alpha + \lambda(x)} I(x, y) + \frac{\lambda(x)}{\alpha + \lambda(x)} Q \alpha U_\alpha(x, y) \quad (16.17.25)$$

As before, we can get stronger results if we assume that λ is bounded, or equivalently, the transition semigroup \mathbf{P} is uniform.

Suppose that λ is bounded and $\alpha \in (0, \infty)$. Then as operators on \mathcal{B} (and hence also as matrices),

1. $I + GU_\alpha = \alpha U_\alpha$
2. $I + U_\alpha G = \alpha U_\alpha$

Proof

Since λ is bounded, G is a bounded operator on \mathcal{B} . The proof of (a) then proceeds as before. For (b) we know from the forward and backward equations that $GP_t = P_t G$ for $t \in [0, \infty)$ and hence $GU_\alpha = U_\alpha G$ for $\alpha \in (0, \infty)$.

As matrices, the equation in (a) holds with more generality than the equation in (b), much as the Kolmogorov backward equation holds with more generality than the forward equation. Note that

$$U_\alpha G(x, y) = \sum_{z \in S} U_\alpha(x, z) G(z, y) = -\lambda(y) U_\alpha(x, y) + \sum_{z \in S} U_\alpha(x, z) \lambda(z) Q(z, y), \quad (x, y) \in S^2 \quad (16.17.26)$$

If λ is unbounded, it's not clear that the second sum is finite.

Suppose that λ is bounded and $\alpha \in (0, \infty)$. Then as operators on \mathcal{B} (and hence also as matrices),

1. $U_\alpha = (\alpha I - G)^{-1}$
2. $G = \alpha I - U_\alpha^{-1}$

Proof

1. This follows immediately from the previous result, since $U_\alpha(\alpha I - G) = I$ and $(\alpha I - G)U_\alpha = I$
2. This follows from (a): $\alpha I - G = U_\alpha^{-1}$ so $G = \alpha I - U_\alpha^{-1}$

So the potential operator U_α and the generator G have a simple, elegant inverse relationship. Of course, these results hold in particular if S is finite, so that all of the various matrices really are matrices in the elementary sense.

Examples and Exercises

The Two-State Chain

Let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be the Markov chain on the set of states $S = \{0, 1\}$, with transition rate $a \in [0, \infty)$ from 0 to 1 and transition rate $b \in [0, \infty)$ from 1 to 0. To avoid the trivial case with both states absorbing, we will assume that $a + b > 0$. The first two results below are a review from the previous two sections.

The generator matrix G is

$$G = \begin{bmatrix} -a & a \\ b & -b \end{bmatrix} \quad (16.17.27)$$

The transition matrix at time $t \in [0, \infty)$ is

$$P_t = \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix} - \frac{1}{a+b} e^{-(a+b)t} \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}, \quad t \in [0, \infty) \quad (16.17.28)$$

Now we can find the potential matrix in two ways.

For $\alpha \in (0, \infty)$, show that the potential matrix U_α is

$$U_\alpha = \frac{1}{\alpha(a+b)} \begin{bmatrix} b & a \\ b & a \end{bmatrix} - \frac{1}{(\alpha+a+b)(a+b)} \begin{bmatrix} -a & a \\ b & -b \end{bmatrix} \quad (16.17.29)$$

1. From the definition.
2. From the relation $U_\alpha = (\alpha I - G)^{-1}$.

Computational Exercises

Consider the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1, 2\}$ with exponential parameter function $\lambda = (4, 1, 3)$ and jump transition matrix

$$Q = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ \frac{1}{3} & \frac{2}{3} & 0 \end{bmatrix} \quad (16.17.30)$$

1. Draw the state graph and classify the states.
2. Find the generator matrix G .
3. Find the potential matrix U_α for $\alpha \in (0, \infty)$.

Answer

1. The edge set is $E = \{(0, 1), (0, 2), (1, 0), (2, 0), (2, 1)\}$. All states are stable.
2. The generator matrix is

$$G = \begin{bmatrix} -4 & 2 & 2 \\ 1 & -1 & 0 \\ 1 & 2 & -3 \end{bmatrix} \quad (16.17.31)$$

3. For $\alpha \in (0, \infty)$,

$$U_\alpha = (\alpha I - G)^{-1} = \frac{1}{15\alpha + 8\alpha^2 + \alpha^3} \begin{bmatrix} 3 + 4\alpha + \alpha^2 & 10 + 2\alpha & 2 + 2\alpha \\ 3 + \alpha & 10 + 7\alpha + \alpha^2 & 2 \\ 3 + \alpha & 10 + 2\alpha & 2 + 5\alpha + \alpha^2 \end{bmatrix} \quad (16.17.32)$$

Special Models

Read the discussion of potential matrices for chains subordinate to the Poisson process.

This page titled [16.17: Potential Matrices](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.18: Stationary and Limiting Distributions of Continuous-Time Chains

In this section, we study the limiting behavior of continuous-time Markov chains by focusing on two interrelated ideas: invariant (or stationary) distributions and limiting distributions. In some ways, the limiting behavior of continuous-time chains is simpler than the limiting behavior of discrete-time chains, in part because the complications caused by periodicity in the discrete-time case do not occur in the continuous-time case. Nonetheless as we will see, the limiting behavior of a continuous-time chain is closely related to the limiting behavior of the embedded, discrete-time jump chain.

Review

Once again, our starting point is a time-homogeneous, continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and with discrete state space (S, \mathcal{S}) . By definition, this means that S is countable with the discrete topology, so that \mathcal{S} is the σ -algebra of all subsets of S .

Let's review what we have so far. We assume that the Markov chain \mathbf{X} is *regular*. Among other things, this means that the basic structure of \mathbf{X} is determined by the *transition times* $\tau = (\tau_0, \tau_1, \tau_2, \dots)$ and the *jump chain* $\mathbf{Y} = (Y_0, Y_1, Y_2, \dots)$. First, $\tau_0 = 0$ and $\tau_1 = \tau = \inf\{t > 0 : X_t \neq X_0\}$. The time-homogeneous and Markov properties imply that the distribution of τ given $X_0 = x$ is exponential with parameter $\lambda(x) \in [0, \infty)$. Part of regularity is that \mathbf{X} is right continuous so that there are no *instantaneous states* where $\lambda(x) = \infty$, which would mean $\mathbb{P}(\tau = 0 | X_0 = x) = 1$. On the other hand, $\lambda(x) \in (0, \infty)$ means that x is a *stable state* so that τ has a proper exponential distribution given $X_0 = x$, with $\mathbb{P}(0 < \tau < \infty | X_0 = x) = 1$. Finally, $\lambda(x) = 0$ means that x is an *absorbing state* so that $\mathbb{P}(\tau = \infty | X_0 = x) = 1$. The remaining transition times are defined recursively: $\tau_{n+1} = \inf\{t > \tau_n : X_t \neq X_{\tau_n}\}$ if $\tau_n < \infty$ and $\tau_{n+1} = \infty$ if $\tau_n = \infty$. Another component of regularity is that with probability 1, $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$, ruling out the *explosion event* of infinitely many jumps in finite time. The jump chain \mathbf{Y} is formed by sampling \mathbf{X} at the transition times (until the chain is sucked into an absorbing state, if that happens). That is, with $M = \sup\{n : \tau_n < \infty\}$ and for $n \in \mathbb{N}$, we define $Y_n = X_{\tau_n}$ if $n \leq M$ and $Y_n = X_{\tau_M}$ if $n > M$. Then \mathbf{Y} is a discrete-time Markov chain with one-step transition matrix Q given $Q(x, y) = \mathbb{P}(X_{\tau} = y | X_0 = x)$ if $(x, y) \in S^2$ with x stable and $Q(x, x) = 1$ if $x \in S$ is absorbing.

The transition matrix P_t at time $t \in [0, \infty)$ is given by $P_t(x, y) = \mathbb{P}(X_t = y | X_0 = x)$ for $(x, y) \in S^2$. The time-homogeneous and Markov properties imply that the collection of transition matrices $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ satisfies the *Chapman-Kolmogorov equations* $P_s P_t = P_{s+t}$ for $s, t \in [0, \infty)$, and hence is a *semigroup* of transition matrices. The transition semigroup \mathbf{P} and the initial distribution of X_0 determine all of the finite-dimensional distributions of \mathbf{X} . Since there are no instantaneous states, P is *standard* which means that $P_t \rightarrow I$ as $t \downarrow 0$ (as matrices, and so pointwise). The fundamental relationship between \mathbf{P} on the one hand, and λ and Q on the other, is

$$P_t(x, y) = I(x, y)e^{-\lambda(x)t} + \int_0^t \lambda(x)e^{-\lambda(x)s} Q P_{t-s}(x, y) ds, \quad (x, y) \in S^2 \quad (16.18.1)$$

From this, it follows that the matrix function $t \mapsto P_t$ is differentiable (again, pointwise) and satisfies the *Kolmogorov backward equation* $\frac{d}{dt} P_t = G P_t$, where the *infinitesimal generator matrix* G is given by $G(x, y) = -\lambda(x)I(x, y) + \lambda(x)Q(x, y)$ for $(x, y) \in S^2$. If we impose the stronger assumption that \mathbf{P} is *uniform*, which means that $P_t \rightarrow I$ as $t \downarrow 0$ as operators on \mathcal{B} (so with respect to the supremum norm), then the backward equation as well as the companion *Kolmogorov forward equation* $\frac{d}{dt} P_t = P_t G$ hold as operators on \mathcal{B} . In addition, we have the matrix exponential representation $P_t = e^{tG}$ for $t \in [0, \infty)$. The uniform assumption is equivalent to the exponential parameter function being bounded.

Finally, for $\alpha \in [0, \infty)$, the α potential matrix U_α of \mathbf{X} is $U_\alpha = \int_0^\infty e^{-\alpha t} P_t dt$. The *resolvent* $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$ is the Laplace transform of \mathbf{P} and hence gives the same information as \mathbf{P} . From this point of view, the time-homogeneous and Markov properties lead to the *resolvent equation* $U_\alpha = U_\beta + (\beta - \alpha)U_\alpha U_\beta$ for $\alpha, \beta \in (0, \infty)$ with $\alpha \leq \beta$. For $\alpha \in (0, \infty)$, the α potential matrix is related to the generator by the fundamental equation $\alpha U_\alpha = I + G U_\alpha$. If \mathbf{P} is uniform, then this equation, as well as the companion $\alpha U_\alpha = I + U_\alpha G$ hold as operators on \mathcal{B} , which leads to $U_\alpha = (\alpha I - G)^{-1}$.

Basic Theory

Relations and Classification

We start our discussion with relations among states and classifications of states. These are the same ones that we studied for discrete-time chains in our study of recurrence and transience, applied here to the jump chain \mathbf{Y} . But as we will see, the relations and classifications make sense for the continuous-time chain \mathbf{X} as well. The discussion is complicated slightly when there are absorbing states. Only when \mathbf{X} is in an absorbing state can we not interpret the values of \mathbf{Y} as the values of \mathbf{X} at the transition times (because of course, there are no transitions when \mathbf{X} is in an absorbing state). But $x \in S$ is absorbing for the continuous-time chain \mathbf{X} if and only if x is absorbing for the jump chain \mathbf{Y} , so this trivial exception is easily handled.

For $y \in S$ let $\rho_y = \inf\{n \in \mathbb{N}_+ : Y_n = y\}$, the (discrete) hitting time to y for the jump chain \mathbf{Y} , where as usual, $\inf(\emptyset) = \infty$. That is, ρ_y is the first positive (discrete) time that \mathbf{Y} is in state y . The analogous random time for the continuous-time chain \mathbf{X} is τ_{ρ_y} , where naturally we take $\tau_\infty = \infty$. This is the first time that \mathbf{X} is in state y , not counting the possible initial period in y . Specifically, suppose $X_0 = x$. If $x \neq y$ then $\tau_{\rho_y} = \inf\{t > 0 : X_t = y\}$. If $x = y$ then $\tau_{\rho_y} = \inf\{t > \tau_1 : X_t = y\}$.

Define the *hitting matrix* H by

$$H(x, y) = \mathbb{P}(\rho_y < \infty \mid Y_0 = x), \quad (x, y) \in S^2 \quad (16.18.2)$$

Then $H(x, y) = \mathbb{P}(\tau_{\rho_y} < \infty \mid X_0 = x)$ except when x is absorbing and $y = x$.

So for the continuous-time chain, if $x \in S$ is stable then $H(x, x)$ is the probability that, starting in x , the chain \mathbf{X} returns to x after its initial period in x . If $x, y \in S$ are distinct, then $H(x, y)$ is simply the probability that \mathbf{X} , starting in x , eventually reaches y . It follows that the basic relation among states makes sense for either the continuous-time chain \mathbf{X} as well as its jump chain \mathbf{Y} .

Define the relation \rightarrow on S^2 by $x \rightarrow y$ if $x = y$ or $H(x, y) > 0$.

The *leads to* relation \rightarrow is *reflexive* by definition: $x \rightarrow x$ for every $x \in S$. From our previous study of discrete-time chains, we know it's also *transitive*: if $x \rightarrow y$ and $y \rightarrow z$ then $x \rightarrow z$ for $x, y, z \in S$. We also know that $x \rightarrow y$ if and only if there is a directed path in the state graph from x to y , if and only if $Q^n(x, y) > 0$ for some $n \in \mathbb{N}$. For the continuous-time transition matrices, we have a stronger result that in turn makes a stronger case that the *leads to* relation is fundamental for \mathbf{X} .

Suppose $(x, y) \in S^2$.

1. If $x \rightarrow y$ then $P_t(x, y) > 0$ for all $t \in (0, \infty)$.
2. If $x \not\rightarrow y$ then $P_t(x, y) = 0$ for all $t \in (0, \infty)$.

Proof

This result is proved in the section on transition matrices and generators.

This result is known as the *Lévy dichotomy*, and is named for Paul Lévy. Let's recall a couple of other definitions:

Suppose that A is a nonempty subset of S .

1. A is *closed* if $x \in A$ and $x \rightarrow y$ imply $y \in A$.
2. A is *irreducible* if A is closed and has no proper closed subset.

If S is irreducible, we also say that the chain \mathbf{X} itself is irreducible.

If A is a nonempty subset of S , then $\text{cl}(A) = \{y \in S : x \rightarrow y \text{ for some } x \in A\}$ is the smallest closed set containing A , and is called the *closure* of A .

Suppose that $A \subseteq S$ is closed. Then

1. P_t^A , the restriction of P_t to $A \times A$, is a transition probability matrix on A for every $t \in [0, \infty)$.
2. \mathbf{X} restricted to A is a continuous-time Markov chain with transition semigroup $\mathbf{P}^A = \{P_t^A : t \in [0, \infty)\}$.

Proof

1. If $x \in A$ and $y \notin A$, then x does not lead to y so in particular $P_t(x, y) = 0$. It follows that $\sum_{y \in A} P_t(x, y) = 1$ for $x \in A$ so P_t^A is a transition probability matrix.
2. This follows from (a). If the chain starts in A , then the chain remains in A for all time, and of course, the Markov property still holds.

Define the relation \leftrightarrow on S by $x \leftrightarrow y$ if $x \rightarrow y$ and $y \rightarrow x$ for $(x, y) \in S^2$.

The *to and from* relation \leftrightarrow defines an equivalence relation on S and hence partitions S into mutually disjoint equivalence classes. Recall from our study of discrete-time chains that a closed set is not necessarily an equivalence class, nor is an equivalence class necessarily closed. However, an irreducible set is an equivalence class, but an equivalence class may not be irreducible. The importance of the relation \leftrightarrow stems from the fact that many important properties of Markov chains (in discrete or continuous time) turn out to be *class properties*, shared by all states in an equivalence class. The following definition is fundamental, and once again, makes sense for either the continuous-time chain \mathbf{X} or its jump chain \mathbf{Y} .

Let $x \in S$.

1. State x is *transient* if $H(x, x) < 1$
2. State x is *recurrent* if $H(x, x) = 1$.

Recall from our study of discrete-time chains that if x is recurrent and $x \rightarrow y$ then y is recurrent and $y \rightarrow x$. Thus, recurrence and transience are class properties, shared by all states in an equivalence class.

Time Spent in a State

For $x \in S$, let N_x denote the number of visits to state x by the jump chain \mathbf{Y} , and let T_x denote the total time spent in state x by the continuous-time chain \mathbf{X} . Thus

$$N_x = \sum_{n=0}^{\infty} \mathbf{1}(Y_n = x), \quad T_x = \int_0^{\infty} \mathbf{1}(X_t = x) dt \quad (16.18.3)$$

The expected values $R(x, y) = \mathbb{E}(N_y | Y_0 = x)$ and $U(x, y) = \mathbb{E}(T_y | X_0 = x)$ for $(x, y) \in S^2$ define the *potential matrices* of \mathbf{Y} and \mathbf{X} , respectively. From our previous study of discrete-time chains, we know the distribution and mean of N_y given $Y_0 = x$ in terms of the hitting matrix H . The next two results give a review:

Suppose that $x, y \in S$ are distinct. Then

1. $\mathbb{P}(N_y = n | Y_0 = y) = H^{n-1}(y, y)[1 - H(y, y)]$ for $n \in \mathbb{N}_+$
2. $\mathbb{P}(N_y = 0 | Y_0 = x) = 1 - H(x, y)$ and $\mathbb{P}(N_y = n | Y_0 = x) = H(x, y)H^{n-1}(y, y)[1 - H(y, y)]$ for $n \in \mathbb{N}_+$

Let's take cases. First suppose that y is recurrent. In part (a), $\mathbb{P}(N_y = n | Y_0 = y) = 0$ for all $n \in \mathbb{N}_+$, and consequently $\mathbb{P}(N_y = \infty | Y_0 = y) = 1$. In part (b), $\mathbb{P}(N_y = n | Y_0 = x) = 0$ for $n \in \mathbb{N}_+$, and consequently $\mathbb{P}(N_y = 0 | Y_0 = x) = 1 - H(x, y)$ while $\mathbb{P}(N_y = \infty | Y_0 = x) = H(x, y)$. Suppose next that y is transient. Part (a) specifies a proper geometric distribution on \mathbb{N}_+ while in part (b), probability $1 - H(x, y)$ is assigned to 0 and the remaining probability $H(x, y)$ is geometrically distributed over \mathbb{N}_+ as in (a). In both cases, N_y is finite with probability 1. Next we consider the expected value, that is, the (discrete) potential. To state the results succinctly we will use the convention that $a/0 = \infty$ if $a > 0$ and $0/0 = 0$.

Suppose again that $x, y \in S$ are distinct. Then

1. $R(y, y) = 1/[1 - H(y, y)]$
2. $R(x, y) = H(x, y)/[1 - H(y, y)]$

Let's take cases again. If $y \in S$ is recurrent then $R(y, y) = \infty$, and for $x \in S$ with $x \neq y$, either $R(x, y) = \infty$ if $x \rightarrow y$ or $R(x, y) = 0$ if $x \nrightarrow y$. If $y \in S$ is transient, $R(y, y)$ is finite, as is $R(x, y)$ for every $x \in S$ with $x \neq y$. Moreover, there is an inverse relationship of sorts between the potential and the hitting probabilities.

Naturally, our next goal is to find analogous results for the continuous-time chain \mathbf{X} . For the distribution of T_y it's best to use the right distribution function.

Suppose that $x, y \in S$ are distinct. Then for $t \in [0, \infty)$

1. $\mathbb{P}(T_y > t \mid X_0 = y) = \exp\{-\lambda(y)[1 - H(y, y)]t\}$
2. $\mathbb{P}(T_y > t \mid X_0 = x) = H(x, y) \exp\{-\lambda(y)[1 - H(y, y)]t\}$

Proof

The proof is by conditioning on N_y .

1. First, if $H(y, y) = 1$ (so that y is recurrent), then either y is absorbing with $\mathbb{P}(\tau_1 = \infty \mid X_0 = y) = 1$ or y is stable and recurrent, so that $\mathbb{P}(N_y = \infty \mid X_0 = y) = 1$. In the second case, starting in state y , T_y is the sum of infinitely many independent variables, each with the exponential distribution with parameter $\lambda(y) \in (0, \infty)$. In both cases, $\mathbb{P}(T_y = \infty \mid X_0 = y) = 1$ and so $\mathbb{P}(T_y > t \mid X_0 = y) = 1$ for every $t \in [0, \infty)$. So suppose that $H(y, y) < 1$ so that y is transient. Then

$$\mathbb{P}(T_y > t \mid X_0 = y) = \sum_{n=1}^{\infty} \mathbb{P}(T_y > t \mid X_0 = y, N_y = n) \mathbb{P}(N_y = n \mid X_0 = y) \quad (16.18.4)$$

Given $N_y = n$, T_y is the sum of n independent variables, each having the exponential distribution with parameter $\lambda(y)$. So T_y has the gamma distribution with parameters n and $\lambda(y)$ and hence

$$\mathbb{P}(T_y > t \mid X_0 = y, N_y = n) = \sum_{k=0}^{n-1} e^{-\lambda(y)t} \frac{[\lambda(y)t]^k}{k!} \quad (16.18.5)$$

From the previous result, $\mathbb{P}(N_y = n \mid X_0 = y) = \mathbb{P}(N_y = n \mid Y_0 = y) = H^{n-1}(y, y)[1 - H(y, y)]$. We substitute, change the order of summation, use geometric series and then exponential series:

$$\begin{aligned} \mathbb{P}(T_y > t \mid X_0 = y) &= \sum_{n=1}^{\infty} \left(\sum_{k=0}^{n-1} e^{-\lambda(y)t} \frac{[\lambda(y)t]^k}{k!} \right) H^{n-1}(y, y)[1 - H(y, y)] \\ &= e^{-\lambda(y)t} [1 - H(y, y)] \sum_{k=0}^{\infty} \frac{[\lambda(y)t]^k}{k!} \sum_{n=k+1}^{\infty} H^{n-1}(y, y) \\ &= e^{-\lambda(y)t} \sum_{k=0}^{\infty} \frac{[\lambda(y)t]^k}{k!} H^k(y, y) = e^{-\lambda(y)t} \exp[\lambda(y)H(y, y)t] \end{aligned}$$

Simplifying gives the result.

2. The proof is similar. If $H(y, y) = 1$ so that y is recurrent, then starting in state x , either $T_y = 0$ if $N_y = 0$, which occurs with probability $1 - H(x, y)$ or $T_y = \infty$ if $N_y = \infty$, which occurs with probability $H(x, y)$. If $H(y, y) < 1$ so that y is transient, then the result follows from conditioning on N_y as in (a), except that $\mathbb{P}(T_y = 0 \mid X_0 = x) = \mathbb{P}(N_y = 0 \mid Y_0 = x) = 1 - H(x, y)$.

Let's take cases as before. Suppose first that y is recurrent. In part (a), $\mathbb{P}(T_y > t \mid X_0 = y) = 1$ for every $t \in [0, \infty)$ and hence $\mathbb{P}(T_y = \infty \mid X_0 = y) = 1$. In part (b), $\mathbb{P}(T_y > t \mid X_0 = x) = H(x, y)$ for every $t \in [0, \infty)$ and consequently $\mathbb{P}(T_y = 0 \mid X_0 = x) = 1 - H(x, y)$ while $\mathbb{P}(T_y = \infty \mid X_0 = x) = H(x, y)$. Suppose next that y is transient. From part (a), the distribution of T_y given $X_0 = y$ is exponential with parameter $\lambda(y)[1 - H(y, y)]$. In part (b), the distribution assigns probability $1 - H(x, y)$ to 0 while the remaining probability $H(x, y)$ is exponentially distributed over $(0, \infty)$ as in (a). Taking expected value, we get a very nice relationship between the potential matrix U of the continuous-time chain \mathbf{X} and the potential matrix R of the discrete-time jump chain \mathbf{Y} :

For every $(x, y) \in S^2$,

$$U(x, y) = \frac{R(x, y)}{\lambda(y)} \quad (16.18.6)$$

Proof

If y is recurrent, then $U(x, y) = R(x, y)$ and the common value is either 0 if $H(x, y) = 0$ or ∞ if $H(x, y) = 1$. So suppose that y is transient. We can compute the expected value of T_y by integrating the right distribution function in the previous theorem. In case $x = y$, we have

$$U(y, y) = \int_0^\infty \exp\{-\lambda(y)[1 - H(y, y)]t\} dt = \frac{1}{\lambda(y)[1 - H(y, y)]} = \frac{R(y, y)}{\lambda(y)} \quad (16.18.7)$$

In the case that x and y are distinct,

$$U(x, y) = \int_0^\infty H(x, y) \exp\{-\lambda(y)[1 - H(y, y)]t\} dt = \frac{H(x, y)}{\lambda(y)[1 - H(y, y)]} = \frac{R(x, y)}{\lambda(y)} \quad (16.18.8)$$

In particular, $y \in S$ is transient if and only if $R(x, y) < \infty$ for every $x \in S$, if and only if $U(x, y) < \infty$ for every $x \in S$. On the other hand, y is recurrent if and only if $R(x, y) = U(x, y) = \infty$ if $x \rightarrow y$ and $R(x, y) = U(x, y) = 0$ if $x \nrightarrow y$.

Null and Positive Recurrence

Unlike transience and recurrence, the definitions of null and positive recurrence of a state $x \in S$ are different for the continuous-time chain \mathbf{X} and its jump chain \mathbf{Y} . This is because these definitions depend on the *expected* hitting time to x , starting in x , and not just the *finiteness* of this hitting time. For $x \in S$, let $\nu(x) = \mathbb{E}(\rho_x \mid Y_0 = x)$, the expected (discrete) return time to x starting in x . Recall that x is *positive recurrent* for \mathbf{Y} if $\nu(x) < \infty$ and x is *null recurrent* if x is recurrent but not positive recurrent, so that $H(x, x) = 1$ but $\nu(x) = \infty$. The definitions are similar for \mathbf{X} , but using the continuous hitting time τ_{ρ_x} .

For $x \in S$, let $\mu(x) = 0$ if x is absorbing and $\mu(x) = \mathbb{E}(\tau_{\rho_x} \mid X_0 = x)$ if x is stable. So if x is stable, $\mu(x)$ is the expected return time to x starting in x (after the initial period in x).

1. State x is *positive recurrent* for \mathbf{X} if $\mu(x) < \infty$.
2. State x is *null recurrent* for \mathbf{X} if x recurrent but not positive recurrent, so that $H(x, x) = 1$ but $\mu(x) = \infty$.

A state $x \in S$ can be positive recurrent for \mathbf{X} but null recurrent for its jump chain \mathbf{Y} or can be null recurrent for \mathbf{X} but positive recurrent for \mathbf{Y} . But like transience and recurrence, positive and null recurrence are *class properties*, shared by all states in an equivalence class under the *to and from* equivalence relation \leftrightarrow .

Invariant Functions

Our next discussion concerns functions that are invariant for the transition matrix Q of the jump chain \mathbf{Y} and functions that are invariant for the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ of the continuous-time chain \mathbf{X} . For both discrete-time and continuous-time chains, there is a close relationship between invariant functions and the limiting behavior in time.

First let's recall the definitions. A function $f : S \rightarrow [0, \infty)$ is *invariant* for Q (or for the chain \mathbf{Y}) if $fQ = f$. It then follows that $fQ^n = f$ for every $n \in \mathbb{N}$. In continuous time we must *assume* invariance at each time. That is, a function $f : S \rightarrow [0, \infty)$ is invariant for \mathbf{P} (or for the chain \mathbf{X}) if $fP_t = f$ for all $t \in [0, \infty)$. Our interest is in nonnegative functions, because we can think of such a function as the density function, with respect to counting measure, of a positive measure on S . We are particularly interested in the special case that f is a *probability density function*, so that $\sum_{x \in S} f(x) = 1$. If Y_0 has a probability density function f that is invariant for Q , then Y_n has probability density function f for all $n \in \mathbb{N}$ and hence \mathbf{Y} is *stationary*. Similarly, if X_0 has a probability density function f that is invariant for \mathbf{P} then X_t has probability density function f for every $t \in [0, \infty)$ and once again, the chain \mathbf{X} is *stationary*.

Our first result shows that there is a one-to-one correspondence between invariant functions for Q and zero functions for the generator G .

Suppose $f : S \rightarrow [0, \infty)$. Then $fG = 0$ if and only if $(\lambda f)Q = \lambda f$, so that λf is invariant for Q .

Proof

This is a simple consequence of the definition of the generator:

$$fG(y) = \sum_{x \in S} f(x)G(x, y) = -\lambda(y)f(y) + \sum_{x \in S} f(x)\lambda(x)Q(x, y), \quad y \in S \quad (16.18.9)$$

or in functional form, $fG = -\lambda f + (\lambda f)Q$

If our chain \mathbf{X} has no absorbing states, then $f : S \rightarrow [0, \infty)$ is invariant for Q if and only if $(f/\lambda)G = 0$.

Suppose that $f : S \rightarrow [0, \infty)$. Then f is invariant for \mathbf{P} if and only if $fG = 0$.

Proof 1

Assume that λ is bounded, so that the transition semigroup \mathbf{P} is uniform. Then $P_t = e^{tG}$ for $t \in [0, \infty)$. So if $f : S \rightarrow [0, \infty)$ then

$$fP_t = f(e^{tG}) = f \sum_{n=0}^{\infty} \frac{t^n}{n!} G^n = f + \sum_{n=1}^{\infty} \frac{t^n}{n!} fG^n \quad (16.18.10)$$

Since f is nonnegative, $fP_t = f$ if and only if $fG = 0$ (in which case $fG^n = 0$ for every $n \in \mathbb{N}_+$).

Proof 2

Suppose that $fP_t = f$ for $t \in [0, \infty)$. Then $\frac{d}{dt}(fP_t) = 0$ for $t \in [0, \infty)$. But using the Kolmogorov backward equation, $\frac{d}{dt}(fP_t) = f \frac{d}{dt}P_t = fGP_t = 0$. Letting $t = 0$ we conclude that $fG = 0$. Conversely, if $fG = 0$ then $\frac{d}{dt}(fP_t) = f \frac{d}{dt}P_t = fGP_t = 0$ for $t \in [0, \infty)$. It follows that fP_t is constant in $t \in [0, \infty)$. Since $fP_0 = f$ it follows that $fP_t = f$ for all $t \in [0, \infty)$.

So putting the two main results together we see that f is invariant for the continuous-time chain \mathbf{X} if and only if λf is invariant for the jump chain \mathbf{Y} . Our next result shows how functions that are invariant for \mathbf{X} are related to the resolvent $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$. To appreciate the result, recall that for $\alpha \in (0, \infty)$ the matrix αU_α is a *probability* matrix, and in fact $\alpha U_\alpha(x, \cdot)$ is the conditional probability density function of X_T , given $X_0 = x$, where T is independent of \mathbf{X} and has the exponential distribution with parameter α . So αU_α is a transition matrix just as P_t is a transition matrix, but corresponding to the exponentially distributed random time T with parameter $\alpha \in (0, \infty)$ rather than the deterministic time $t \in [0, \infty)$.

Suppose that $f : S \rightarrow [0, \infty)$. If $fG = 0$ then $f(\alpha U_\alpha) = f$ for $\alpha \in (0, \infty)$. Conversely if $f(\alpha U_\alpha) = f$ for $\alpha \in (0, \infty)$ then $fG = 0$.

Proof

Recall that $I + GU_\alpha = \alpha U_\alpha$ for $\alpha \in (0, \infty)$. Hence if $fG = 0$ then

$$f(\alpha U_\alpha) = f + fGU_\alpha = f \quad (16.18.11)$$

Conversely, suppose that $f(\alpha U_\alpha) = f$. Then

$$fGU_\alpha = \int_0^\infty e^{-\alpha t} fGP_t dt = 0 \quad (16.18.12)$$

As a function of $\alpha \in (0, \infty)$, the integral on the right side is the Laplace transform of the time function $t \mapsto fGP_t$. Hence we must have $fGP_t = 0$ for $t \in (0, \infty)$, and letting $t \downarrow 0$ gives $fG = 0$.

So extending our summary, $f : S \rightarrow [0, \infty)$ is invariant for the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ if and only if λf is invariant for jump transition matrices $\{Q^n : n \in \mathbb{N}\}$ if and only if $fG = 0$ if and only if f is invariant for the collection of probability matrices $\{\alpha U_\alpha : \alpha \in (0, \infty)\}$. From our knowledge of the theory for discrete-time chains, we now have the following fundamental result:

Suppose that \mathbf{X} is irreducible and recurrent.

1. There exists $g : S \rightarrow (0, \infty)$ that is invariant for \mathbf{X} .
2. If f is invariant for \mathbf{X} , then $f = cg$ for some constant $c \in [0, \infty)$.

Proof

The result is trivial if S consists of a single, necessarily absorbing, state. Otherwise, there are no absorbing states, since \mathbf{X} is irreducible and so $\lambda(x) > 0$ for $x \in S$. From the result above, f is invariant for \mathbf{X} if and only if λf is invariant for \mathbf{Y} . But \mathbf{Y} is also irreducible and recurrent, so we know that there exists a strictly positive function that is invariant for \mathbf{Y} , and every other function that is invariant for \mathbf{Y} is a nonnegative multiple of this one. Hence the same is true for \mathbf{X} .

Invariant functions have a nice interpretation in terms of *occupation times*, an interpretation that parallels the discrete case. The potential gives the expected total time in a state, starting in another state, but here we need to consider the expected time in a state during a *cycle* that starts and ends in another state.

For $x \in S$, define the function γ_x by

$$\gamma_x(y) = \mathbb{E} \left(\int_0^{\tau_{P_x}} \mathbf{1}(X_s = y) ds \mid X_0 = x \right), \quad y \in S \quad (16.18.13)$$

so that $\gamma_x(y)$ is the expected *occupation time* in state y before the first return to x , starting in x .

Suppose again that \mathbf{X} is irreducible and recurrent. For $x \in S$,

1. $\gamma_x : S \rightarrow (0, \infty)$
2. γ_x is invariant for \mathbf{X}
3. $\gamma_x(x) = 1/\lambda(x)$
4. $\mu(x) = \sum_{y \in S} \gamma_x(y)$

Proof

As is often the case, the proof is based on results that we already have for the embedded jump chain. For $x \in S$, define

$$\delta_x(y) = \mathbb{E} \left(\sum_{n=0}^{\rho_x-1} \mathbf{1}(Y_n = y) \mid Y_0 = x \right), \quad y \in S \quad (16.18.14)$$

so that $\delta_x(y)$ is the expected number of visits to y before the first return to x , starting in x , for the jump chain $\mathbf{Y} = (Y_0, Y_1, \dots)$. Since \mathbf{X} is irreducible and recurrent, so is \mathbf{Y} . From our results in the discrete case we know that

1. $\delta_x : S \rightarrow (0, \infty)$
2. δ_x is invariant for \mathbf{Y}
3. $\delta_x(x) = 1$

From our results above, it follows that the function $y \mapsto \delta_x(y)/\lambda(y)$ satisfies properties (a), (b), and (c) in the theorem. But each visit to y by the jump chain \mathbf{Y} has expected length $1/\lambda(y)$ for the continuous-time chain \mathbf{X} . It follows that $\gamma_x(y) = \delta_x(y)/\lambda(y)$ for $x, y \in S$. By definition, $\gamma_x(y)$ is the expected occupation time in y before the first return to x , starting in x . Hence, summing over $y \in S$ gives the expected return time to x , starting in x , so (d) holds.

So now we have some additional insight into positive and null recurrence for the continuous-time chain \mathbf{X} and the associated jump chain \mathbf{Y} . Suppose again that the chains are irreducible and recurrent. There exist $g : S \rightarrow (0, \infty)$ that is invariant for \mathbf{Y} , and then g/λ is invariant for \mathbf{X} . The invariant functions are unique up to multiplication by positive constants. The jump chain \mathbf{Y} is positive recurrent if and only if $\sum_{x \in S} g(x) < \infty$ while the continuous-time chain \mathbf{X} is positive recurrent if and only if $\sum_{x \in S} g(x)/\lambda(x) < \infty$. Note that if λ is bounded (which is equivalent to the transition semigroup \mathbf{P} being uniform), then \mathbf{X} is positive recurrent if and only if \mathbf{Y} is positive recurrent.

Suppose again that \mathbf{X} is irreducible and recurrent.

1. If \mathbf{X} is null recurrent then \mathbf{X} does not have an invariant probability density function.
2. If \mathbf{X} is positive recurrent then \mathbf{X} has a unique, positive invariant probability density function.

Proof

From the previous result, there exists $g : S \rightarrow (0, \infty)$ that is invariant for \mathbf{X} , and every other invariant function is a nonnegative multiple of this one. The function f given by

$$f(y) = \frac{g(y)}{\sum_{x \in S} g(x)}, \quad y \in S \quad (16.18.15)$$

is uniquely defined (that is, unchanged if we replace g by cg where $c > 0$).

1. If $\sum_{x \in S} g(x) = \infty$ then $f(y) = 0$ for every $y \in S$.
2. If $\sum_{x \in S} g(x) < \infty$ then $f(y) > 0$ for every $y \in S$ and $\sum_{y \in S} f(y) = 1$.

Limiting Behavior

Our next discussion focuses on the limiting behavior of the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. Our first result is a simple corollary of the result above for potentials.

If $y \in S$ is transient, then $P_t(x, y) \rightarrow 0$ as $t \rightarrow \infty$ for every $x \in S$.

Proof

This follows from the previous result. If $y \in S$ is transient, then for any $x \in S$,

$$U(x, y) = \int_0^\infty P_t(x, y) dt < \infty \quad (16.18.16)$$

and so we must have $P_t(x, y) \rightarrow 0$ as $t \rightarrow \infty$.

So we should turn our attention to the recurrent states. The set of recurrent states partitions into equivalent classes under \leftrightarrow , and each of these classes is irreducible. Hence we can assume without loss of generality that our continuous-time chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is irreducible and recurrent. To avoid trivialities, we will also assume that S has at least two states. Thus, there are no absorbing states and so $\lambda(x) > 0$ for $x \in S$. Here is the main result.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is irreducible and recurrent. Then $f(y) = \lim_{t \rightarrow \infty} P_t(x, y)$ exists for each $y \in S$, independently of $x \in S$. The function f is invariant for \mathbf{X} and

$$f(y) = \frac{\gamma_x(y)}{\mu(x)}, \quad y \in S \quad (16.18.17)$$

1. If \mathbf{X} is null recurrent then $f(y) = 0$ for all $y \in S$.
2. If \mathbf{X} is positive recurrent then $f(y) > 0$ for all $y \in S$ and $\sum_{y \in S} f(y) = 1$.

Proof sketch

The basic idea is that

$$\lim_{t \rightarrow \infty} P_t(x, y) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t P_s(x, y) ds \quad (16.18.18)$$

The expression on the right is the limiting proportion of time spent in $y \in S$, starting in $x \in S$. This proportion is $\gamma_x(y)/\mu(x)$, so the results follow from the theorem [above](#).

The limiting function f can be computed in a number of ways. First we find a function $g : S \rightarrow (0, \infty)$ that is invariant for \mathbf{X} . We can do this by solving

- $gP_t = g$ for $t \in (0, \infty)$
- $gG = 0$
- $g(\alpha U_\alpha) = g$ for $\alpha \in (0, \infty)$
- $hQ = h$ and then $g = h/\lambda$

The function g is unique up to multiplication by positive constants. If $\sum_{x \in S} g(x) < \infty$, then we are in the positive recurrent case and so f is simply g normalized:

$$f(y) = \frac{g(y)}{\sum_{x \in S} g(x)}, \quad y \in S \quad (16.18.19)$$

The following result is known as the *ergodic theorem* for continuous-time Markov chains. It can also be thought of as a strong law of large numbers for continuous-time Markov chains.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is irreducible and positive recurrent, with (unique) invariant probability density function f . If $h : S \rightarrow \mathbb{R}$ then

$$\frac{1}{t} \int_0^t h(X_s) ds \rightarrow \sum_{x \in S} f(x)h(x) \text{ as } t \rightarrow \infty \quad (16.18.20)$$

with probability 1, assuming that the sum on the right converges absolutely.

Notes

First, let $x, y \in S$ and let $h = \mathbf{1}_y$, the indicator function of y . Then given $X_0 = x$, $\frac{1}{t} \int_0^t h(X_s) ds$ is the average occupation time in state y , starting in state x , over the time interval $[0, t]$. In expected value, this is $\frac{1}{t} \int_0^t P_s(x, y) ds$ which we know converges to $f(y)$ as $t \rightarrow \infty$, independently of x . So in this special case, the ergodic theorem states that the convergence is with probability 1 also. A general function $h : S \rightarrow \mathbb{R}$ is a linear combination of the indicator functions of the points in S , so the ergodic theorem is plausible.

Note that no assumptions are made about X_0 , so the limit is independent of the initial state. By now, this should come as no surprise. After a long period of time, the Markov chain \mathbf{X} “forgets” about the initial state. Note also that $\sum_{x \in S} f(x)h(x)$ is the expected value of h , thought of as a random variable on S with probability measure defined by f . On the other hand, $\frac{1}{t} \int_0^t h(X_s) ds$ is the average of the time function $s \mapsto h(X_s)$ on the interval $[0, t]$. So the ergodic theorem states that the limiting time average on the left is the same as the spatial average on the right.

Applications and Exercises

The Two-State Chain

The continuous-time, two-state chain has been studied in the last several sections. The following result puts the pieces together and completes the picture.

Consider the continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1\}$ with transition rate $a \in (0, \infty)$ from 0 to 1 and transition rate $b \in (0, \infty)$ from 1 to 0. Give each of the following

1. The transition matrix Q^n for \mathbf{Y} at $n \in \mathbb{N}$.
2. The infinitesimal generator G .
3. The transition matrix P_t for \mathbf{X} at $t \in [0, \infty)$.
4. The invariant probability density function for \mathbf{Y} .
5. The invariant probability density function for \mathbf{X} .
6. The limiting behavior of Q^n as $n \rightarrow \infty$.
7. The limiting behavior of P_t as $t \rightarrow \infty$.

Answer

Note that since the transition rates a and b are positive, the chain is irreducible.

1. First, $Q = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ and then for $n \in \mathbb{N}$, $Q^n = Q$ if n is odd and $Q^n = I$ if n is even.
2. $G = \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}$.
3. $P_t = \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix} - \frac{1}{a+b} e^{-(a+b)t} \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}$ for $t \in [0, \infty)$.
4. $f_d = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \end{bmatrix}$
5. $f_c = \begin{bmatrix} \frac{b}{a+b} & \frac{a}{a+b} \end{bmatrix}$
6. As in (a), $Q^{2n} = I$ and $Q^{2n+1} = Q$ for $n \in \mathbb{N}$. So there are two sub-sequential limits. The jump chain \mathbf{Y} is periodic with period 2.
7. $P_t \rightarrow \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix}$ as $t \rightarrow \infty$. Each row is f_c .

Computational Exercises

The following continuous-time chain has also been studied in the previous three sections.

Consider the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1, 2\}$ with exponential parameter function $\lambda = (4, 1, 3)$ and jump transition matrix

$$Q = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ \frac{1}{3} & \frac{2}{3} & 0 \end{bmatrix} \quad (16.18.21)$$

1. Recall the generator matrix G .
2. Find the invariant probability density function f_d for \mathbf{Y} by solving $f_d Q = f_d$.
3. Find the invariant probability density function f_c for \mathbf{X} by solving $f_c G = 0$.
4. Verify that λf_c is a multiple of f_d .
5. Describe the limiting behavior of Q^n as $n \rightarrow \infty$.
6. Describe the limiting behavior of P_t as $t \rightarrow \infty$.
7. Verify the result in (f) by recalling the transition matrix P_t for \mathbf{X} at $t \in [0, \infty)$.

Answer

1. $G = \begin{bmatrix} -4 & 2 & 2 \\ 1 & -1 & 0 \\ 1 & 2 & -3 \end{bmatrix}$
2. $f_d = \frac{1}{14} [6 \ 5 \ 3]$
3. $f_c = \frac{1}{15} [3 \ 10 \ 2]$
4. $\lambda f_c = \frac{1}{15} [12 \ 10 \ 6] = \frac{28}{15} f_d$
5. $Q^n \rightarrow \frac{1}{14} \begin{bmatrix} 6 & 5 & 3 \\ 6 & 5 & 3 \\ 6 & 5 & 3 \end{bmatrix}$ as $n \rightarrow \infty$
6. $P_t \rightarrow \frac{1}{15} \begin{bmatrix} 3 & 10 & 2 \\ 3 & 10 & 2 \\ 3 & 10 & 2 \end{bmatrix}$ as $t \rightarrow \infty$
7. $P_t = \frac{1}{15} \begin{bmatrix} 3 + 12e^{-5t} & 10 - 10e^{-3t} & 2 - 12e^{-5t} + 10e^{-3t} \\ 3 - 3e^{-5t} & 10 + 5e^{-3t} & 2 + 3e^{-5t} - 5e^{-3t} \\ 3 - 3e^{-5t} & 10 - 10e^{-3t} & 2 + 3e^{-5t} + 10e^{-3t} \end{bmatrix}$ for $t \in [0, \infty)$

Special Models

Read the discussion of stationary and limiting distributions for chains subordinate to the Poisson process.

Read the discussion of stationary and limiting distributions for continuous-time birth-death chains.

Read the discussion of classification and limiting distributions for continuous-time queuing chains.

This page titled [16.18: Stationary and Limiting Distributions of Continuous-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.19: Time Reversal in Continuous-Time Chains

Earlier, we studied time reversal of discrete-time Markov chains. In continuous time, the issues are basically the same. First, the Markov property stated in the form that the past and future are independent given the present, essentially treats the past and future symmetrically. However, there is a lack of symmetry in the fact that in the usual formulation, we have an *initial* time 0, but not a *terminal* time. If we introduce a terminal time, then we can run the process backwards in time. In this section, we are interested in the following questions:

- Is the new process still Markov?
- If so, how are the various parameters of the reversed Markov chain related to those of the original chain?
- Under what conditions are the forward and backward Markov chains stochastically the same?

Consideration of these questions leads to *reversed chains*, an important and interesting part of the theory of continuous-time Markov chains. As always, we are also interested in the relationship between properties of a continuous-time chain and the corresponding properties of its discrete-time jump chain. In this section we will see that there are simple and elegant connections between the time reversal of a continuous-time chain and the time-reversal of the jump chain.

Basic Theory

Reversed Chains

Our starting point is a (homogeneous) continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with (countable) state space S . We will assume that \mathbf{X} is irreducible, so that every state in S leads to every other state, and to avoid trivialities, we will assume that there are at least two states. The irreducibility assumption involves no serious loss of generality since otherwise we could simply restrict our attention to an irreducible equivalence class of states. With our usual notation, we will let $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ denote the semigroup of transition matrices of \mathbf{X} and G the infinitesimal generator. Let $\lambda(x)$ denote the exponential parameter for the holding time in state $x \in S$ and Q the transition matrix for the discrete-time jump chain $\mathbf{Y} = (Y_0, Y_1, \dots)$. Finally, let $\mathbf{U} = \{U_\alpha : \alpha \in [0, \infty)\}$ denote the collection of potential matrices of \mathbf{X} . We will assume that the chain \mathbf{X} is regular, which gives us the following properties:

- $P_t(x, x) \rightarrow 1$ as $t \downarrow 0$ for $x \in S$.
- There are no instantaneous states, so $\lambda(x) < \infty$ for $x \in S$.
- The transition times (τ_1, τ_2, \dots) satisfy $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$.
- We may assume that the chain \mathbf{X} is right continuous and has left limits.

The assumption of regularity rules out various types of weird behavior that, while mathematically possible, are usually not appropriate in applications. If \mathbf{X} is uniform, a stronger assumption than regularity, we have the following additional properties:

- $P_t(x, x) \rightarrow 1$ as $t \downarrow 0$ uniformly in $x \in S$.
- λ is bounded.
- $P_t = e^{tG}$ for $t \in [0, \infty)$.
- $U_\alpha = (\alpha I - G)^{-1}$ for $\alpha \in (0, \infty)$.

Now let $h \in (0, \infty)$. We will think of h as the *terminal time* or *time horizon* so the chains in our first discussion will be defined on the time interval $[0, h]$. Notationally, we won't bother to indicate the dependence on h , since ultimately the time horizon won't matter. Define $\hat{X}_t = X_{h-t}$ for $t \in [0, h]$. Thus, the process *forward in time* is $\mathbf{X} = \{X_t : t \in [0, h]\}$ while the process *backwards in time* is

$$\hat{\mathbf{X}} = \{\hat{X}_t : t \in [0, h]\} = \{X_{h-t} : t \in [0, h]\} \quad (16.19.1)$$

Similarly let

$$\hat{\mathcal{F}}_t = \sigma\{\hat{X}_s : s \in [0, t]\} = \sigma\{X_{h-s} : s \in [0, t]\} = \sigma\{X_r : r \in [h-t, h]\}, \quad t \in [0, h] \quad (16.19.2)$$

So $\hat{\mathcal{F}}_t$ is the σ -algebra of events of the process $\hat{\mathbf{X}}$ up to time t , which of course, is also the σ -algebra of events of \mathbf{X} from time $h-t$ forward. Our first result is that the chain reversed in time is still Markov

The process $\hat{\mathbf{X}} = \{\hat{X}_t : t \in [0, h]\}$ is a Markov chain, but is not time homogeneous in general. For $s, t \in [0, h]$ with $s < t$, the transition matrix from s to t is

$$\hat{P}_{s,t}(x, y) = \frac{\mathbb{P}(X_{h-t} = y)}{\mathbb{P}(X_{h-s} = x)} P_{t-s}(y, x), \quad (x, y) \in S^2 \quad (16.19.3)$$

Proof

Let $A \in \hat{\mathcal{F}}_s$ and $x, y \in S$. Then

$$\begin{aligned} \mathbb{P}(\hat{X}_t = y \mid \hat{X}_s = x, A) &= \frac{\mathbb{P}(\hat{X}_t = y, \hat{X}_s = x, A)}{\mathbb{P}(\hat{X}_s = x, A)} = \frac{\mathbb{P}(X_{h-t} = y, X_{h-s} = x, A)}{\mathbb{P}(X_{h-s} = x, A)} \\ &= \frac{\mathbb{P}(A \mid X_{h-t} = y, X_{h-s} = x) \mathbb{P}(X_{h-s} = x \mid X_{h-t} = y) \mathbb{P}(X_{h-t} = y)}{\mathbb{P}(A \mid X_{h-s} = x) \mathbb{P}(X_{h-s} = x)} \end{aligned}$$

But $A \in \sigma\{X_r : r \in [h-s, h]\}$ and $h-t < h-s$, so by the Markov property for \mathbf{X} ,

$$\mathbb{P}(A \mid X_{h-t} = y, X_{h-s} = x) = \mathbb{P}(A \mid X_{h-s} = x) \quad (16.19.4)$$

By the time homogeneity of \mathbf{X} , $\mathbb{P}(X_{h-s} = x \mid X_{h-t} = y) = P_{t-s}(y, x)$. Substituting and simplifying gives

$$\mathbb{P}(\hat{X}_t = y \mid \hat{X}_s = x, A) = \frac{\mathbb{P}(X_{h-t} = y)}{\mathbb{P}(X_{h-s} = x)} P_{t-s}(y, x) \quad (16.19.5)$$

However, the backwards chain will be time homogeneous if X_0 has an invariant distribution.

Suppose that \mathbf{X} is positive recurrent, with (unique) invariant probability density function f . If X_0 has the invariant distribution, then $\hat{\mathbf{X}}$ is a time-homogeneous Markov chain. The transition matrix at time $t \in [0, \infty)$ (for every terminal time $h \geq t$), is given by

$$\hat{P}_t(x, y) = \frac{f(y)}{f(x)} P_t(y, x), \quad (x, y) \in S^2 \quad (16.19.6)$$

Proof

This follows from the result [above](#). Recall that if X_0 has PDF f , then X_{h-t} and X_{h-s} also have PDF f .

The previous result holds in the limit of the terminal time, regardless of the initial distribution.

Suppose again that \mathbf{X} is positive recurrent, with (unique) invariant probability density function f . Regardless of the distribution of X_0 ,

$$\mathbb{P}(\hat{X}_{s+t} = y \mid \hat{X}_s = x) \rightarrow \frac{f(y)}{f(x)} P_t(y, x) \text{ as } h \rightarrow \infty \quad (16.19.7)$$

Proof

This follows from the conditional probability [above](#) and our study of the limiting behavior of continuous-time Markov chains. Since \mathbf{X} is irreducible and positive recurrent, $\mathbb{P}(X_{h-s} = x) \rightarrow f(x)$ and $\mathbb{P}(X_{h-t} = y) \rightarrow f(y)$ as $h \rightarrow \infty$ for every $x, y \in S$.

These three results are motivation for the definition that follows. We can generalize by defining the reversal of an irreducible Markov chain, as long as there is a positive, invariant function. Recall that a positive invariant function defines a positive measure on S , but of course not in general a probability measure.

Suppose that $g : S \rightarrow (0, \infty)$ is invariant for \mathbf{X} . The *reversal* of \mathbf{X} with respect to g is the Markov chain $\hat{\mathbf{X}} = \{\hat{X}_t : t \in [0, \infty)\}$ with transition semigroup $\hat{\mathbf{P}}$ defined by

$$\hat{P}_t(x, y) = \frac{g(y)}{g(x)} P_t(y, x), \quad (x, y) \in S^2, t \in [0, \infty) \quad (16.19.8)$$

Justification

We need to show that the definition makes sense, namely that $\hat{\mathbf{P}}$ defines a transition semigroup for a Markov chain $\hat{\mathbf{X}}$ satisfying the same assumptions that we have imposed on \mathbf{X} . First let $t \in [0, \infty)$. Since g is invariant for \mathbf{X} ,

$$\sum_{y \in S} \hat{P}_t(x, y) = \frac{1}{g(x)} \sum_{y \in S} g(y) P_t(y, x) = \frac{g(x)}{g(x)} = 1, \quad x \in S \quad (16.19.9)$$

Hence \hat{P}_t is a valid transition matrix. Next we show that the Chapman-Kolmogorov equations (the semigroup property) holds. Let $s, t \in [0, \infty)$ and $x, z \in S$. Then

$$\begin{aligned} \hat{P}_s \hat{P}_t(x, z) &= \sum_{y \in S} \hat{P}_s(x, y) \hat{P}_t(y, z) = \sum_{y \in S} \frac{g(y)}{g(x)} P_s(y, x) \frac{g(z)}{g(y)} P_t(z, y) \\ &= \frac{g(z)}{g(x)} \sum_{y \in S} P_t(z, y) P_s(y, x) = \frac{g(z)}{g(x)} P_{s+t}(z, x) = \hat{P}_{s+t}(x, z) \end{aligned}$$

Next note that $\hat{P}_t(x, x) = P_t(x, x)$ for every $x \in S$. Hence $\hat{P}_t(x, x) \rightarrow 1$ as $t \downarrow 0$ for x , so \hat{P} is also a standard transition semigroup. Note also that if P is uniform, then so is \hat{P} . Finally, since X is irreducible, $P_t(x, y) > 0$ for every $(x, y) \in S^2$ and $t \in (0, \infty)$. Since g is positive, it follows that $\hat{P}_t(y, x) > 0$ for every $(x, y) \in S^2$ and $t \in (0, \infty)$, and hence \hat{X} is also irreducible.

Recall that if g is a positive invariant function for X then so is cg for every constant $c \in (0, \infty)$. Note that g and cg generate the same reversed chain. So let's consider the cases:

Suppose again that X is a Markov chain satisfying the assumptions above.

1. If X is recurrent, then X always has a positive invariant function g , unique up to multiplication by positive constants. Hence the reversal of a recurrent chain X always exists and is unique, and so we can refer to the reversal of X without reference to the invariant function.
2. Even better, if X is positive recurrent, then there exists a unique invariant probability density function, and the reversal of X can be interpreted as the *time reversal* (relative to a time horizon) when X has the invariant distribution, as in the motivating [result](#) above.
3. If X is transient, then there may or may not exist a positive invariant function, and if one does exist, it may not be unique (up to multiplication by positive constants). So a transient chain may have no reversals or more than one.

Nonetheless, the general definition is natural, because most of the important properties of the reversed chain follow from the *basic balance equation* relating the transition semigroups P and \hat{P} , and the invariant function g :

$$g(x) \hat{P}_t(x, y) = g(y) P_t(y, x), \quad (x, y) \in S^2, \quad t \in [0, \infty) \quad (16.19.10)$$

We will see the balance equation repeated for other objects associated with the Markov chains.

Suppose again that $g: S \rightarrow (0, \infty)$ is invariant for X , and that \hat{X} is the time reversal of X with respect to g . Then

1. g is also invariant for \hat{X} .
2. X is the time reversal of \hat{X} with respect to g .

Proof

1. For $y \in S$,

$$g \hat{P}_t(y) = \sum_{x \in S} g(x) \hat{P}_t(x, y) = \sum_{x \in S} g(y) P_t(y, x) = g(y) \sum_{x \in S} P_t(y, x) = g(y) \quad (16.19.11)$$

2. This follows from the symmetry of the fundamental equation: $g(x) \hat{P}_t(x, y) = g(y) P_t(y, x)$ for $(x, y) \in S^2$ and $t \in [0, \infty)$.

In the balance equation for the transition semigroups, it's not really necessary to know a-priori that the function g is invariant, if we know the two transition semigroups.

Suppose that $g: S \rightarrow (0, \infty)$. Then g is invariant and the Markov chains X and \hat{X} are time reversals with respect to g if and only if

$$g(x) \hat{P}_t(x, y) = g(y) P_t(y, x), \quad (x, y) \in S^2, \quad t \in [0, \infty) \quad (16.19.12)$$

Proof

All that is left to show is that the balance equation implies that g is invariant. The computation is exactly the same as in the last result:

$$gP_t(x) = \sum_{y \in S} g(y)P_t(y, x) = \sum_{y \in S} g(x)\hat{P}_t(x, y) = g(x) \sum_{y \in S} \hat{P}_t(x, y) = g(x), \quad x \in S \quad (16.19.13)$$

Here is a slightly more complicated (but equivalent) version of the balance equation for the transition probabilities.

Suppose again that $g: S \rightarrow (0, \infty)$. Then g is invariant and the chains \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals with respect to g if and only if

$$g(x_1)\hat{P}_{t_1}(x_1, x_2)\hat{P}_{t_2}(x_2, x_3)\cdots\hat{P}_{t_n}(x_n, x_{n+1}) = g(x_{n+1})P_{t_n}(x_{n+1}, x_n)P_{t_{n-1}}(x_n, x_{n-1})\cdots P_{t_1}(x_2, x_1) \quad (16.19.14)$$

for all $n \in \mathbb{N}_+$, $(t_1, t_2, \dots, t_n) \in [0, \infty)^n$, and $(x_1, x_2, \dots, x_{n+1}) \in S^{n+1}$.

Proof

All that is necessary is to show that the basic balance equation implies the balance equation in the theorem. When $n = 1$, we have the basic balance equation itself:

$$g(x_1)\hat{P}_{t_1}(x_1, x_2) = g(x_2)P_{t_1}(x_2, x_1) \quad (16.19.15)$$

For $n = 2$,

$$g(x_1)\hat{P}_{t_1}(x_1, x_2)\hat{P}_{t_2}(x_2, x_3) = g(x_2)P_{t_1}(x_2, x_1)\hat{P}_{t_2}(x_2, x_3) = g(x_3)P_{t_2}(x_3, x_2)P_{t_1}(x_2, x_1) \quad (16.19.16)$$

Continuing in this manner (or using induction) gives the general result.

The balance equation holds for the potential matrices.

Suppose again that $g: S \rightarrow (0, \infty)$. Then g is invariant and the chains \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals with respect to g if and only if the potential matrices satisfy

$$g(x)\hat{U}_\alpha(x, y) = g(y)U_\alpha(y, x), \quad (x, y) \in S^2, \quad \alpha \in [0, \infty) \quad (16.19.17)$$

Proof

We just need to show that the balance equation for the transition semigroups is equivalent to the balance equation above for the potential matrices. Suppose first $g(x)\hat{P}_t(x, y) = g(y)P_t(y, x)$ for $t \in [0, \infty)$ and $(x, y) \in S^2$. Then

$$\begin{aligned} g(x)\hat{U}_\alpha(x, y) &= g(x) \int_0^\infty e^{-\alpha t} \hat{P}_t(x, y) dt = \int_0^\infty e^{-\alpha t} g(x)\hat{P}_t(x, y) dt \\ &= \int_0^\infty e^{-\alpha t} g(y)P_t(y, x) dt = g(y) \int_0^\infty e^{-\alpha t} P_t(y, x) dt = g(y)U_\alpha(y, x) \end{aligned}$$

Conversely, suppose that $g(x)\hat{U}_\alpha(x, y) = g(y)U_\alpha(y, x)$ for $(x, y) \in S^2$ and $\alpha \in [0, \infty)$. As above,

$$g(x)\hat{U}_\alpha(x, y) = \int_0^\infty e^{-\alpha t} g(x)\hat{P}_t(x, y) dt \quad (16.19.18)$$

So for fixed $(x, y) \in S^2$, the function $\alpha \mapsto g(x)\hat{U}_\alpha(x, y)$ is the Laplace transform of the time function $t \mapsto g(x)\hat{P}_t(x, y)$. Similarly, $\alpha \mapsto g(y)U_\alpha(y, x)$ is the Laplace transform of the $t \mapsto g(y)P_t(y, x)$. The Laplace transform of a continuous function uniquely determines the function so it follows that $g(x)\hat{P}_t(x, y) = g(y)P_t(y, x)$ for $t \in [0, \infty)$ and $(x, y) \in S^2$.

As a corollary, continuous-time chains that are time reversals are of the same type.

If \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals, then \mathbf{X} and $\hat{\mathbf{X}}$ are of the same type: transient, null recurrent, or positive recurrent.

Proof

Suppose that \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals with respect to the invariant function $g : S \rightarrow (0, \infty)$. Then from the previous result, $\hat{U}(x, x) = U(x, x)$ for $x \in S$. The chains are transient if the common potential is finite for each $x \in S$ and recurrent if the potential is infinite for each $x \in S$. Suppose that the chains are recurrent. Then g is unique up to multiplication by positive constants and the chains are both positive recurrent if $\sum_{x \in S} g(x) < \infty$ and both null recurrent if $\sum_{x \in S} g(x) = \infty$.

The balance equation extends to the infinitesimal generator matrices.

Suppose again that $g : S \rightarrow (0, \infty)$. Then g is invariant and the Markov chains \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals if and only if the infinitesimal generators satisfy

$$g(x)\hat{G}(x, y) = g(y)G(y, x), \quad (x, y) \in S^2 \quad (16.19.19)$$

Proof

We need to show that the balance equation for the transition semigroups is equivalent to the balance equation for the generators. Suppose first that $g(x)\hat{P}_t(x, y) = g(y)P_t(y, x)$ for $t \in [0, \infty)$ and $(x, y) \in S^2$. Taking derivatives with respect to t and using Kolmogorov's backward equation gives $g(x)\hat{G}\hat{P}_t(x, y) = g(y)GP_t(y, x)$ for $t \in [0, \infty)$ and $(x, y) \in S^2$. Evaluating at $t = 0$ gives $g(x)\hat{G}(x, y) = g(y)G(y, x)$. Conversely, suppose that $g(x)\hat{G}(x, y) = g(y)G(y, x)$ for $(x, y) \in S^2$. Then repeated application (or induction) shows that $g(x)\hat{G}^n(x, y) = g(y)G^n(y, x)$ for every $n \in \mathbb{N}$ and $(x, y) \in S^2$. If the transition matrices are uniform, we can express them as exponentials of the generators. Hence for $t \in [0, \infty)$ and $(x, y) \in S^2$,

$$\begin{aligned} g(x)\hat{P}_t(x, y) &= g(x)e^{t\hat{G}(x, y)} = g(x) \sum_{n=0}^{\infty} \frac{t^n}{n!} \hat{G}^n(x, y) = \sum_{n=0}^{\infty} \frac{t^n}{n!} g(x)\hat{G}^n(x, y) \\ &= \sum_{n=0}^{\infty} \frac{t^n}{n!} g(y)G^n(y, x) = g(y) \sum_{n=0}^{\infty} \frac{t^n}{n!} G^n(y, x) = g(y)e^{tG(y, x)} = g(y)P_t(y, x) \end{aligned}$$

This leads to further results and connections:

Suppose again that $g : S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} and $\hat{\mathbf{X}}$ are time reversals with respect to g if and only if

1. $\hat{\mathbf{X}}$ and \mathbf{X} have the same exponential parameter function λ .
2. The jump chains \mathbf{Y} and $\hat{\mathbf{Y}}$ are (discrete) time reversals with respect to λg .

Proof

The exponential parameter functions are related to the generator matrices by $\lambda(x) = -G(x, x)$ and $\hat{\lambda}(x) = -\hat{G}(x, x)$ for $x \in S$. The transition matrices for the jump chains are related to the generator matrices by $Q(x, y) = G(x, y)/\lambda(x)$ and $\hat{Q}(x, y) = \hat{G}(x, y)/\hat{\lambda}(x)$ for $(x, y) \in S^2$ with $x \neq y$. Hence conditions (a) and (b) are equivalent to

$$g(x)\hat{G}(x, y) = g(y)G(y, x), \quad (x, y) \in S^2 \quad (16.19.20)$$

Recall also from the general theory, that if g is invariant for \mathbf{X} then λg is invariant for the jump chain \mathbf{Y} .

In our original discussion of time reversal in the positive recurrent case, we could have argued that the previous results must be true. If we run the positive recurrent chain $\mathbf{X} = \{X_t : t \in [0, h]\}$ backwards in time to obtain the time reversed chain $\hat{\mathbf{X}} = \{\hat{X}_t : t \in [0, h]\}$, then the exponential parameters for $\hat{\mathbf{X}}$ must be the same as those for \mathbf{X} , and the jump chain $\hat{\mathbf{Y}}$ for $\hat{\mathbf{X}}$ must be the time reversal of the jump chain \mathbf{Y} for \mathbf{X} .

Reversible Chains

Clearly an interesting special case is when the time reversal of a continuous-time Markov chain is stochastically the same as the original chain. Once again, we assume that we have a regular Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ that is irreducible on the state space S , with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. As before, $\mathbf{U} = \{U_\alpha : \alpha \in [0, \infty)\}$ denotes the collection of potential matrices, and G the infinitesimal generator. Finally, λ denotes the exponential parameter function, $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ the jump chain, and Q the transition matrix of \mathbf{Y} . Here is the definition of reversibility:

Suppose that $g : S \rightarrow (0, \infty)$ is invariant for \mathbf{X} . Then \mathbf{X} is *reversible* with respect to g if the time reversed chain $\hat{\mathbf{X}} = \{\hat{X}_t : t \in [0, \infty)\}$ also has transition semigroup \mathbf{P} . That is,

$$g(x)P_t(x, y) = g(y)P_t(y, x), \quad (x, y) \in S^2, \quad t \in [0, \infty) \quad (16.19.21)$$

Clearly if \mathbf{X} is reversible with respect to g then \mathbf{X} is reversible with respect to cg for every $c \in (0, \infty)$. So here is another review of the cases:

Suppose that \mathbf{X} is a Markov chain satisfying the assumptions above.

1. If \mathbf{X} is recurrent, then there exists an invariant function $g: S \rightarrow (0, \infty)$ that is unique up to multiplication by positive constants. So \mathbf{X} is either reversible or not, and we do not have to reference the invariant function.
2. Even better, if \mathbf{X} is positive recurrent, then there exists a unique invariant probability density function f . Again, \mathbf{X} is either reversible or not, but if it is, then with the invariant distribution, the chain \mathbf{X} is stochastically the same, forward in time or backward in time.
3. If \mathbf{X} is transient, then a positive invariant function may or may not exist. If such a function does exist, it may not be unique, up to multiplication by positive constants. So in the transient case, \mathbf{X} may be reversible with respect to one invariant function but not with respect to others.

The following results are corollaries of the results above for time reversals. First, we don't need to know a priori that the function g is invariant.

Suppose that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if

$$g(x)P_t(x, y) = g(y)P_t(y, x), \quad (x, y) \in S^2, \quad t \in [0, \infty) \quad (16.19.22)$$

Suppose again that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if

$$g(x_1)P_{t_1}(x_1, x_2)P_{t_2}(x_2, x_3) \cdots P_{t_n}(x_n, x_{n+1}) = g(x_{n+1})P_{t_n}(x_{n+1}, x_n)P_{t_{n-1}}(x_n, x_{n-1}) \cdots P_{t_1}(x_2, x_1) \quad (16.19.23)$$

for all $n \in \mathbb{N}_+$, $(t_1, t_2, \dots, t_n) \in [0, \infty)^n$, and $(x_1, x_2, \dots, x_{n+1}) \in S^{n+1}$.

Suppose again that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if

$$g(x)U_\alpha(x, y) = g(y)U_\alpha(y, x), \quad (x, y) \in S^2, \quad \alpha \in [0, \infty) \quad (16.19.24)$$

Suppose again that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible with respect to g if and only if

$$g(x)G(x, y) = g(y)G(y, x), \quad (x, y) \in S^2 \quad (16.19.25)$$

Suppose again that $g: S \rightarrow (0, \infty)$. Then g is invariant and \mathbf{X} is reversible if and only if the jump chain \mathbf{Y} is reversible with respect to λg .

Recall that \mathbf{X} is recurrent if and only if the jump chain \mathbf{Y} is recurrent. In this case, the invariant functions for \mathbf{X} and \mathbf{Y} exist and are unique up to positive constants. So in this case, the previous theorem states that \mathbf{X} is reversible if and only if \mathbf{Y} is reversible. In the positive recurrent case (the most important case), the following theorem gives a condition for reversibility that does not directly reference the invariant distribution. The condition is known as the *Kolmogorov cycle condition*, and is named for Andrei Kolmogorov

Suppose that \mathbf{X} is positive recurrent. Then \mathbf{X} is reversible if and only if for every sequence of distinct states (x_1, x_2, \dots, x_n) ,

$$G(x_1, x_2)G(x_2, x_3) \cdots G(x_{n-1}, x_n)G(x_n, x_1) = G(x_1, x_n)G(x_n, x_{n-1}) \cdots G(x_3, x_2)G(x_2, x_1) \quad (16.19.26)$$

Proof

Suppose that \mathbf{X} is reversible, and let f denote the invariant PDF of \mathbf{X} . Then $G(x, y) = \frac{f(y)}{f(x)}G(x, y)$ for $(x, y) \in S^2$. Substituting gives the Kolmogorov cycle condition. Conversely, suppose that the Kolmogorov cycle condition holds for \mathbf{X} . Recall that $G(x, y) = \lambda(x)Q(x, y)$ for $(x, y) \in S^2$. Substituting into the cycle condition for \mathbf{X} gives the cycle condition for \mathbf{Y} . Hence \mathbf{Y} is reversible and therefore so is \mathbf{X} .

Note that the Kolmogorov cycle condition states that the transition rate of visiting states $(x_2, x_3, \dots, x_n, x_1)$ in sequence, starting in state x_1 is the same as the transition rate of visiting states $(x_n, x_{n-1}, \dots, x_2, x_1)$ in sequence, starting in state x_1 . The cycle

condition is also known as the *balance equation for cycles*.




Figure 16.19.1: The Kolmogorov cycle condition

Applications and Exercises

The Two-State Chain

The continuous-time, two-state chain has been studied in our previous sections on continuous-time chains, so naturally we are interested in time reversal.

Consider the continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1\}$ with transition rate $a \in (0, \infty)$ from 0 to 1 and transition rate $b \in (0, \infty)$ from 1 to 0. Show that \mathbf{X} is reversible

1. Using the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$.
2. Using the resolvent $\mathbf{U} = \{U_\alpha : \alpha \in (0, \infty)\}$.
3. Using the generator matrix G .

Solutions

First note that \mathbf{X} is irreducible since $a > 0$ and $b > 0$. Since S is finite, \mathbf{X} is positive recurrent.

1. Recall that

$$P_t = \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix} - \frac{1}{a+b} e^{-(a+b)t} \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}, \quad t \in [0, \infty) \quad (16.19.27)$$

All we have to do is find a positive function g on S with the property that $g(0)P_t(0, 1) = g(1)P_t(1, 0)$. The other conditions are trivially satisfied. Note that $g(0) = b$, $g(1) = a$ satisfies the property. It follows that g is invariant for \mathbf{X} , unique up to multiplication by positive constants, and that \mathbf{X} is reversible.

2. Recall that

$$U_\alpha = \frac{1}{\alpha(a+b)} \begin{bmatrix} b & a \\ b & a \end{bmatrix} - \frac{1}{(\alpha+a+b)(a+b)} \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}, \quad \alpha \in (0, \infty) \quad (16.19.28)$$

Again, we just need to find a positive function g on S with the property that $g(0)U_\alpha(0, 1) = g(1)U_\alpha(1, 0)$. The other conditions are trivially satisfied. The function g in part (a) satisfies, the condition, which of course must be the case.

3. Recall that $G = \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}$. Once again, we just need to find a positive function g on S with the property that $g(0)G(0, 1) = g(1)G(1, 0)$. The function g given in (a) satisfies the condition. Note that this procedure is the easiest of the three.

Of course, the invariant PDF f is $f(0) = b/(a+b)$, $f(1) = a/(a+b)$.

Computational Exercises

The Markov chain in the following exercise has also been studied in previous sections.

Consider the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1, 2\}$ with exponential parameter function $\lambda = (4, 1, 3)$ and jump transition matrix

$$Q = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 1 & 0 & 0 \\ \frac{1}{3} & \frac{2}{3} & 0 \end{bmatrix} \quad (16.19.29)$$

Give each of the following for the time reversed chain $\hat{\mathbf{X}}$:

1. The state graph.

2. The semigroup of transition matrices $\hat{P} = \{\hat{P}_t : t \in [0, \infty)\}$.
3. The resolvent of potential matrices $\hat{U} = \{\hat{U}_\alpha : \alpha \in (0, \infty)\}$.
4. The generator matrix \hat{G} .
5. The transition matrix of the jump chain \hat{Y} .

Solutions

Note that the chain is irreducible, and since S is finite, positive recurrent. We found previously that an invariant function (unique up to multiplication by positive constants) is $g = (3, 10, 2)$.

1. The edge set is $\hat{E} = \{(1, 0), (2, 0), (0, 1), (0, 2), (1, 2)\}$. The exponential parameter function $\lambda = (4, 1, 3)$ is the same as for \mathbf{X} .
2. The transition matrix at $t \in [0, \infty)$ is

$$\hat{P}_t = \frac{1}{15} \begin{bmatrix} 3 + 12e^{-5t} & 10 - 10e^{-5t} & 2 - 2e^{-5t} \\ 3 - 3e^{-3t} & 10 + 5e^{-3t} & 2 - 2e^{-3t} \\ 3 - 18e^{-5t} + 15e^{-3t} & 10 + 15e^{-5t} - 25e^{-3t} & 2 + 3e^{-5t} + 10e^{-3t} \end{bmatrix} \quad (16.19.30)$$

3. The potential matrix at $\alpha \in (0, \infty)$ is

$$\hat{U}_\alpha = \frac{1}{15\alpha + 8\alpha^2 + \alpha^3} \begin{bmatrix} 3 + 4\alpha + \alpha^2 & 10 + \frac{10}{3}\alpha & 2 + \frac{2}{3}\alpha \\ 3 + \frac{3}{5}\alpha & 10 + 7\alpha + \alpha^2 & 2 + \frac{2}{5}\alpha \\ 3 + 3\alpha & 10 & 2 + 5\alpha + \alpha^2 \end{bmatrix} \quad (16.19.31)$$

4. The generator matrix is

$$\hat{G} = \begin{bmatrix} -4 & \frac{10}{3} & \frac{2}{3} \\ \frac{3}{5} & -1 & \frac{2}{5} \\ 3 & 0 & -3 \end{bmatrix} \quad (16.19.32)$$

5. The transition matrix of the jump chain is

$$\hat{Q} = \begin{bmatrix} 0 & \frac{5}{6} & \frac{1}{6} \\ \frac{3}{5} & 0 & \frac{2}{5} \\ 1 & 0 & 0 \end{bmatrix} \quad (16.19.33)$$

Special Models

Read the discussion of time reversal for chains subordinate to the Poisson process.

Read the discussion of time reversal for continuous-time birth-death chains.

This page titled [16.19: Time Reversal in Continuous-Time Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.20: Chains Subordinate to the Poisson Process

Basic Theory

Introduction

Recall that the standard Poisson process with rate parameter $r \in (0, \infty)$ involves three interrelated stochastic processes. First the sequence of *interarrival times* $\mathbf{T} = (T_1, T_2, \dots)$ is independent, and each variable has the exponential distribution with parameter r . Next, the sequence of *arrival times* $\boldsymbol{\tau} = (\tau_0, \tau_1, \dots)$ is the partial sum sequence associated with the interarrival sequence \mathbf{T} :

$$\tau_n = \sum_{i=1}^n T_i, \quad n \in \mathbb{N} \quad (16.20.1)$$

For $n \in \mathbb{N}_+$, the arrival time τ_n has the gamma distribution with parameters n and r . Finally, the Poisson counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is defined by

$$N_t = \max\{n \in \mathbb{N} : \tau_n \leq t\}, \quad t \in [0, \infty) \quad (16.20.2)$$

so that N_t is the number of arrivals in $(0, t]$ for $t \in [0, \infty)$. The counting variable N_t has the Poisson distribution with parameter rt for $t \in [0, \infty)$. The counting process \mathbf{N} and the arrival time process $\boldsymbol{\tau}$ are inverses in the sense that $\tau_n \leq t$ if and only if $N_t \geq n$ for $t \in [0, \infty)$ and $n \in \mathbb{N}$. The Poisson counting process can be viewed as a continuous-time Markov chain.

Suppose that X_0 takes values in \mathbb{N} and is independent of \mathbf{N} . Define $X_t = X_0 + N_t$ for $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on \mathbb{N} with exponential parameter function given by $\lambda(x) = r$ for $x \in \mathbb{N}$ and jump transition matrix Q given by $Q(x, x+1) = 1$ for $x \in S$.

Proof

This follows directly from the basic structure of a continuous-time Markov chain. Given $X_t = x$, the holding time in state $x \in \mathbb{N}$ is exponential with parameter r , and the next state is deterministically $x+1$. Note that the addition of the variable X_0 is just to allow us the freedom of arbitrary initial distributions on the state space, as is routine with Markov processes.

Note that the Poisson process, viewed as a Markov chain is a pure birth chain. Clearly we can generalize this continuous-time Markov chain in a simple way by allowing a general embedded jump chain.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Markov chain with (countable) state space S , and with constant exponential parameter $\lambda(x) = r \in (0, \infty)$ for $x \in S$, and jump transition matrix Q . Then \mathbf{X} is said to be *subordinate* to the Poisson process with rate parameter r .

1. The transition times (τ_1, τ_2, \dots) are the arrival times of the Poisson process with rate r .
2. The inter-transition times $(\tau_1, \tau_2 - \tau_1, \dots)$ are the inter-arrival times of the Poisson process with rate r (independent, and each with the exponential distribution with rate r).
3. $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is the Poisson counting process, where N_t is the number of transitions in $(0, t]$ for $t \in [0, \infty)$.
4. The Poisson process and the jump chain $\mathbf{Y} = (Y_0, Y_1, \dots)$ are independent, and $X_t = Y_{N_t}$ for $t \in [0, \infty)$.

Proof

These results all follow from the basic structure of a continuous-time Markov chain.

Since all states are stable, note that we must have $Q(x, x) = 0$ for $x \in S$. Note also that for $x, y \in S$ with $x \neq y$, the exponential rate parameter for the transition from x to y is $\mu(x, y) = rQ(x, y)$. Conversely suppose that $\mu : S^2 \rightarrow (0, \infty)$ satisfies $\mu(x, x) = 0$ and $\sum_{y \in S} \mu(x, y) = r$ for every $x \in S$. Then the Markov chain with transition rates given by μ is subordinate to the Poisson process with rate r . It's easy to construct a Markov chain subordinate to the Poisson process.

Suppose that $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is a Poisson counting process with rate $r \in (0, \infty)$ and that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a discrete-time Markov chain on S , independent of \mathbf{N} , whose transition matrix satisfies $Q(x, x) = 0$ for every $x \in S$. Let $X_t = Y_{N_t}$ for $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain subordinate to the Poisson process.

Generator and Transition Matrices

Next let's find the generator matrix and the transition semigroup. Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on S subordinate to the Poisson process with rate $r \in (0, \infty)$ and with jump transition matrix Q . As usual, let $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ denote the transition semigroup and G the infinitesimal generator.

The generator matrix G of \mathbf{X} is $G = r(Q - I)$. Hence for $t \in [0, \infty)$

1. The Kolmogorov backward equation is $P'_t = r(Q - I)P_t$
2. The Kolmogorov forward equation is $P'_t = rP_t(Q - I)$

Proof

This follows directly from the general theory since $G(x, x) = -\lambda(x) = -r$ for $x \in S$ and $G(x, y) = \lambda(x)Q(x, y) = rQ(x, y)$ for distinct $x, y \in S$.

There are several ways to find the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$. The best way is a probabilistic argument using the underlying Poisson process.

For $t \in [0, \infty)$, the transition matrix P_t is given by

$$P_t = \sum_{n=0}^{\infty} e^{-rt} \frac{(rt)^n}{n!} Q^n \quad (16.20.3)$$

Proof from the underlying Poisson process

Let N_t denote the number of transitions in $(0, t]$ for $t \in [0, \infty)$, so that $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is the Poisson counting process. Let $\mathbf{Y} = (Y_0, Y_1, \dots)$ denote the jump chain, with transition matrix Q . Then \mathbf{N} and \mathbf{Y} are independent, and $X_t = Y_{N_t}$ for $t \in [0, \infty)$. Conditioning we have

$$\begin{aligned} P_t(x, y) &= \mathbb{P}(X_t = y \mid X_0 = x) = \mathbb{P}(Y_{N_t} = y \mid Y_0 = x) \\ &= \sum_{n=0}^{\infty} \mathbb{P}(Y_{N_t} = y \mid N_t = n, Y_0 = y) \mathbb{P}(N_t = n \mid Y_0 = y) \\ &= \sum_{n=0}^{\infty} \mathbb{P}(Y_n = y \mid Y_0 = x) \mathbb{P}(N_t = n) = \sum_{n=0}^{\infty} e^{-rt} \frac{(rt)^n}{n!} Q^n(x, y) \end{aligned}$$

Proof using the generator matrix

Note first that for $n \in \mathbb{N}$,

$$G^n = [r(Q - I)]^n = r^n \sum_{k=0}^n \binom{n}{k} (-1)^{n-k} Q^k \quad (16.20.4)$$

Hence

$$\begin{aligned} P_t &= e^{tG} = \sum_{n=0}^{\infty} \frac{t^n}{n!} G^n = \sum_{n=0}^{\infty} \frac{t^n}{n!} r^n \sum_{k=0}^n \binom{n}{k} (-1)^{n-k} Q^k \\ &= \sum_{n=0}^{\infty} \sum_{k=0}^n \frac{(rt)^n}{k!(n-k)!} (-1)^{n-k} Q^k = \sum_{k=0}^{\infty} \sum_{n=k}^{\infty} \frac{(rt)^n}{k!(n-k)!} (-1)^{n-k} Q^k \\ &= \sum_{k=0}^{\infty} \frac{(rt)^k}{k!} Q^k \sum_{n=k}^{\infty} \frac{1}{(n-k)!} (-rt)^{n-k} = \sum_{k=0}^{\infty} e^{-rt} \frac{(rt)^k}{k!} Q^k \end{aligned}$$

Potential Matrices

Next let's find the potential matrices. As with the transition matrices, we can do this in (at least) two different ways.

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on S subordinate to the Poisson process with rate $r \in (0, \infty)$ and with jump transition matrix Q . For $\alpha \in (0, \infty)$, the potential matrix U_α of \mathbf{X} is

$$U_\alpha = \frac{1}{\alpha + r} \sum_{n=0}^{\infty} \left(\frac{r}{\alpha + r} \right)^n Q^n \quad (16.20.5)$$

Proof from the definition

Using the previous result,

$$\begin{aligned} U_\alpha(x, y) &= \int_0^\infty e^{-\alpha t} P_t(x, y) dt = \int_0^\infty e^{-\alpha t} \sum_{n=0}^{\infty} e^{-rt} \frac{(rt)^n}{n!} Q^n(x, y) dt \\ &= \sum_{n=0}^{\infty} Q^n(x, y) \frac{r^n}{n!} \int_0^\infty e^{-(r+\alpha)t} t^n dt \end{aligned}$$

The interchange of sum and integral is justified since the terms are nonnegative. Using the change of variables $s = (r + \alpha)t$ gives

$$U_\alpha(x, y) = \frac{1}{\alpha + r} \sum_{n=0}^{\infty} \left(\frac{r}{\alpha + r} \right)^n \frac{1}{n!} Q^n(x, y) \int_0^\infty e^{-st} s^n ds \quad (16.20.6)$$

The last integral is $n!$.

Proof using the generator

From the result above,

$$\alpha I - G = \alpha I - r(Q - I) = (\alpha + r)I - rQ = (\alpha + r) \left(I - \frac{r}{\alpha + r} Q \right) \quad (16.20.7)$$

Since $\left\| \frac{r}{\alpha + r} Q \right\| = \frac{r}{\alpha + r} < 1$ we have

$$(\alpha I - G)^{-1} = \frac{1}{\alpha + r} \left(I - \frac{r}{\alpha + r} Q \right)^{-1} = \frac{1}{\alpha + r} \sum_{n=0}^{\infty} \left(\frac{r}{\alpha + r} \right)^n Q^n \quad (16.20.8)$$

Recall that for $p \in (0, 1)$, the p -potential matrix of the jump chain \mathbf{Y} is $R_p = \sum_{n=0}^{\infty} p^n Q^n$. Hence we have the following nice relationship between the potential matrix of \mathbf{X} and the potential matrix of \mathbf{Y} :

$$U_\alpha = \frac{1}{\alpha + r} R_{r/(\alpha + r)} \quad (16.20.9)$$

Next recall that $\alpha U_\alpha(x, \cdot)$ is the probability density function of X_T given $X_0 = x$, where T has the exponential distribution with parameter α and is independent of \mathbf{X} . On the other hand, $\alpha U_\alpha(x, \cdot) = (1 - p)R_p(x, \cdot)$ where $p = r/(\alpha + r)$. We know from our study of discrete potentials that $(1 - p)R_p(x, \cdot)$ is the probability density function of Y_M where M has the geometric distribution on \mathbb{N} with parameter $1 - p$ and is independent of \mathbf{Y} . But also $X_T = Y_{N_T}$. So it follows that if T has the exponential distribution with parameter α , $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is a Poisson process with rate r , and is independent of T , then N_T has the geometric distribution on \mathbb{N} with parameter $\alpha/(\alpha + r)$. Of course, we could easily verify this directly, but it's still fun to see such connections.

Limiting Behavior and Stationary Distributions

Once again, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on S subordinate to the Poisson process with rate $r \in (0, \infty)$ and with jump transition matrix Q . Let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ denote the jump process. The limiting behavior and stationary distributions of \mathbf{X} are closely related to those of \mathbf{Y} .

Suppose that \mathbf{X} (and hence \mathbf{Y}) are irreducible and positive recurrent

1. $g : S \rightarrow (0, \infty)$ is invariant for \mathbf{X} if and only if g is invariant for \mathbf{Y} .
2. f is an invariant probability density function for \mathbf{X} if and only if f is an invariant probability density function for \mathbf{Y} .
3. \mathbf{X} is null recurrent if and only if \mathbf{Y} is null recurrent, and in this case, $\lim_{n \rightarrow \infty} Q^n(x, y) = \lim_{t \rightarrow \infty} P_t(x, y) = 0$ for $(x, y) \in S^2$.
4. \mathbf{X} is positive recurrent if and only if \mathbf{Y} is positive recurrent. If \mathbf{Y} is aperiodic, then $\lim_{n \rightarrow \infty} Q^n(x, y) = \lim_{t \rightarrow \infty} P_t(x, y) = f(y)$ for $(x, y) \in S^2$, where f is the invariant probability density function.

Proof

All of these results follow from the basic theory of stationary and limiting distributions for continuous-time chains, and the fact that the exponential parameter function λ is constant.

Time Reversal

Once again, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on S subordinate to the Poisson process with rate $r \in (0, \infty)$ and with jump transition matrix Q . Let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ denote the jump process. We assume that \mathbf{X} (and hence \mathbf{Y}) are irreducible. The time reversal of \mathbf{X} is closely related to that of \mathbf{Y} .

Suppose that $g : S \rightarrow (0, \infty)$ is invariant for \mathbf{X} . The time reversal $\hat{\mathbf{X}}$ with respect to g is also subordinate to the Poisson process with rate r . The jump chain $\hat{\mathbf{Y}}$ of $\hat{\mathbf{X}}$ is the (discrete) time reversal of \mathbf{Y} with respect to g .

Proof

From the previous result, g is also invariant for \mathbf{Y} . From the general theory of time reversal, $\hat{\mathbf{X}}$ has the same exponential parameter function as \mathbf{X} (namely the constant function r) and so is also subordinate to the Poisson process with rate r . Finally, the jump chain $\hat{\mathbf{Y}}$ of $\hat{\mathbf{X}}$ is the reversal of \mathbf{Y} with respect to rg and hence also with respect to g .

In particular, \mathbf{X} is reversible with respect to g if and only if \mathbf{Y} is reversible with respect to g . As noted earlier, \mathbf{X} and \mathbf{Y} are of the same type: both transient or both null recurrent or both positive recurrent. In the recurrent case, there exists a positive invariant function that is unique up to multiplication by constants. In this case, the reversal of \mathbf{X} is unique, and is the chain subordinate to the Poisson process with rate r whose jump chain is the reversal of \mathbf{Y} .

Uniform Chains

In the [construction above](#) for a Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ that is subordinate to the Poisson process with rate r and jump transition kernel Q , we assumed of course that $Q(x, x) = 0$ for every $x \in S$. So there are no absorbing states and the sequence (τ_1, τ_2, \dots) of arrival times of the Poisson process are the jump times of the chain \mathbf{X} . However in our introduction to continuous-time chains, we saw that the general construction of a chain starting with the function λ and the transition matrix Q works without this assumption on Q , although the exponential parameters and transition probabilities change. The same idea works here.

Suppose that $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ is a counting Poisson process with rate $r \in (0, \infty)$ and that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a discrete-time Markov chain with transition matrix Q on $S \times S$ satisfying $Q(x, x) < 1$ for $x \in S$. Assume also that \mathbf{N} and \mathbf{Y} are independent. Define $X_t = Y_{N_t}$ for $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-Markov chain with exponential parameter function $\lambda(x) = r[1 - Q(x, x)]$ for $x \in S$ and jump transition matrix \tilde{Q} given by

$$\tilde{Q}(x, y) = \frac{Q(x, y)}{1 - Q(x, x)}, \quad (x, y) \in S^2, x \neq y \quad (16.20.10)$$

Proof

This follows from the result in the introduction.

The Markov chain constructed above is no longer a chain subordinate to the Poisson process by our [definition above](#), since the exponential parameter function is not constant, and the transition times of \mathbf{X} are no longer the arrival times of the Poisson process. Nonetheless, many of the basic results above still apply.

Let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be the Markov chain constructed in the previous theorem. Then

1. For $t \in [0, \infty)$, the transition matrix P_t is given by

$$P_t = \sum_{n=0}^{\infty} e^{-rt} \frac{(rt)^n}{n!} Q^n \quad (16.20.11)$$

2. For $\alpha \in (0, \infty)$, the α potential matrix is given by

$$U_\alpha = \frac{1}{\alpha + r} \sum_{n=0}^{\infty} \left(\frac{r}{\alpha + r} \right)^n Q^n \quad (16.20.12)$$

3. The generator matrix is $G = r(Q - I)$
4. $g : S \rightarrow (0, \infty)$ is invariant for \mathbf{X} if and only if g is invariant for \mathbf{Y} .

Proof

The proofs are just as before.

It's a remarkable fact that every continuous-time Markov chain with bounded exponential parameters can be constructed as in the last theorem, a process known as *uniformization*. The name comes from the fact that in the construction, the exponential parameters become constant, but at the expense of allowing the embedded discrete-time chain to jump from a state back to that state. To review the definition, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on S with transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$, exponential parameter function λ and jump transition matrix Q . Then \mathbf{P} is *uniform* if $P_t(x, x) \rightarrow 1$ as $t \downarrow 0$ uniformly in x , or equivalently if λ is bounded.

Suppose that $\lambda : S \rightarrow (0, \infty)$ is bounded and that Q is a transition matrix on S with $Q(x, x) = 0$ for every $x \in S$. Let $r \in (0, \infty)$ be an upper bound on λ and $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ a Poisson counting process with rate r . Define the transition matrix \hat{Q} on S by

$$\begin{aligned} \hat{Q}(x, x) &= 1 - \frac{\lambda(x)}{r} \quad x \in S \\ \hat{Q}(x, y) &= \frac{\lambda(x)}{r} Q(x, y) \quad (x, y) \in S^2, x \neq y \end{aligned}$$

and let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ be a discrete-time Markov chain with transition matrix \hat{Q} , independent of \mathbf{N} . Define $X_t = Y_{N_t}$ for $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain with exponential parameter function λ and jump transition matrix Q .

Proof

Note that $\hat{Q}(x, y) \geq 0$ for every $(x, y) \in S^2$ and $\sum_{y \in S} \hat{Q}(x, y) = 1$ for every $x \in S$. Thus \hat{Q} is a transition matrix on S . Note also that $\hat{Q}(x, x) < 1$ for every $x \in S$. By construction, $\lambda(x) = r[1 - \hat{Q}(x, x)]$ for $x \in S$ and

$$Q(x, y) = \frac{\hat{Q}(x, y)}{1 - \hat{Q}(x, x)}, \quad (x, y) \in S^2, x \neq y \quad (16.20.13)$$

So the result now follows from the [theorem above](#).

Note in particular that if the state space S is finite then of course λ is bounded so the previous theorem applies. The theorem is useful for simulating a continuous-time Markov chain, since the Poisson process and discrete-time chains are simple to simulate. In addition, we have nice representations for the transition matrices, potential matrices, and the generator matrix.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on S with bounded exponential parameter function $\lambda : S \rightarrow (0, \infty)$ and jump transition matrix Q . Define r and \hat{Q} as in the last theorem. Then

1. For $t \in [0, \infty)$, the transition matrix P_t is given by

$$P_t = \sum_{n=0}^{\infty} e^{-rt} \frac{(rt)^n}{n!} \hat{Q}^n \quad (16.20.14)$$

2. For $\alpha \in (0, \infty)$, the α potential matrix is given by

$$U_\alpha = \frac{1}{\alpha + r} \sum_{n=0}^{\infty} \left(\frac{r}{\alpha + r} \right)^n \hat{Q}^n \quad (16.20.15)$$

3. The generator matrix is $G = r(\hat{Q} - I)$
4. $g : S \rightarrow (0, \infty)$ is invariant for \mathbf{X} if and only if g is invariant for \hat{Q} .

Proof

These results follow from the [theorem above](#).

Examples

The Two-State Chain

The following exercise applies the uniformization method to the two-state chain.

Consider the continuous-time Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{0, 1\}$ with exponential parameter function $\lambda = (a, b)$, where $a, b \in (0, \infty)$. Thus, states 0 and 1 are stable and the jump chain has transition matrix

$$Q = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad (16.20.16)$$

Let $r = a + b$, an upper bound on λ . Show that

1. $\hat{Q} = \frac{1}{a+b} \begin{bmatrix} b & a \\ b & a \end{bmatrix}$
2. $G = \begin{bmatrix} -a & a \\ b & -b \end{bmatrix}$
3. $P_t = \hat{Q} - \frac{1}{a+b} e^{-(a+b)t} G$ for $t \in [0, \infty)$
4. $U_\alpha = \frac{1}{\alpha} \hat{Q} - \frac{1}{(\alpha+a+b)(a+b)} G$ for $\alpha \in (0, \infty)$

Proof

The form of \hat{Q} follows easily from the definition [above](#). Note that the rows of \hat{Q} are the invariant PDF. It then follows that $\hat{Q}^n = \hat{Q}$ for $n \in \mathbb{N}_+$. The results for the transition matrix P_t and the potential U_α then follow easily from the [theorem above](#).

Although we have obtained all of these results for the two-state chain before, the derivation based on uniformization is the easiest.

This page titled [16.20: Chains Subordinate to the Poisson Process](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.21: Continuous-Time Birth-Death Chains

Basic Theory

Introduction

A continuous-time *birth-death chain* is a simple class of Markov chains on a subset of \mathbb{Z} with the property that the only possible transitions are to increase the state by 1 (*birth*) or decrease the state by 1 (*death*). It's easiest to define the birth-death process in terms of the exponential transition rates, part of the basic structure of continuous-time Markov chains.

Suppose that S is an integer interval (that is, a set of consecutive integers), either finite or infinite. The *birth-death chain* with birth rate function $\alpha : S \rightarrow [0, \infty)$ and death rate function $\beta : S \rightarrow [0, \infty)$ is the Markov chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on S with transition rate $\alpha(x)$ from x to $x + 1$ and transition rate $\beta(x)$ from x to $x - 1$, for $x \in S$.

If S has a minimum element m , then of course we must have $\beta(m) = 0$. If $\alpha(m) = 0$ also, then the boundary point m is absorbing. Similarly, if S has a maximum element n then we must have $\alpha(n) = 0$. If $\beta(n) = 0$ also then the boundary point n is absorbing. If $x \in S$ is not a boundary point, then typically we have $\alpha(x) + \beta(x) > 0$, so that x is stable. If $\beta(x) = 0$ for all $x \in S$, then \mathbf{X} is a *pure birth process*, and similarly if $\alpha(x) = 0$ for all $x \in S$ then \mathbf{X} is a *pure death process*. From the transition rates, it's easy to compute the parameters of the exponential holding times in a state and the transition matrix of the embedded, discrete-time jump chain.

Consider again the birth-death chain \mathbf{X} on S with birth rate function α and death rate function β . As usual, let λ denote the exponential parameter function and Q the transition matrix for the jump chain.

1. $\lambda(x) = \alpha(x) + \beta(x)$ for $x \in S$
2. If $x \in S$ is stable, so that $\alpha(x) + \beta(x) > 0$, then

$$Q(x, x+1) = \frac{\alpha(x)}{\alpha(x) + \beta(x)}, \quad Q(x, x-1) = \frac{\beta(x)}{\alpha(x) + \beta(x)} \quad (16.21.1)$$

Note that jump chain $\mathbf{Y} = (Y_0, Y_1, \dots)$ is a discrete-time birth death chain. The probability functions p , q , and r of \mathbf{Y} are given as follows: If $x \in S$ is stable then

$$\begin{aligned} p(x) &= Q(x, x+1) = \frac{\alpha(x)}{\alpha(x) + \beta(x)} \\ q(x) &= Q(x, x-1) = \frac{\beta(x)}{\alpha(x) + \beta(x)} \\ r(x) &= Q(x, x) = 0 \end{aligned}$$

If x is absorbing then of course $p(x) = q(x) = 0$ and $r(x) = 1$. Except for the initial state, the jump chain \mathbf{Y} is deterministic for a pure birth process, with $Q(x, x) = 1$ if x is absorbing and $Q(x, x+1) = 1$ if x is stable. Similarly, except for the initial state, \mathbf{Y} is deterministic for a pure death process, with $Q(x, x) = 1$ if x is absorbing and $Q(x, x-1) = 1$ if x is stable. Note that the Poisson process with rate parameter $r \in (0, \infty)$, viewed as a continuous-time Markov chain, is a pure birth process on \mathbb{N} with birth function $\alpha(x) = r$ for each $x \in \mathbb{N}$. More generally, a birth death process with $\lambda(x) = \alpha(x) + \beta(x) = r$ for all $x \in S$ is also subordinate to the Poisson process with rate r .

Note that λ is bounded if and only if α and β are bounded (always the case if S is finite), and in this case the birth-death chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is uniform. If λ is unbounded, then \mathbf{X} may not even be regular, as an [example below](#) shows. Recall that a sufficient condition for \mathbf{X} to be regular when S is infinite is

$$\sum_{x \in S_+} \frac{1}{\lambda(x)} = \sum_{x \in S_+} \frac{1}{\alpha(x) + \beta(x)} = \infty \quad (16.21.2)$$

where $S_+ = \{x \in S : \lambda(x) = \alpha(x) + \beta(x) > 0\}$ is the set of stable states. Except for the aforementioned example, we will restrict our study to regular birth-death chains.

Infinitesimal Generator and Transition Matrices

Suppose again that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time birth-death chain on an interval $S \subseteq \mathbb{Z}$ with birth rate function α and death rate function β . As usual, we will let P_t denote the transition matrix at time $t \in [0, \infty)$ and G the infinitesimal generator. As always, the infinitesimal generator gives the same information as the exponential parameter function and the jump transition matrix, but in a more compact and useful form.

The generator matrix G is given by

$$G(x, x) = -[\alpha(x) + \beta(x)], \quad G(x, x+1) = \alpha(x), \quad G(x, x-1) = \beta(x), \quad x \in S \quad (16.21.3)$$

Proof

This follows from the general theory, since $G(x, x) = -\lambda(x)$ for $x \in S$ and $G(x, y) = \lambda(x)Q(x, y)$ for $(x, y) \in S^2$ with $x \neq y$.

The Kolmogorov backward and forward equations are

1. $\frac{d}{dt}P_t(x, y) = -[\alpha(x) + \beta(x)]P_t(x, y) + \alpha(x)P_t(x+1, y) + \beta(x)P_t(x-1, y)$ for $(x, y) \in S^2$.
2. $\frac{d}{dt}P_t(x, y) = -[\alpha(y) + \beta(y)]P_t(x, y) + \alpha(y-1)P_t(x, y-1) + \beta(y+1)P_t(x, y+1)$ for $(x, y) \in S^2$

Proof

These results follow from the generator matrix G above.

1. The backward equation is $\frac{d}{dt}P_t = GP_t$.
2. The forward equation is $\frac{d}{dt}P_t = P_tG$.

Limiting Behavior and Stationary Distributions

For our discussion of limiting behavior, we will consider first the important special case of a continuous-time birth-death chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \mathbb{N}$ and with $\alpha(x) > 0$ for all $x \in \mathbb{N}$ and $\beta(x) > 0$ for all $x \in \mathbb{N}_+$. For the jump chain $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$, recall that

$$p(x) = Q(x, x+1) = \frac{\alpha(x)}{\alpha(x) + \beta(x)}, \quad q(x) = Q(x, x-1) = \frac{\beta(x)}{\alpha(x) + \beta(x)}, \quad x \in \mathbb{N} \quad (16.21.4)$$

The jump chain \mathbf{Y} is a discrete-time birth-death chain, and our notation here is consistent with the notation that we used in that section. Note that \mathbf{X} and \mathbf{Y} are irreducible. We first consider transience and recurrence.

The chains \mathbf{X} and \mathbf{Y} are recurrent if and only if

$$\sum_{x=0}^{\infty} \frac{\beta(1) \cdots \beta(x)}{\alpha(1) \cdots \alpha(x)} = \infty \quad (16.21.5)$$

Proof

Recall that \mathbf{X} is recurrent if and only if \mathbf{Y} is recurrent. In our study of discrete-time birth-death chains we saw that \mathbf{Y} is recurrent if and only if

$$\sum_{x=0}^{\infty} \frac{q(1) \cdots q(x)}{p(1) \cdots p(x)} = \infty \quad (16.21.6)$$

But trivially,

$$\frac{q(1) \cdots q(x)}{p(1) \cdots p(x)} = \frac{\beta(1) \cdots \beta(x)}{\alpha(1) \cdots \alpha(x)} \quad (16.21.7)$$

Next we consider positive recurrence and invariant distributions. It's nice to look at this from different points of view.

The function $g : \mathbb{N} \rightarrow (0, \infty)$ defined by

$$g(x) = \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)}, \quad x \in \mathbb{N} \quad (16.21.8)$$

is invariant for \mathbf{X} , and is the only invariant function, up to multiplication by constants. Hence \mathbf{X} is positive recurrent if and only if $B = \sum_{x=0}^{\infty} g(x) < \infty$, in which case the (unique) invariant probability density function f is given by $f(x) = \frac{1}{B} g(x)$ for $x \in \mathbb{N}$. Moreover, $P_t(x, y) \rightarrow f(y)$ as $t \rightarrow \infty$ for every $x, y \in \mathbb{N}$.

Proof using the jump chain

From our study of discrete-time birth-death chains, we know that the function $h : \mathbb{N} \rightarrow (0, \infty)$ defined by

$$h(x) = \frac{p(0) \cdots p(x-1)}{q(1) \cdots q(x)}, \quad x \in \mathbb{N} \quad (16.21.9)$$

is invariant for \mathbf{Y} , and is the only positive invariant function up to multiplication by positive constants. It then follows from our study of invariant functions for continuous-time chains that the function h/λ is invariant for \mathbf{X} , and again is the only positive invariant function up to multiplication by positive constants. But it's simple to see that

$$\frac{h(x)}{\lambda(x)} = \frac{h(x)}{\alpha(x) + \beta(x)} = \frac{\alpha(1) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)} = \frac{1}{\alpha(0)} g(x) \quad (16.21.10)$$

where g is the function given in the theorem. The remaining parts of the theorem follow from the general theory.

Proof from the balance equation

A function $g : \mathbb{N} \rightarrow (0, \infty)$ is invariant for \mathbf{X} if and only if it satisfies the balance equation $gG = 0$. For our birth-death chain, this reduces to

$$\begin{aligned} \alpha(0)g(0) &= \beta(1)g(1) \\ [\alpha(x) + \beta(x)]g(x) &= \alpha(x-1)g(x-1) + \beta(x+1)g(x+1), \quad x \in \mathbb{N}_+ \end{aligned}$$

Substituting the equation with $x = 0$ on the left into the equation with $x = 1$ on the left gives $\alpha(1)g(1) = \beta(2)g(2)$. Substituting this into the equation with $x = 2$ on the left gives $\alpha(2)g(2) = \beta(3)g(3)$. In general, the balance equations imply

$$\alpha(x)g(x) = \beta(x+1)g(x+1), \quad x \in \mathbb{N} \quad (16.21.11)$$

Solving these new balance equations recursively gives

$$g(x) = g(0) \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)} \quad (16.21.12)$$

Letting $g(0) = 1$ gives the particular invariant function in the theorem. Again, the remaining parts follow from the general theory.

Here is a summary of the classification:

For the continuous-time birth-death chain \mathbf{X} , let

$$A = \sum_{x=0}^{\infty} \frac{\beta(1) \cdots \beta(x)}{\alpha(1) \cdots \alpha(x)}, \quad B = \sum_{x=0}^{\infty} \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)} \quad (16.21.13)$$

1. \mathbf{X} is transient if $A < \infty$.
2. \mathbf{X} is null recurrent if $A = \infty$ and $B = \infty$.
3. \mathbf{X} is positive recurrent if $B < \infty$.

Suppose now that $n \in \mathbb{N}_+$ and that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time birth-death chain on the integer interval $\mathbb{N}_n = \{0, 1, \dots, n\}$. We assume that $\alpha(x) > 0$ for $x \in \{0, 1, \dots, n-1\}$ while $\beta(x) > 0$ for $x \in \{1, 2, \dots, n\}$. Of course, we must have $\beta(0) = \alpha(n) = 0$. With these assumptions, \mathbf{X} is irreducible, and since the state space is finite, positive recurrent. So all that remains is to find the invariant distribution. The result is essentially the same as when the state space is \mathbb{N} .

The invariant probability density function f_n is given by

$$f_n(x) = \frac{1}{B_n} \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)} \text{ for } x \in \mathbb{N}_n \text{ where } B_n = \sum_{x=0}^n \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)} \quad (16.21.14)$$

Proof

Define

$$g_n(x) = \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)}, \quad x \in \mathbb{N}_n \quad (16.21.15)$$

The proof that g_n is invariant for \mathbf{X} is the same as [before](#). The constant B_n is the normalizing constant.

Note that $B_n \rightarrow B$ as $n \rightarrow \infty$, and if $B < \infty$, $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$ for $x \in \mathbb{N}$. We will see this type of behavior again. Results for the birth-death chain on \mathbb{N}_n often converge to the corresponding results for the birth-death chain on \mathbb{N} as $n \rightarrow \infty$.

Absorption

Often when the state space $S = \mathbb{N}$, the state of a birth-death chain represents a *population* of individuals of some sort (and so the terms *birth* and *death* have their usual meanings). In this case state 0 is absorbing and means that the population is extinct. Specifically, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular birth-death chain on \mathbb{N} with $\alpha(0) = \beta(0) = 0$ and with $\alpha(x), \beta(x) > 0$ for $x \in \mathbb{N}_+$. Thus, state 0 is absorbing and all positive states lead to each other and to 0. Let $T = \min\{t \in [0, \infty) : X_t = 0\}$ denote the time until absorption, where as usual, $\min \emptyset = \infty$. Many of the results concerning extinction of the continuous-time birth-death chain follow easily from corresponding results for the discrete-time birth-death jump chain.

One of the following events will occur:

1. *Population extinction*: $T < \infty$ or equivalently, $X_s = 0$ for some $s \in [0, \infty)$ and hence $X_t = 0$ for all $t \in [s, \infty)$.
2. *Population explosion*: $T = \infty$ or equivalently $X_t \rightarrow \infty$ as $t \rightarrow \infty$.

Proof

Part (b) follows from the general theory, since 0 is absorbing, and all positive states lead to each other and to 0. Thus the positive states are transient and we know that with probability 1, the jump chain will visit a transient state only finitely often. Thus $T = \infty$ is equivalent to $X_t \rightarrow \infty$ as $t \rightarrow \infty$. Without the assumption that the chain is regular, population explosion could occur in finite time.

Naturally we would like to find the probability of these complementary events, and happily we have already done so in our study of discrete-time birth-death chains. The absorption probability function v is defined by

$$v(x) = \mathbb{P}(T < \infty) = \mathbb{P}(X_t = 0 \text{ for some } t \in [0, \infty) \mid X_0 = x), \quad x \in \mathbb{N} \quad (16.21.16)$$

As before, let

$$A = \sum_{i=0}^{\infty} \frac{\beta(1) \cdots \beta(i)}{\alpha(1) \cdots \alpha(i)} \quad (16.21.17)$$

1. If $A = \infty$ then $v(x) = 1$ for all $x \in \mathbb{N}$.
2. If $A < \infty$ then

$$v(x) = \frac{1}{A} \sum_{i=x}^{\infty} \frac{\beta(1) \cdots \beta(i)}{\alpha(1) \cdots \alpha(i)}, \quad x \in \mathbb{N} \quad (16.21.18)$$

Proof

The continuous-time chain is absorbed into 0 if and only if the discrete-time jump chain is absorbed into 0. So the result follows from the corresponding result for discrete-time birth-death chains. Recall again that $q(x)/p(x) = \beta(x)/\alpha(x)$ for $x \in \mathbb{N}_+$

The mean time to extinction is considered next, so let $m(x) = \mathbb{E}(T \mid X_0 = x)$ for $x \in \mathbb{N}$. Unlike the probability of extinction, computing the mean time to extinction cannot be easily reduced to the corresponding discrete-time computation. However, the

method of computation does extend.

The mean absorption function is given by

$$m(x) = \sum_{j=1}^x \sum_{k=j-1}^{\infty} \frac{\alpha(j) \cdots \alpha(k)}{\beta(j) \cdots \beta(k+1)}, \quad x \in \mathbb{N} \quad (16.21.19)$$

Probabilistic Proof

The time required to go from state $x \in \mathbb{N}_+$ to $x-1$ has the same distribution as the time required to go from state 1 to 0, except with parameters $\alpha(y), \beta(y)$ for $y \in \{x, x+1, \dots\}$ instead of parameters $\alpha(y), \beta(y)$ for $y \in \{1, 2, \dots\}$. So by the additivity of expected value, we just need to compute $m(1)$ as a function of the parameters. Starting in state 1, the chain will be absorbed in state 0 after a random number of returns to state 1 without absorption. Whenever the chain is in state 1, absorption occurs at the next transition with probability $q(1)$ so it follows that the number of times that the chain is in state 1 before absorption has the geometric distribution on \mathbb{N}_+ with success parameter $q(1)$. The mean of this distribution is $1/q(1) = [\alpha(1) + \beta(1)]/\beta(1)$. On the other hand, starting in state 1, time until the chain is in state 1 again (without absorption) has the same distribution as the return time to state 0, starting in state 0 for the irreducible birth-death chain on \mathbb{N} with birth and death rates α' and β' given by $\alpha'(x) = \alpha(x+1)$ for $x \in \mathbb{N}$ and $\beta'(x) = \beta(x+1)$ for $x \in \mathbb{N}_+$. Thus, let

$$\mu = \frac{1}{\alpha(1) + \beta(1)} \sum_{k=0}^{\infty} \frac{\alpha(1) \cdots \alpha(k)}{\beta(2) \cdots \beta(k+1)} \quad (16.21.20)$$

Then μ is the mean return time to state 0 for the chain \mathbf{X}' . Specifically, note that if $\mu = \infty$ then \mathbf{X}' is either transient or null recurrent. If $\mu < \infty$ then $1/\mu$ is the invariant PDF at 0. So, it follows that

$$m(1) = \frac{1}{q(1)} \mu = \sum_{k=0}^{\infty} \frac{\alpha(1) \cdots \alpha(k)}{\beta(1) \cdots \beta(k+1)} \quad (16.21.21)$$

By our argument above, the mean time to go from state x to $x-1$ is

$$\sum_{k=x-1}^{\infty} \frac{\alpha(x) \cdots \alpha(k)}{\beta(x) \cdots \beta(k+1)} \quad (16.21.22)$$

In particular, note that

$$m(1) = \sum_{k=0}^{\infty} \frac{\alpha(1) \cdots \alpha(k)}{\beta(1) \cdots \beta(k+1)} \quad (16.21.23)$$

If $m(1) = \infty$ then $m(x) = \infty$ for all $x \in S$. If $m(1) < \infty$ then $m(x) < \infty$ for all $x \in S$

Next we will consider a birth-death chain on a finite integer interval with both endpoints absorbing. Our interest is in the probability of absorption in one endpoint rather than the other, and in the mean time to absorption. Thus suppose that $n \in \mathbb{N}_+$ and that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time birth-death chain on $\mathbb{N}_n = \{0, 1, \dots, n\}$ with $\alpha(0) = \beta(0) = 0$, $\alpha(n) = \beta(n) = 0$, and $\alpha(x) > 0, \beta(x) > 0$ for $x \in \{1, 2, \dots, n-1\}$. So the endpoints 0 and n are absorbing, and all other states lead to each other and to the endpoints. Let $T = \inf\{t \in [0, \infty) : X_t \in \{0, n\}\}$, the time until absorption, and for $x \in S$ let $v_n(x) = \mathbb{P}(X_T = 0 \mid X_0 = x)$ and $m_n(x) = \mathbb{E}(T \mid X_0 = x)$. The definitions make sense since T is finite with probability 1.

The absorption probability function for state 0 is given by

$$v_n(x) = \frac{1}{A_n} \sum_{i=x}^{n-1} \frac{\beta(1) \cdots \beta(i)}{\alpha(1) \cdots \alpha(i)} \text{ for } x \in \mathbb{N}_n \text{ where } A_n = \sum_{i=0}^{n-1} \frac{\beta(1) \cdots \beta(i)}{\alpha(1) \cdots \alpha(i)} \quad (16.21.24)$$

Proof

The jump chain $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a discrete-time birth-death chain on \mathbb{N}_n with 0 and n absorbing. Also, \mathbf{X} is absorbed into 0 or n if and only if \mathbf{Y} is absorbed into 0 or n , respectively. So the result follows from the corresponding result for \mathbf{Y} , since $q(x)/p(x) = \beta(x)/\alpha(x)$ for $x \in \{1, 2, \dots, n-1\}$.

Note that $A_n \rightarrow A$ as $n \rightarrow \infty$ where A is the constant above for the absorption probability at 0 with the infinite state space \mathbb{N} . If $A < \infty$ then $v_n(x) \rightarrow v(x)$ as $n \rightarrow \infty$ for $x \in \mathbb{N}$.

Time Reversal

Essentially, every irreducible continuous-time birth-death chain is reversible.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a positive recurrent birth-death chain on an integer interval $S \subseteq \mathbb{Z}$ with birth rate function $\alpha : S \rightarrow [0, \infty)$ and death rate function $\beta : S \rightarrow \infty$. Assume that $\alpha(x) > 0$, except at the maximum value of S , if there is one, and similarly that $\beta(x) > 0$, except at the minimum value of X , if there is one. Then \mathbf{X} is reversible.

Proof

Note that \mathbf{X} is irreducible. As usual, let G denote the generator matrix. It's easy to see that under the assumptions, $G(x, y) = 0$ implies $G(y, x) = 0$ for $(x, y) \in S^2$, and that the Kolmogorov cycle condition is satisfied: For every $n \in \mathbb{N}_+$ and every sequence $(x_1, x_2, \dots, x_n) \in S^n$,

$$G(x_1, x_2) \cdots G(x_{n-1}, x_n) G(x_n, x_1) = G(x_1, x_n) G(x_n, x_{n-1}) \cdots G(x_2, x_1) \quad (16.21.25)$$

In the important special case of a birth-death chain on \mathbb{N} , we can verify the balance equations directly.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time birth-death chain on $S = \mathbb{N}$ and with birth rate $\alpha(x) > 0$ for all $x \in \mathbb{N}$ and death rate $\beta(x) > 0$ for all $x \in \mathbb{N}_+$. Then \mathbf{X} is reversible.

Proof

We just need to show that the balance equation for a reversible chain holds, and this was actually done in the result [above](#). As before, let $g : \mathbb{N} \rightarrow (0, \infty)$ be the function given by

$$g(x) = \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)}, \quad x \in \mathbb{N} \quad (16.21.26)$$

The only nontrivial case of the balance equation $g(x)G(x, y) = g(y)G(y, x)$ for $(x, y) \in S^2$ is

$$g(x)G(x, x+1) = g(x+1)G(x+1, x) = \frac{\alpha(0) \cdots \alpha(x)}{\beta(1) \cdots \beta(x)}, \quad x \in \mathbb{N} \quad (16.21.27)$$

It follows from the general theory that g is invariant for \mathbf{X} and that \mathbf{X} is reversible with respect to g . Since we actually know from our work above that g is the only positive invariant function, up to multiplication by positive constants, we can simply say that \mathbf{X} is reversible.

In the positive recurrent case, it follows that the birth-death chain is stochastically the same, forward or backward in time, if the chain has the invariant distribution.

Examples and Special Cases

Regular and Irregular Chains

Our first exercise gives two pure birth chains, each with an unbounded exponential parameter function. One is regular and one is irregular.

Consider the pure birth process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on \mathbb{N}_+ with birth rate function α .

1. If $\alpha(x) = x^2$ for $x \in \mathbb{N}_+$, then \mathbf{X} is not regular.
2. If $\alpha(x) = x$ for $x \in \mathbb{N}_+$, then \mathbf{X} is regular.

Proof

The jump chain \mathbf{Y} is deterministic, except for the initial state. Given $Y_0 = x \in \mathbb{N}_+$, we have $Y_n = n + x$. Hence

1. $\sum_{n=0}^{\infty} \frac{1}{\lambda(Y_n)} = \sum_{n=0}^{\infty} \frac{1}{(n+x)^2} < \infty$
2. $\sum_{n=0}^{\infty} \frac{1}{\lambda(Y_n)} = \sum_{n=0}^{\infty} \frac{1}{n+x} = \infty$

So the results follow from the general theory.

Constant Birth and Death Rates

Our next examples consider birth-death chains with constant birth and death rates, except perhaps at the endpoints. Note that such chains will be regular since the exponential parameter function λ is bounded.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is the birth-death chain on \mathbb{N} , with constant birth rate $\alpha \in (0, \infty)$ on \mathbb{N} and constant death rate $\beta \in (0, \infty)$ on \mathbb{N}_+ .

1. \mathbf{X} is transient if $\beta < \alpha$.
2. \mathbf{X} is null recurrent if $\beta = \alpha$.
3. \mathbf{X} is positive recurrent if $\beta > \alpha$. The invariant distribution is the geometric distribution on \mathbb{N} with parameter α/β

$$f(x) = \left(1 - \frac{\alpha}{\beta}\right) \left(\frac{\alpha}{\beta}\right)^x, \quad x \in \mathbb{N} \quad (16.21.28)$$

Proof

Note that \mathbf{X} is irreducible since the birth rate is positive on \mathbb{N} and the death rate is positive on \mathbb{N}_+ . The series in the results above are geometric series:

$$\frac{\beta(1) \cdots \beta(x)}{\alpha(1) \cdots \alpha(x)} = \left(\frac{\beta}{\alpha}\right)^x, \quad \frac{\alpha(0) \cdots \alpha(x-1)}{\beta(1) \cdots \beta(x)} = \left(\frac{\alpha}{\beta}\right)^x, \quad x \in \mathbb{N} \quad (16.21.29)$$

Next we consider the chain with 0 absorbing. As in the general discussion above, let v denote the function that gives the probability of absorption and m the function that gives the mean time to absorption.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is the birth-death chain in \mathbb{N} with constant birth rate $\alpha \in (0, \infty)$ on \mathbb{N}_+ , constant death rate $\beta \in (0, \infty)$ on \mathbb{N}_+ , and with 0 absorbing. Then

1. If $\beta \geq \alpha$ then $v(x) = 1$ for $x \in \mathbb{N}$. If $\beta < \alpha$ then $v(x) = (\beta/\alpha)^x$ for $x \in \mathbb{N}$.
2. If $\alpha \geq \beta$ then $m(x) = \infty$. If $\alpha < \beta$ then $m(x) = x/(\beta - \alpha)$ for $x \in \mathbb{N}$.

Next let's look at chains on a finite state space. Let $n \in \mathbb{N}_+$ and define $\mathbb{N}_n = \{0, 1, \dots, n\}$.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time birth-death chain on \mathbb{N}_n with constant birth rate $\alpha \in (0, \infty)$ on $\{0, 1, \dots, n-1\}$ and constant death rate $\beta \in (0, \infty)$ on $\{1, 2, \dots, n\}$. The invariant probability density function f_n is given as follows:

1. If $\alpha \neq \beta$ then

$$f_n(x) = \frac{(\alpha/\beta)^x (1 - \alpha/\beta)}{1 - (\alpha/\beta)^{n+1}}, \quad x \in \mathbb{N}_n \quad (16.21.30)$$

2. If $\alpha = \beta$ then $f_n(x) = 1/(n+1)$ for $x \in \mathbb{N}_n$

Note that when $\alpha = \beta$, the invariant distribution is uniform on \mathbb{N}_n . Our final exercise considers the absorption probability at 0 when both endpoints are absorbing. Let v_n denote the function that gives the probability of absorption into 0, rather than n .

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is the birth-death chain on \mathbb{N}_n with constant birth rate α and constant death rate β on $\{1, 2, \dots, n-1\}$, and with 0 and n absorbing.

1. If $\alpha \neq \beta$ then

$$v_n(x) = \frac{(\beta/\alpha)^x - (\beta/\alpha)^n}{1 - (\beta/\alpha)^n}, \quad x \in \mathbb{N}_n \quad (16.21.31)$$

2. If $\alpha = \beta$ then $v_n(x) = (n-x)/n$ for $x \in \mathbb{N}_n$.

Linear Birth and Death Rates

For our next discussion, consider individuals that act identically and independently. Each individual splits into two at exponential rate $a \in (0, \infty)$ and dies at exponential rate $b \in (0, \infty)$.

Let X_t denote the population at time $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular, continuous-time birth-death chain with birth and death rate functions given by $\alpha(x) = ax$ and $\beta(x) = bx$ for $x \in \mathbb{N}$.

Proof

The fact that \mathbf{X} is a continuous-time Markov chain follows from the assumptions. Moreover, since the individuals act independently, the overall birth and death rates when the population is $x \in \mathbb{N}$ is simply x times the individual birth and death rates. The chain is regular since

$$\sum_{x=1}^{\infty} \frac{1}{(a+b)x} = \infty \quad (16.21.32)$$

Note that 0 is absorbing since the population is extinct, so as usual, our interest is in the probability of absorption and the mean time to absorption as functions of the initial state. The probability of absorption is the same as for the chain with [constant birth and death rates](#) discussed above.

The absorption probability function v is given as follows:

1. $v(x) = 1$ for all $x \in \mathbb{N}$ if $b \geq a$.
2. $v(x) = (b/a)^x$ for $x \in \mathbb{N}$ if $b < a$.

Proof

These results follow from the general results [above](#) since $\beta(x)/\alpha(x) = b/a$ for $x \in \mathbb{N}_+$. Hence for $x \in \mathbb{N}$,

$$\sum_{i=x}^{\infty} (b/a)^i = \begin{cases} \infty & b \geq a \\ \frac{(b/a)^x}{1-b/a} & b < a \end{cases} \quad (16.21.33)$$

The mean time to absorption is more interesting.

The mean time to absorption function m is given as follows:

1. If $a \geq b$ then $m(x) = \infty$ for $x \in \mathbb{N}_+$.
2. If $a < b$ then

$$m(x) = \sum_{j=1}^x \frac{b^{j-1}}{a^j} \int_0^{a/b} \frac{u^{j-1}}{1-u} du, \quad x \in \mathbb{N} \quad (16.21.34)$$

Proof

1. From the general results [above](#), note that

$$m(1) = \sum_{k=0}^{\infty} \frac{1}{(k+1)b} \left(\frac{a}{b}\right)^k \quad (16.21.35)$$

The sum is infinite if $a \geq b$.

2. If $a < b$ then again from the general formula above,

$$m(x) = \sum_{j=1}^x \sum_{k=j-1}^{\infty} \frac{1}{(k+1)b} \left(\frac{a}{b}\right)^{k-j+1} = \sum_{j=1}^x \frac{1}{b} \left(\frac{b}{a}\right)^j \sum_{k=j-1}^{\infty} \frac{1}{k+1} \left(\frac{a}{b}\right)^{k+1} \quad (16.21.36)$$

The inner series converges absolutely. Moreover, for $k \in \mathbb{N}$,

$$\frac{1}{k+1} \left(\frac{a}{b}\right)^{k+1} = \int_0^{a/b} u^k du \quad (16.21.37)$$

Substituting and interchanging the sum and integral gives

$$m(x) = \sum_{j=1}^x \frac{b^{j-1}}{a^j} \int_0^{a/b} \left(\sum_{k=j-1}^{\infty} u^k \right) du = \sum_{j=1}^x \frac{b^{j-1}}{a^j} \int_0^{a/b} \frac{u^{j-1}}{1-u} du \quad (16.21.38)$$

For small values of $x \in \mathbb{N}$, the integrals in the case $a < b$ can be done by elementary methods. For example,

$$\begin{aligned} m(1) &= -\frac{1}{a} \ln\left(1 - \frac{a}{b}\right) \\ m(2) &= m(1) - \frac{1}{a} - \frac{b}{a^2} \ln\left(1 - \frac{a}{b}\right) \\ m(3) &= m(2) - \frac{1}{2a} - \frac{b}{a^2} - \frac{b^2}{a^3} \ln\left(1 - \frac{a}{b}\right) \end{aligned}$$

However, a general formula requires the introduction of a special function that is not much more helpful than the integrals themselves. The Markov chain \mathbf{X} is actually an example of a branching chain. We will revisit this chain in that section.

Linear Birth and Death with Immigration

We continue our previous discussion but generalizing a bit. Suppose again that we have individuals that act identically and independently. An individual splits into two at exponential rate $a \in [0, \infty)$ and dies at exponential rate $b \in [0, \infty)$. Additionally, new individuals enter the population at exponential rate $c \in [0, \infty)$. This is the *immigration effect*, and when $c = 0$ we have the birth-death chain in the previous discussion.

Let X_t denote the population at time $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a regular, continuous-time birth-death chain with birth and death rate functions given by $\alpha(x) = ax + c$ and $\beta(x) = bx$ for $x \in \mathbb{N}$.

Proof

The fact that \mathbf{X} is a continuous-time Markov chain follows from the assumptions. Moreover, since the individuals act independently, the overall birth rate when the population is $x \in \mathbb{N}$ is $ax + c$ while the death rate is bx . The chain is regular since

$$\sum_{x=1}^{\infty} \frac{1}{((a+b)x + c)} = \infty \quad (16.21.39)$$

The infinitesimal matrix G is given as follows, for $x \in \mathbb{N}$:

1. $G(x, x) = -[(a+b)x + c]$
2. $G(x, x+1) = ax + c$
3. $G(x, x-1) = bx$

The backward and forward equations are given as follows, for $(x, y) \in \mathbb{N}^2$ and $t \in (0, \infty)$

1. $\frac{d}{dt} P_t(x, y) = -[(a+b)x + c]P_t(x, y) + (ax + c)P_t(x+1, y) + bxP_t(x-1, y)$
2. $\frac{d}{dt} P_t(x, y) = -[(a+b)y + c]P_t(x, y) + [a(y-1) + c]P_t(x, y-1) + b(y+1)P_t(x, y+1)$

We can use the forward equation to find the expected population size. Let $M_t(x) = \mathbb{E}(X_t, | X_0 = x)$ for $t \in [0, \infty)$ and $x \in \mathbb{N}$.

For $t \in [0, \infty)$ and $x \in \mathbb{N}$, the mean population size $M_t(x)$ is given as follows:

1. If $a = b$ then $M_t(x) = ct + x$.
2. If $a \neq b$ then

$$M_t(x) = \frac{c}{a-b} \left[e^{(a-b)t} - 1 \right] + x e^{(a-b)t} \quad (16.21.40)$$

Proof

First note that $M_t(x) = \sum_{y=0}^{\infty} y P_t(x, y)$ for $x \in \mathbb{N}$. Multiplying the forward equation above by y and summing over $y \in \mathbb{N}$ gives

$$\begin{aligned} \sum_{y=0}^{\infty} y \frac{d}{dt} P_t(x, y) &= a \sum_{y=2}^{\infty} y(y-1) P_t(x, y-1) + c \sum_{y=1}^{\infty} y P_t(x, y-1) \\ &\quad - (a+b) \sum_{y=0}^{\infty} y^2 P_t(x, y) - c \sum_{y=0}^{\infty} y P_t(x, y) + b \sum_{y=0}^{\infty} y(y+1) P_t(x, y+1) \end{aligned}$$

Re-indexing the sums and using some algebra gives the first-order differential equation

$$\frac{d}{dt} M_t(x) = c + (a-b)M_t(x), \quad x \in \mathbb{N}, t \in (0, \infty) \quad (16.21.41)$$

with initial condition $M_0(x) = x$. Solving the differential equation gives the result.

Note that $b > a$, so that the individual death rate exceeds the birth rate, then $M_t(x) \rightarrow c/(b-a)$ as $t \rightarrow \infty$ for $x \in \mathbb{N}$. If $a \geq b$ so that the birth rate equals or exceeds the death rate, then $M_t(x) \rightarrow \infty$ as $t \rightarrow \infty$ for $x \in \mathbb{N}_+$.

Next we will consider the special case with no births, but only death and immigration. In this case, the invariant distribution is easy to compute, and is one of our favorites.

Suppose that $a = 0$ and that $b, c > 0$. Then \mathbf{X} is positive recurrent. The invariant distribution is Poisson with parameter c/b :

$$f(x) = e^{-c/b} \frac{(c/b)^x}{x!}, \quad x \in \mathbb{N} \quad (16.21.42)$$

Proof

In terms of the general theory [above](#), note that the invariant function g , unique up to multiplication by positive constants, is given by

$$g(x) = \frac{\alpha(0) \cdots \alpha(x)}{\beta(1) \cdots \beta(x)} = \frac{c^x}{b^x x!} = \frac{(c/b)^x}{x!}, \quad x \in \mathbb{N} \quad (16.21.43)$$

Hence $B = \sum_{x=0}^{\infty} g(x) = e^{c/b} < \infty$ and therefore the chain is positive recurrent with invariant PDF

$$f(x) = \frac{1}{B} g(x) = e^{-c/b} \frac{(c/b)^x}{x!}, \quad x \in \mathbb{N} \quad (16.21.44)$$

This is the PDF of the Poisson distribution with parameter c/b .

The Logistics Chain

Consider a population that fluctuates between a minimum value $m \in \mathbb{N}_+$ and a maximum value $n \in \mathbb{N}_+$, where of course, $m < n$. Given the population size, the individuals act independently and identically. Specifically, if the population is $x \in \{m, m+1, \dots, n\}$ then an individual splits in two at exponential rate $a(n-x)$ and dies at exponential rate $b(x-m)$, where $a, b \in (0, \infty)$. Thus, an individual's birth rate decreases linearly with the population size from $a(n-m)$ to 0 while the death rate increases linearly with the population size from 0 to $b(n-m)$. These assumptions lead to the following definition.

The continuous-time birth-death chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on $S = \{m, m+1, \dots, n\}$ with birth rate function α and death rate function β given by

$$\alpha(x) = ax(n-x), \quad \beta(x) = bx(x-m), \quad x \in S \quad (16.21.45)$$

is the *logistic chain* on S with parameters a and b .

Justification

Since the individuals act independently and identically, the overall birth rate and death rates when the population is $x \in S$ is simply x times the birth and death rate for an individual.

Note that the logistics chain is a stochastic counterpart of the *logistics differential equation*, which typically has the form

$$\frac{dx}{dt} = c(x-m)(n-x) \quad (16.21.46)$$

where $m, n, c \in (0, \infty)$ and $m < n$. Starting in $x(0) \in (m, n)$, the solution remains in (m, n) for all $t \in [0, \infty)$. Of course, the logistics differential equation models a system that is continuous in time and space, whereas the logistics Markov chain models a system that is continuous in time and discrete in space.

For the logistics chain

1. The exponential parameter function λ is given by

$$\lambda(x) = ax(n-x) + bx(x-m), \quad x \in S \quad (16.21.47)$$

2. The transition matrix Q of the jump chain is given by

$$Q(x, x-1) = \frac{b(x-m)}{a(n-x) + b(x-m)}, \quad Q(x, x+1) = \frac{a(n-x)}{a(n-x) + b(x-m)}, \quad x \in S \quad (16.21.48)$$

In particular, m and n are reflecting boundary points, and so the chain is irreducible.

The generator matrix G for the logistics chain is given as follows, for $x \in S$:

1. $G(x, x) = -x[a(n-x) + b(x-m)]$
2. $G(x, x-1) = bx(x-m)$
3. $G(x, x+1) = ax(n-x)$

Since S is finite, \mathbf{X} is positive recurrent. The invariant distribution is given next.

Define $g: S \rightarrow (0, \infty)$ by

$$g(x) = \frac{1}{x} \binom{n-m}{x-m} \left(\frac{a}{b}\right)^{x-m}, \quad x \in S \quad (16.21.49)$$

Then g is invariant for \mathbf{X} .

Proof

Since we know that \mathbf{X} is reversible, we just need to show that $g(x)G(x, y) = g(y)G(y, x)$ for $(x, y) \in S^2$. For the logistics chain, the only non-trivial equation is $g(x)G(x, x+1) = g(x+1)G(x+1, x)$ for $x \in S$. Simple substitution and algebra show that both sides reduce to

$$\frac{(n-m)!}{(x-m)!(n-x-1)!} \frac{a^{x-m+1}}{b^{x-m}} \quad (16.21.50)$$

Of course it now follows that the invariant probability density function f for \mathbf{X} is given by $f(x) = g(x)/c$ for $x \in S$ where c is the normalizing constant

$$c = \sum_{x=m}^n \frac{1}{x} \binom{n-m}{x-m} \left(\frac{a}{b}\right)^{x-m} \quad (16.21.51)$$

The limiting distribution of \mathbf{X} has probability density function f .

Other Special Birth-Death Chains

There are a number of special birth-death chains that are studied in other sections, because the models are important and lead to special insights and analytic tools. These include

- Queuing chains
- The pure death branching chain
- The Yule process, a pure birth branching chain
- The general birth-death branching chain

This page titled [16.21: Continuous-Time Birth-Death Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.22: Continuous-Time Queuing Chains

Basic Theory

Introduction

In a *queuing model*, customers arrive at a station for service. As always, the terms are generic; here are some typical examples:

- The customers are persons and the service station is a store.
- The customers are file requests and the service station is a web server.



Figure 16.22.1: Ten customers and a server

Queuing models can be quite complex, depending on such factors as the probability distribution that governs the arrival of customers, the probability distribution that governs the service of customers, the number of servers, and the behavior of the customers when all servers are busy. Indeed, queuing theory has its own lexicon to indicate some of these factors. In this section, we will discuss a few of the basic, continuous-time queuing chains. In a general sense, the main interest in any queuing model is the number of customers in the system as a function of time, and in particular, whether the servers can adequately handle the flow of customers. This section parallels the section on discrete-time queuing chains.

Our main assumptions are as follows:

1. There are $k \in \mathbb{N}_+ \cup \{\infty\}$ servers.
2. The customers arrive according to a Poisson process with rate $\mu \in (0, \infty)$.
3. If all of the servers are busy, a new customer goes to the end of a single line of customers waiting service.
4. The time required to service a customer has an exponential distribution with parameter $\nu \in (0, \infty)$.
5. The service times are independent from customer to customer, and are independent of the arrival process.

Assumption (b) means that the times between arrivals of customers are independent and exponentially distributed, with parameter μ . Assumption (c) means that we have a *first-in, first-out* model, often abbreviated *FIFO*. Note that there are three parameters in the model: the number of servers k , the exponential parameter μ that governs the arrivals, and the exponential parameter ν that governs the service times. The special cases $k = 1$ (a single server) and $k = \infty$ (infinitely many servers) deserve special attention. As you might guess, the assumptions lead to a continuous-time Markov chain.

Let X_t denote the number of customers in the system (waiting in line or being served) at time $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on \mathbb{N} , known as the *M/M/k queuing chain*.

In terms of the basic structure of the chain, the important quantities are the exponential parameters for the states and the transition matrix for the embedded jump chain.

For the M/M/k chain \mathbf{X} ,

1. The exponential parameter function λ is given by $\lambda(x) = \mu + \nu x$ if $x \in \mathbb{N}$ and $x < k$ and $\lambda(x) = \mu + \nu k$ if $x \in \mathbb{N}$ and $x \geq k$.
2. The transition matrix Q for the jump chain is given by

$$Q(x, x-1) = \frac{\nu x}{\mu + \nu x}, \quad Q(x, x+1) = \frac{\mu}{\mu + \nu x}, \quad x \in \mathbb{N}, x < k$$

$$Q(x, x-1) = \frac{\nu k}{\mu + \nu k}, \quad Q(x, x+1) = \frac{\mu}{\mu + \nu k}, \quad x \in \mathbb{N}, x \geq k$$

So the M/M/k chain is a birth-death chain with 0 as a reflecting boundary point. That is, in state $x \in \mathbb{N}_+$, the next state is either $x-1$ or $x+1$, while in state 0, the next state is 1. When $k = 1$, the single-server queue, the exponential parameter in state $x \in \mathbb{N}_+$ is $\mu + \nu$ and the transition probabilities for the jump chain are

$$Q(x, x-1) = \frac{\nu}{\mu + \nu}, \quad Q(x, x+1) = \frac{\mu}{\mu + \nu} \quad (16.22.1)$$

When $k = \infty$, the infinite server queue, the cases above for $x \geq k$ are vacuous, so the exponential parameter in state $x \in \mathbb{N}$ is $\mu + \nu x$ and the transition probabilities are

$$Q(x, x-1) = \frac{\nu x}{\mu + \nu x}, \quad Q(x, x+1) = \frac{\mu}{\mu + \nu x} \quad (16.22.2)$$

Infinitesimal Generator

The infinitesimal generator of the chain gives the same information as the exponential parameter function and the jump transition matrix, but in a more compact form.

For the M/M/k queuing chain \mathbf{X} , the infinitesimal generator G is given by

$$\begin{aligned} G(x, x) &= -(\mu + \nu x), \quad G(x, x-1) = \nu x, \quad G(x, x+1) = \mu; & x \in \mathbb{N}, \quad x < k \\ G(x, x) &= -(\mu + \nu k), \quad G(x, x-1) = \nu k, \quad G(x, x+1) = \mu; & x \in \mathbb{N}, \quad x \geq k \end{aligned}$$

So for $k=1$, the single server queue, the generator G is given by $G(0,0) = -\mu$, $G(0,1) = \mu$, while for $x \in \mathbb{N}_+$, $G(x,x) = -(\mu + \nu)$, $G(x,x-1) = \nu$, $G(x,x+1) = \mu$. For $k = \infty$, the infinite server case, the generator G is given by $G(x,x) = -(\mu + \nu x)$, $G(x,x-1) = \nu x$, and $G(x,x+1) = \mu$ for all $x \in \mathbb{N}$.

Classification and Limiting Behavior

Again, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ denote the M/M/k queuing chain with arrival rate μ , service rate ν and with $k \in \mathbb{N}_+ \cup \{\infty\}$ servers. As noted in the introduction, of fundamental importance is the question of whether the servers can handle the flow of customers, so that the queue eventually empties, or whether the length of the queue grows without bound. To understand the limiting behavior, we need to classify the chain as transient, null recurrent, or positive recurrent, and find the invariant functions. This will be easy to do using our results for more general continuous-time birth-death chains. Note first that \mathbf{X} is irreducible. It's best to consider the single server and infinite server cases individually.

The single server queuing chain \mathbf{X} is

1. Transient if $\nu < \mu$.
2. Null recurrent if $\nu = \mu$.
3. Positive recurrent if $\nu > \mu$. The invariant distribution is the geometric distribution on \mathbb{N} with parameter μ/ν . The invariant probability density function f is given by

$$f(x) = \left(1 - \frac{\mu}{\nu}\right) \left(\frac{\mu}{\nu}\right)^x, \quad x \in \mathbb{N} \quad (16.22.3)$$

Proof

This follows directly from results for the continuous-time birth-death chain, with constant birth rate μ on \mathbb{N} and constant death rate ν on \mathbb{N}_+ .

The result makes intuitive sense. If the service rate is less than the arrival rate, the chain is transient and the length of the queue grows to infinity. If the service rate is greater than the arrival rate, the chain is positive recurrent. At the boundary between these two cases, when the arrival and service rates are the same, the chain is null recurrent.

The infinite server queuing chain \mathbf{X} is positive recurrent. The invariant distribution is the Poisson distribution with parameter μ/ν . The invariant probability density function f is given by

$$f(x) = e^{-\mu/\nu} \frac{(\mu/\nu)^x}{x!}, \quad x \in \mathbb{N} \quad (16.22.4)$$

Proof

This also follows from results for the continuous-time birth-death chain. In the notation of that section, the birth rate is constant, $\mu(x) = \mu$ for $x \in \mathbb{N}$ and the death rate is proportional to the number of customers in the system: $\nu(x) = \nu x$ for $x \in \mathbb{N}_+$. Hence the invariant function (unique up to multiplication by constants) is

$$x \mapsto \frac{\mu(0) \cdots \mu(x-1)}{\nu(1) \cdots \nu(x)} = \frac{\mu^x}{\nu^x x!} \quad (16.22.5)$$

Normalized, this is the Poisson distribution with parameter μ/ν .

This result also makes intuitive sense.

This page titled [16.22: Continuous-Time Queuing Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

16.23: Continuous-Time Branching Chains

Basic Theory

Introduction

Generically, suppose that we have a system of *particles* that can generate or split into other particles of the same type. Here are some typical examples:

- The particles are biological organisms that reproduce.
- The particles are neutrons in a chain reaction.
- The particles are electrons in an electron multiplier.

We assume that the lifetime of each particle is exponentially distributed with parameter $\alpha \in (0, \infty)$, and at the end of its life, is replaced by a random number of new particles that we will refer to as *children* of the original particle. The number of children N of a particle has probability density function f on \mathbb{N} . The particles act independently, so in addition to being identically distributed, the lifetimes and the number of children are independent from particle to particle. Finally, we assume that $f(1) = 0$, so that a particle cannot simply die and be replaced by a single new particle. Let μ and σ^2 denote the mean and variance of the number of offspring of a single particle. So

$$\mu = \mathbb{E}(N) = \sum_{n=0}^{\infty} n f(n), \quad \sigma^2 = \text{var}(N) = \sum_{n=0}^{\infty} (n - \mu)^2 f(n) \quad (16.23.1)$$

We assume that μ is finite and so σ^2 makes sense. In our study of discrete-time Markov chains, we studied branching chains in terms of *generational time*. Here we want to study the model in real time.

Let X_t denote the number of particles at time $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a continuous-time Markov chain on \mathbb{N} , known as a *branching chain*. The exponential parameter function λ and jump transition matrix Q are given by

1. $\lambda(x) = \alpha x$ for $x \in \mathbb{N}$
2. $Q(x, x+k-1) = f(k)$ for $x \in \mathbb{N}_+$ and $k \in \mathbb{N}$.

Proof

That \mathbf{X} is a continuous-time Markov chain follows from the assumptions and the basic structure of continuous-time Markov chains. It turns out that the assumption that $\mu < \infty$ implies that \mathbf{X} is regular, so that $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$, where τ_n is the time of the n th jump for $n \in \mathbb{N}_+$.

1. Starting with x particles, the time of the first state change is the minimum of x independent variables, each exponentially distributed with parameter α . As we know, this minimum is also exponentially distributed with parameter αx .
2. Starting in state $x \in \mathbb{N}_+$, the next state will be $x+k-1$ for $k \in \mathbb{N}$, if the particle dies and leaves k children in her place. This happens with probability $f(k)$.

Of course 0 is an absorbing state, since this state means extinction with no particles. (Note that $\lambda(0) = 0$ and so by default, $Q(0, 0) = 1$.) So with a branching chain, there are essentially two types of behavior: population extinction or population explosion.

For the branching chain $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ one of the following events occurs with probability 1:

1. Extinction: $X_t = 0$ for some $t \in [0, \infty)$ and hence $X_s = 0$ for all $s \geq t$.
2. Explosion: $X_t \rightarrow \infty$ as $t \rightarrow \infty$.

Proof

If $f(0) > 0$ then all states lead to the absorbing state 0 and hence the set of positive states \mathbb{N}_+ is transient. With probability 1, the jump chain \mathbf{Y} visits a transient state only finitely many times, so with probability 1 either $Y_n = 0$ for some $n \in \mathbb{N}$ or $Y_n \rightarrow \infty$ as $n \rightarrow \infty$. If $f(0) = 0$ then Y_n is strictly increasing in n , since $f(1) = 0$ by assumption. Hence with probability 1, $Y_n \rightarrow \infty$ as $n \rightarrow \infty$.

Without the assumption that $\mu < \infty$, explosion can actually occur in finite time. On the other hand, the assumption that $f(1) = 0$ is for convenience. Without this assumption, \mathbf{X} would still be a continuous-time Markov chain, but as discussed in the Introduction, the exponential parameter function would be $\lambda(x) = \alpha f(1)x$ for $x \in \mathbb{N}$ and the jump transition matrix would be

$$Q(x, x+k-1) = \frac{f(k)}{1-f(1)}, \quad x \in \mathbb{N}_+, k \in \{0, 2, 3, \dots\} \quad (16.23.2)$$

Because all particles act identically and independently, the branching chain starting with $x \in \mathbb{N}_+$ particles is essentially x independent copies of the branching chain starting with 1 particle. In many ways, this is the fundamental insight into branching chains, and in particular, means that we can often condition on $X(0) = 1$.

Generator and Transition Matrices

As usual, we will let $\mathbf{P} = \{P_t : t \in [0, \infty)\}$ denote the semigroup of transition matrices of \mathbf{X} , so that $P_t(x, y) = \mathbb{P}(X_t = y \mid X = x)$ for $(x, y) \in \mathbb{N}^2$. Similarly, G denotes the infinitesimal generator matrix of \mathbf{X} .

The infinitesimal generator G is given by

$$\begin{aligned} G(x, x) &= -\alpha x, & x \in \mathbb{N} \\ G(x, x+k-1) &= \alpha x f(k), & x \in \mathbb{N}_+, k \in \mathbb{N} \end{aligned}$$

Proof

This follows immediately from the exponential parameter function and the jump transition matrix [above](#).

The Kolmogorov backward equation is

$$\frac{d}{dt} P_t(x, y) = -\alpha x P_t(x, x) + \alpha x \sum_{k=0}^{\infty} f(k) P_t(x+k-1, y), \quad (x, y) \in \mathbb{N}^2 \quad (16.23.3)$$

Proof

The backward equation is $\frac{d}{dt} P_t = G P_t$, so the result follows from the previous theorem.

Unlike some of our other continuous-time models, the jump chain \mathbf{Y} governed by Q is not the discrete-time version of the model. That is, \mathbf{Y} is not a discrete-time branching chain, since in discrete time, the index n represents the n th generation, whereas here it represents the n th time that a particle reproduces. However, there are lots of discrete-time branching chain embedded in the continuous-time chain.

Fix $t \in (0, \infty)$ and define $\mathbf{Z}_t = \{X_{nt} : n \in \mathbb{N}\}$. Then \mathbf{Z}_t is a discrete-time branching chain with offspring probability density function f_t given by $f_t(x) = P_t(1, x)$ for $x \in \mathbb{N}$.

Proof

In general, we know that sampling a (homogeneous) continuous-time Markov chain at multiples of a fixed $t \in (0, \infty)$, results in a (homogeneous) discrete-time Markov chain. For \mathbf{Z}_t to be a branching chain, we just need to note that

$$P_t(x, y) = f_t^{*x}(y), \quad (x, y) \in \mathbb{N}^2 \quad (16.23.4)$$

where f_t^{*x} is the convolution power of f_t of order x . This is a consequence of the fundamental fact that X_t given $X_0 = x$ has the same distribution as the sum of x independent copies of X_t given $X_0 = 1$. Recall that the PDF of a sum of independent variables is the convolution of the individual PDFs.

Probability Generating Functions

As in the discrete case, probability generating functions are an important analytic tool for continuous-time branching chains.

For $t \in [0, \infty)$ let Φ_t denote the probability generating function of X_t given $X_0 = 1$

$$\Phi_t(r) = \mathbb{E}(r^{X_t} \mid X_0 = 1) = \sum_{x=0}^{\infty} r^x P_t(1, x) \quad (16.23.5)$$

Let Ψ denote the probability generating function of N

$$\Psi(r) = \mathbb{E}(r^N) = \sum_{n=0}^{\infty} r^n f(n) \quad (16.23.6)$$

The generating functions are defined (the series are absolutely convergent) at least for $r \in (-1, 1]$.

The collection of generating functions $\Phi = \{\Phi_t : t \in [0, \infty)\}$ gives the same information as the collection of probability density functions $\{P_t(1, \cdot) : t \in [0, \infty)\}$. With the fundamental insight that the branching process starting with one particle determines the branching process in general, Φ actually determines the transition semigroup $\mathbf{P} = \{P_t : t \in [0, \infty)\}$.

For $t \in [0, \infty)$ and $x \in \mathbb{N}$, the probability generating function of X_t given $X_0 = x$ is Φ_t^x :

$$\sum_{y=0}^{\infty} r^y P_t(x, y) = [\Phi_t(r)]^x \quad (16.23.7)$$

Proof

Again, given $X_0 = x$, the number of particles X_t at time t has the same distribution as the sum of x independent copies of X_t given $X_0 = 1$. Recall that the PGF of a sum of independent variables is the product of the PGFs of the variables.

Note that Φ_t is the generating function of the offspring distribution for the embedded discrete-time branching chain $\mathbf{Z}_t = \{X_{nt} : n \in \mathbb{N}\}$ for $t \in (0, \infty)$. On the other hand, Ψ is the generating function of the offspring distribution for the continuous-time chain. So our main goal in this discussion is to see how Φ is built from Ψ . Because \mathbf{P} is a semigroup under matrix multiplication, and because the particles act identically and independently, Φ is a semigroup under composition.

$\Phi_{s+t} = \Phi_s \circ \Phi_t$ for $s, t \in [0, \infty)$.

Proof

Using the semigroup property (the Chapman-Kolmogorov equations) and the previous result we have

$$\begin{aligned} \Phi_{s+t}(r) &= \sum_{y=0}^{\infty} r^y P_{s+t}(1, y) = \sum_{y=0}^{\infty} r^y \sum_{x=0}^{\infty} P_s(1, x) P_t(x, y) = \sum_{x=0}^{\infty} P_s(1, x) \sum_{y=0}^{\infty} r^y P_t(x, y) \\ &= \sum_{x=0}^{\infty} P_s(1, x) [\Phi_t(r)]^x = \Phi_s[\Phi_t(r)] \end{aligned}$$

Note also that $\Phi_0(r) = \mathbb{E}(r^{X_0} | X_0 = 1) = r$ for all $r \in \mathbb{R}$. This also follows from the semigroup property: $\Phi_0 = \Phi_0 \circ \Phi_0$. The fundamental relationship between the collection of generating functions Φ and the generating function Ψ is given in the following theorem:

The mapping $t \mapsto \Phi_t$ satisfies the differential equation

$$\frac{d}{dt} \Phi_t = \alpha(\Psi \circ \Phi_t - \Phi_t) \quad (16.23.8)$$

Proof

Using the Kolmogorov backward equation we have

$$\frac{d}{dt} \Phi_t(r) = \sum_{x=0}^{\infty} r^x \frac{d}{dt} P_t(1, x) = \sum_{x=0}^{\infty} r^x G P_t(1, x) \quad (16.23.9)$$

Using the [generator above](#),

$$G P_t(1, x) = \sum_{y=0}^{\infty} G(1, y) P_t(y, x) = -\alpha P_t(1, x) + \sum_{k=0}^{\infty} \alpha f(k) P_t(k, x), \quad x \in \mathbb{N} \quad (16.23.10)$$

Substituting and using the result [above](#) gives

$$\begin{aligned} \frac{d}{dt} \Phi_t(r) &= \sum_{x=0}^{\infty} r^x \left[-\alpha P_t(1, x) + \sum_{k=0}^{\infty} \alpha f(k) P_t(k, x) \right] = -\alpha \sum_{x=0}^{\infty} r^x P_t(1, x) + \alpha \sum_{x=0}^{\infty} \sum_{k=0}^{\infty} r^x f(k) P_t(k, x) \\ &= -\alpha \Phi_t(r) + \alpha \sum_{k=0}^{\infty} f(k) \sum_{x=0}^{\infty} r^x P_t(k, x) = -\alpha \Psi_t(r) + \alpha \sum_{k=0}^{\infty} f(k) [\Phi_t(r)]^k = -\alpha \Phi_t(r) + \alpha \Psi[\Phi_t(r)] \end{aligned}$$

This differential equation, along with the initial condition $\Phi_0(r) = r$ for all $r \in \mathbb{R}$ determines the collection of generating functions Φ . In fact, an implicit solution for $\Phi_t(r)$ is given by the integral equation

$$\int_r^{\Phi_t(r)} \frac{1}{\Psi(u) - u} du = \alpha t \quad (16.23.11)$$

Another relationship is given in the following theorem. Here, Φ'_t refers to the derivative of the generating function Φ_t with respect to its argument, of course (so r , not t).

For $t \in [0, \infty)$,

$$\Phi'_t = \frac{\Psi \circ \Phi_t - \Phi_t}{\Psi - \Phi_0} \quad (16.23.12)$$

Proof

From the semigroup property, we have $\Phi_{t+s}(r) = \Phi_t[\Phi_s(r)]$ for $s, t \in [0, \infty)$. Differentiating with respect to s and using the chain rule along with the previous theorem gives

$$\frac{d}{ds} \Phi_{t+s}(r) = \Phi'_t[\Phi_s(r)] \frac{d}{ds} \Phi_s(r) = \Phi'_t[\Phi_s(r)] \alpha [\Psi(\Phi_s(r)) - \Phi_s(r)] \quad (16.23.13)$$

Evaluating at $s = 0$ and using the condition $\Phi_0(r) = r$ we have

$$\frac{d}{dt} \Phi_t(r) = \Phi'_t(r) \alpha [\Psi(r) - r] \quad (16.23.14)$$

Using the previous theorem once again gives

$$\alpha [\Psi(\Phi_t(r)) - \Phi_t(r)] = \Phi'_t(r) \alpha [\Psi(r) - r] \quad (16.23.15)$$

Solving for $\Phi'_t(r)$ gives the result.

Moments

In this discussion, we will study the mean and variance of the number of particles at time $t \in [0, \infty)$. Let

$$m_t = \mathbb{E}(X_t \mid X_0 = 1), \quad v_t = \text{var}(X_t \mid X_0 = 1), \quad t \in [0, \infty) \quad (16.23.16)$$

so that m_t and v_t are the mean and variance, starting with a single particle. As always with a branching process, it suffices to consider a single particle:

For $t \in [0, \infty)$ and $x \in \mathbb{N}$,

1. $\mathbb{E}(X_t \mid X_0 = x) = x m_t$
2. $\text{var}(X_t \mid X_0 = x) = x v_t$

Proof

Once again, the distribution of X_t given $X_0 = x$ is the same as the distribution of the sum of x independent copies of X_t given $X_0 = 1$. Recall that the mean of a sum of variables is the sum of the individual means, and the variance of the sum of *independent* variables is the sum of the individual variances.

Recall also that μ and σ^2 are the the mean and variance of the number of offspring of a particle. Here is the connection between the means:

$$m_t = e^{\alpha(\mu-1)t} \text{ for } t \in [0, \infty).$$

1. If $\mu < 1$ then $m_t \rightarrow 0$ as $t \rightarrow \infty$. This is *extinction in the mean*.
2. If $\mu > 1$ then $m_t \rightarrow \infty$ as $t \rightarrow \infty$. This is *explosion in the mean*.
3. If $\mu = 1$ then $m_t = 1$ for all $t \in [0, \infty)$. This is *stability in the mean*.

Proof

From the proof of the [previous theorem](#),

$$\frac{d}{dt} \Phi_t(r) = \alpha \Phi'_t(r) [\Psi(r) - r] \quad (16.23.17)$$

Differentiating with respect to r , interchanging the order of integration on the left, and using the product rule on the right gives

$$\frac{d}{dt}\Phi'_t(r) = \alpha\Phi''_t(r)[\Psi(r) - r] + \alpha\Phi'_t(r)[\Psi'(r) - 1] \quad (16.23.18)$$

Now let $r = 1$ and recall that $\Phi(1) = 1$. We get

$$\frac{d}{dt}\Phi'_t(1) = \alpha\Phi'_t(1)[\Psi'(1) - 1] \quad (16.23.19)$$

From the basic theory of probability generating functions, $m_t = \Phi'_t(1)$ and similarly, $\mu = \Psi'(1)$. Hence we have

$$\frac{d}{dt}m_t = \alpha(\mu - 1)m_t \quad (16.23.20)$$

Of course we have the initial condition $m_0 = 1$.

This result is intuitively very appealing. As a function of time, the expected number of particles either grows or decays exponentially, depending on whether the expected number of offspring of a particle is greater or less than one. The connection between the variances is more complicated. We assume that $\sigma^2 < \infty$.

If $\mu \neq 1$ then

$$v_t = \left[\frac{\sigma^2}{\mu - 1} + (\mu - 1) \right] \left[e^{2\alpha(\mu - 1)t} - e^{\alpha(\mu - 1)t} \right], \quad t \in [0, \infty) \quad (16.23.21)$$

If $\mu = 1$ then $v_t = \alpha\sigma^2 t$.

1. If $\mu < 1$ then $v_t \rightarrow 0$ as $t \rightarrow \infty$
2. If $\mu \geq 1$ then $v_t \rightarrow \infty$ as $t \rightarrow \infty$

Proof

Probability generating functions are naturally connected to factorial moments, so it's best to work with these. Thus, let $w_t = \mathbb{E}[X_t(X_t - 1) \mid X_0 = 1]$ for $t \in [0, \infty)$ and let $\delta = \mathbb{E}[N(N - 1)]$. These are the factorial moments of order 2. In the proof of the last theorem we showed that

$$\frac{d}{dt}\Phi'_t(r) = \alpha\Phi''_t(r)[\Psi(r) - r] + \alpha\Phi'_t(r)[\Psi'(r) - 1] \quad (16.23.22)$$

Differentiating with respect to r again gives

$$\frac{d}{dt}\Phi''_t(r) = \alpha\Phi'''_t(r)[\Psi(r) - r] + 2\alpha\Phi''_t(r)[\Psi'(r) - 1] + \alpha\Phi'_t(r)\Psi''(r) \quad (16.23.23)$$

Now substitute $r = 1$. Recall that $\Phi''_t(1) = w_t$, $\Phi'_t(1) = m_t = e^{\alpha(\mu - 1)t}$, $\Psi''(1) = \delta$, $\Psi'(1) = \mu$, and $\Psi(1) = 1$. We get the differential equation

$$\frac{d}{dt}w_t = 2\alpha(\mu - 1)w_t + \alpha\delta e^{\alpha(\mu - 1)t} \quad (16.23.24)$$

with the initial condition $w_0 = 0$.

Suppose that $\mu \neq 1$. Then using standard methods for a linear, first order differential equations with constant coefficients and an exponential forcing function, the solution is

$$w_t = \frac{\delta}{\mu - 1} \left[e^{2\alpha(\mu - 1)t} - e^{\alpha(\mu - 1)t} \right] \quad (16.23.25)$$

But $\delta = \sigma^2 + \mu^2 - \mu$, and similarly $w_t = v_t + m_t^2 - m_t$ with $m_t = e^{\alpha(\mu - 1)t}$. Substitution and some algebra then gives the result.

Suppose now that $\mu = 1$. Then also $m_t = 1$ for all $t \in [0, \infty)$ and so $\delta = \sigma^2$ and $v_t = w_t$. The differential equation above reduces simply to

$$\frac{d}{dt}v_t = \alpha\sigma^2 \quad (16.23.26)$$

with initial condition $v_0 = 0$ so trivially $v_t = \alpha\sigma^2 t$. Finally, in the context of part (b), note that if $\mu = 1$ we must have $\sigma^2 > 0$ since we have assumed that $f(1) = 0$.

If $\mu < 1$ so that $m_t \rightarrow 0$ as $t \rightarrow \infty$ and we have extinction in the mean, then $v_t \rightarrow 0$ as $t \rightarrow \infty$ also. If $\mu > 1$ so that $m_t \rightarrow \infty$ as $t \rightarrow \infty$ and we have explosion in the mean, then $v_t \rightarrow \infty$ as $t \rightarrow \infty$ also. We would expect these results. On the other hand, if $\mu = 1$ so that $m_t = 1$ for all $t \in [0, \infty)$ and we have stability in the mean, then v_t grows linearly in t . This gives some insight into what to expect next when we consider the *probability* of extinction.

The Probability of Extinction

As shown [above](#), there are two types of behavior for a branching process, either population extinction or population explosion. In this discussion, we study the extinction probability, starting as usual with a single particle:

$$q = \mathbb{P}(X_t = 0 \text{ for some } t \in (0, \infty) \mid X_0 = 1) = \lim_{t \rightarrow \infty} \mathbb{P}(X_t = 0 \mid X_0 = 1) \quad (16.23.27)$$

Need we say it? The extinction probability starting with an arbitrary number of particles is easy.

For $x \in \mathbb{N}$,

$$\mathbb{P}(X_t = 0 \text{ for some } t \in (0, \infty) \mid X_0 = x) = \lim_{t \rightarrow \infty} \mathbb{P}(X_t = 0 \mid X_0 = x) = q^x \quad (16.23.28)$$

Proof

Given $X_0 = x$, extinction has occurred by time t if and only if extinction has occurred by time t for each of the x independent branching chains formed from the descendents of the x initial particles.

We can easily relate extinction for the continuous-time branching chain \mathbf{X} to extinction for any of the embedded discrete-time branching chains.

If extinction occurs for \mathbf{X} then extinction occurs for \mathbf{Z}_t for every $t \in (0, \infty)$. Conversely if extinction occurs for \mathbf{Z}_t for some $t \in (0, \infty)$ then extinction occurs for \mathbf{Z}_t for every $t \in (0, \infty)$ and extinction occurs for \mathbf{X} . Hence q is the minimum solution in $(0, 1]$ of the equation $\Phi_t(r) = r$ for every $t \in (0, \infty)$.

Proof

The statements about the extinction event follow immediately from the fact that 0 is absorbing, so that if $X_t = 0$ for some $t \in (0, \infty)$ then $X_s = 0$ for every $s \in [t, \infty)$. The result for the extinction probability q follows from the theory of discrete-time branching chains.

So whether or not extinction is certain depends on the critical parameter μ .

The extinction probability q and the mean of the offspring distribution μ are related as follows:

1. If $\mu \leq 1$ then $q = 1$, so extinction is certain.
2. If $\mu > 1$ then $0 < q < 1$, so there is a positive probability of extinction and a positive probability of explosion.

Proof

These results follow from the corresponding results for discrete-time branching chains. Fix $t \in (0, \infty)$ and recall that m_t is the mean of the offspring distribution for the discrete-time chain $\mathbf{Z}_t = \{X_{nt} : n \in \mathbb{N}\}$. From the result [above](#),

1. If $\mu \leq 1$ then $m_t \leq 1$.
2. If $\mu > 1$ then $m_t > 1$.

It would be nice to have an equation for q in terms of the offspring probability generating function Ψ . This is also easy

The probability of extinction q is the minimum solution in $(0, 1]$ of the equation $\Psi(r) = r$.

Proof

From the result [above](#), $\Phi_t(q) = 1$ for every $t \in (0, \infty)$. Substituting $r = q$ in the [differential equation above](#), we have $\frac{d}{dt}\Phi_t(q) = 0$ and hence $\Psi(q) = q$. As in the theory of discrete branching chains, the equation $\Psi(r) = r$ has only the solution 1 in $(0, 1]$ if $\mu = \Psi'(1) \leq 1$ or there are two solutions $q \in (0, 1)$ and 1 if $\mu > 1$. In both cases, q is the smaller solution.

Special Models

We now turn our attention to a number of special branching chains that are important in applications or lead to interesting insights. We will use the notation established above, so that α is the parameter of the exponential lifetime of a particle, Q is the transition matrix of the

jump chain, G is the infinitesimal generator matrix, and P_t is the transition matrix at time $t \in [0, \infty)$. Similarly, $m_t = \mathbb{E}(X_t | X_0 = x)$, $v_t = \text{var}(X_t | X_0 = x)$, and Φ_t are the mean, variance, and generating function of the number of particles at time $t \in [0, \infty)$, starting with a single particle. As always, be sure to try these exercises yourself before looking at the proofs and solutions.

The Pure Death Branching Chain

First we consider the branching chain in which each particle simply dies without offspring. Sadly for these particles, extinction is inevitable, but this case is still a good place to start because the analysis is simple and lead to explicit formulas. Thus, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a branching process with lifetime parameter $\alpha \in (0, \infty)$ and offspring probability density function f with $f(0) = 1$.

The transition matrix of the jump chain and the generator matrix are given by

1. $Q(0, 0) = 1$ and $Q(x, x - 1) = 1$ for $x \in \mathbb{N}_+$
2. $G(x, x) = -\alpha x$ for $x \in \mathbb{N}$ and $G(x, x - 1) = \alpha x$ for $x \in \mathbb{N}_+$

The time-varying functions are more interesting.

Let $t \in [0, \infty)$. Then

1. $m_t = e^{-\alpha t}$
2. $v_t = e^{-\alpha t} - e^{-2\alpha t}$
3. $\Phi_t(r) = 1 - (1 - r)e^{-\alpha t}$ for $r \in \mathbb{R}$
4. Given $X_0 = x$ the distribution of X_t is binomial with trial parameter x and success parameter $e^{-\alpha t}$.

$$P_t(x, y) = \binom{x}{y} e^{-\alpha t y} (1 - e^{-\alpha t})^{x-y}, \quad x \in \mathbb{N}, y \in \{0, 1, \dots, x\} \quad (16.23.29)$$

Direct Proof

All of these results follow from the general methods above, with $\mu = \sigma = 0$ and $\Psi(r) = 1$ for $r \in \mathbb{R}$. But it's helpful to give direct proofs. Given $X_0 = 1$, let τ be the time until the first transition, which is simply the lifetime of the particle. So τ has the exponential distribution with parameter α . For $t \in [0, \infty)$, X_t is an indicator random variable (taking just values 0 and 1) with

$$\mathbb{P}(X_t = 1 | X_0 = 1) = \mathbb{P}(\tau > t | X_0 = 1) = e^{-\alpha t} \quad (16.23.30)$$

Part (a), (b), and (c) are standard results for an indicator variable. For part (d), given $X_0 = x$, each of the x particles, independently, is still alive at time t with probability $e^{-\alpha t}$. Hence the number of particles still alive has the binomial distribution with parameters x and $e^{-\alpha t}$.

In particular, note that $P_t(x, 0) = (1 - e^{-\alpha t})^x \rightarrow 1$ as $t \rightarrow \infty$. that is, the probability of extinction by time t increases to 1 exponentially fast. Since we have an explicit formula for the transition matrices, we can find an explicit formula for the potential matrices as well. The result uses the beta function B .

For $\beta \in (0, \infty)$ the potential matrix U_β is given by

$$U_\beta(x, y) = \frac{1}{\alpha} \binom{x}{y} B(y + \beta/\alpha, x - y + 1), \quad x \in \mathbb{N}, y \in \{0, 1, \dots, x\} \quad (16.23.31)$$

For $\beta = 0$, the potential matrix U is given by

1. $U(x, 0) = \infty$ for $x \in \mathbb{N}$
2. $U(x, y) = 1/\alpha y$ for $x, y \in \mathbb{N}_+$ and $x \leq y$.

Proof

Suppose that $\beta > 0$ and that $x, y \in \mathbb{N}$ with $x \leq y$. By definition

$$U_\beta(x, y) = \int_0^\infty e^{-\beta t} P_t(x, y) dt = \int_0^\infty e^{-\beta t} \binom{x}{y} e^{-\alpha t y} (1 - e^{-\alpha t})^{x-y} dt \quad (16.23.32)$$

Substitute $u = e^{-\alpha t}$ so that $du = -\alpha e^{-\alpha t} dt$ or equivalently $dt = -du/\alpha u$. After some algebra, the result is

$$U_\beta(x, y) = \frac{1}{\alpha} \binom{x}{y} \int_0^1 u^{y+\beta/\alpha-1} (1-u)^{x-y} du \quad (16.23.33)$$

By definition, the last integral is $B(y + \beta/\alpha, x - y + 1)$.

1. For $x \in \mathbb{N}$,

$$U(x, 0) = \int_0^\infty (1 - e^{-\alpha t})^x dt = \infty \quad (16.23.34)$$

2. For $x, y \in \mathbb{N}_+$ with $x \leq y$, the derivation above and properties of the beta function give

$$U(x, y) = \frac{1}{\alpha} \binom{x}{y} B(y, x - y + 1) = \frac{1}{\alpha} \binom{x}{y} \frac{(y-1)!(x-y)!}{x!} = \frac{1}{\alpha y} \quad (16.23.35)$$

We could argue the results for the potential U directly. Recall that $U(x, y)$ is the expected time spent in state y starting in state x . Since 0 is absorbing and all states lead to 0, $U(x, 0) = \infty$ for $x \in \mathbb{N}$. If $x, y \in \mathbb{N}_+$ and $x \leq y$, then x leads to y with probability 1. Once in state y the time spent in y has an exponential distribution with parameter $\lambda(y) = \alpha y$, and so the mean is $1/\alpha y$. Of course, when the chain leaves y , it never returns.

Recall that βU_β is a transition probability matrix for $\beta > 0$, and in fact $\beta U_\beta(x, \cdot)$ is the probability density function of X_T given $X_0 = x$ where T is independent of \mathbf{X} has the exponential distribution with parameter β . For the next result, recall the *ascending power* notation

$$a^{[k]} = a(a+1) \cdots (a+k-1), \quad a \in \mathbb{R}, k \in \mathbb{N} \quad (16.23.36)$$

For $\beta > 0$ and $x \in \mathbb{N}_+$, the function $\beta U_\beta(x, \cdot)$ is the beta-binomial probability density function with parameters $x, \beta/\alpha$, and 1.

$$\beta U_\beta(x, y) = \binom{x}{y} \frac{(\beta/\alpha)^{[y]} 1^{[x-y]}}{(1 + \beta/\alpha)^{[x]}}, \quad x \in \mathbb{N}, y \in \{0, 1, \dots, x\} \quad (16.23.37)$$

Proof

From the previous result and properties of the beta function.

$$\beta U_\beta(x, y) = \frac{\beta}{\alpha} \binom{x}{y} B(y + \beta/\alpha, x - y + 1), \quad x \in \mathbb{N}, y \in \{0, 1, \dots, x\} \quad (16.23.38)$$

But from properties of the beta function,

$$B(y + \beta/\alpha, x - y + 1) = B(\beta/\alpha, 1) \frac{(\beta/\alpha)^{[y]} 1^{[x-y]}}{(1 + \beta/\alpha)^{[x]}} = \frac{\alpha}{\beta} \frac{(\beta/\alpha)^{[y]} 1^{[x-y]}}{(1 + \beta/\alpha)^{[x]}} \quad (16.23.39)$$

Substituting gives the result

The Yule Process

Next we consider the pure birth branching chain in which each particle, at the end of its life, is replaced by 2 new particles. Equivalently, we can think of particles that never die, but each particle gives birth to a new particle at a constant rate. This chain could serve as the model for an unconstrained nuclear reaction, and is known as the *Yule process*, named for George Yule. So specifically, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be the branching chain with exponential parameter $\alpha \in (0, \infty)$ and offspring probability density function given by $f(2) = 1$. Explosion is inevitable, starting with at least one particle, but other properties of the Yule process are interesting. In particular, there are fascinating parallels with the [pure death branching chain](#). Since 0 is an isolated, absorbing state, we will sometimes restrict our attention to positive states.

The transition matrix of the jump chain and the generator matrix are given by

1. $Q(0, 0) = 1$ and $Q(x, x+1) = 1$ for $x \in \mathbb{N}_+$
2. $G(x, x) = -\alpha x$ for $x \in \mathbb{N}$ and $G(x, x+1) = \alpha x$ for $x \in \mathbb{N}_+$

Since the Yule process is a pure birth process and the birth rate in state $x \in \mathbb{N}$ is αx , the process is also called the *linear birth chain*. As with the pure death process, we can give the distribution of X_t specifically.

Let $t \in [0, \infty)$. Then

1. $m_t = e^{\alpha t}$
2. $v_t = e^{2\alpha t} - e^{\alpha t}$
3. $\Phi_t(r) = \frac{r e^{-\alpha t}}{1 - r + r e^{-\alpha t}}$ for $|r| < \frac{1}{1 - e^{-\alpha t}}$

4. Given $X_0 = x$, X_t has the negative binomial distribution on \mathbb{N}_+ with stopping parameter x and success parameter $e^{-\alpha t}$.

$$P_t(x, y) = \binom{y-1}{x-1} e^{-x\alpha t} (1 - e^{-\alpha t})^{y-x}, \quad x \in \mathbb{N}_+, y \in \{x, x+1, \dots\} \quad (16.23.40)$$

Proof from the general results

Parts (a) and (b) follow from the general moment results above, with $\mu = 2$ and $\sigma^2 = 0$. For part (c), note that $\Psi(r) = r^2$ for $r \in \mathbb{R}$, so the integral equation for Φ_t is

$$\int_r^{\Phi_t(r)} \frac{1}{u^2 - u} = \alpha t \quad (16.23.41)$$

From partial fractions, $\frac{1}{u^2 - u} = \frac{1}{u-1} - \frac{1}{u}$, so the result follows by standard integration and algebra. We recognize Φ_t as the probability generating function of the geometric distribution on \mathbb{N}_+ with success parameter $e^{-\alpha t}$, so for part (d) we use our standard argument. Given $X_0 = x \in \mathbb{N}_+$, X_t has the same distribution as the sum of x independent copies of X_t given $X_0 = 1$, and so this is the distribution of the sum of x independent variables each with the geometric distribution on \mathbb{N}_+ with parameter $e^{-\alpha t}$. But this is the negative binomial distribution on \mathbb{N}_+ with parameters x and $e^{-\alpha t}$.

Direct proof

As usual, let $\tau_0 = 0$ and let τ_n denote the time of the n th transition (birth) for $n \in \mathbb{N}_+$. Given $X_0 = 1$, the population is n at time τ_{n-1} . So the random interval $\tau_n - \tau_{n-1}$ (the time until the next birth) has the exponential distribution with parameter αn and these intervals are independent as n varies. From a result in the section on the exponential distribution, it follows that $\tau_n = \sum_{k=1}^n (\tau_k - \tau_{k-1})$ has distribution function given by

$$\mathbb{P}(\tau_n \leq t \mid X_0 = 1) = (1 - e^{-\alpha t})^n, \quad t \in [0, \infty) \quad (16.23.42)$$

Curiously, this is also the distribution function of the maximum of n independent variables, each with the exponential distribution with rate α . Hence

$$\mathbb{P}(X_t \geq n \mid X_0 = 1) = \mathbb{P}(\tau_{n-1} \leq t \mid X_0 = 1) = (1 - e^{-\alpha t})^{n-1}, \quad n \in \mathbb{N}_+ \quad (16.23.43)$$

and therefore

$$\mathbb{P}(X_t = n \mid X_0 = 1) = \mathbb{P}(X_t \geq n \mid X_0 = 1) - \mathbb{P}(X_t \geq n+1 \mid X_0 = 1) = (1 - e^{-\alpha t})^{n-1} e^{-\alpha t}, \quad n \in \mathbb{N}_+ \quad (16.23.44)$$

So given $X_0 = 1$, X_t has the geometric distribution with parameter $e^{-\alpha t}$. The other results then follow easily.

Recall that the negative binomial distribution with parameters $k \in \mathbb{N}_+$ and $p \in (0, 1)$ governs the trial number of the k th success in a sequence of Bernoulli trials with success parameter p . So the occurrence of this distribution in the Yule process suggests such an interpretation. However this interpretation is not nearly as obvious as with the binomial distribution in the pure death branching chain. Next we give the potential matrices.

For $\beta \in [0, \infty)$ the potential matrix U_β is given by

$$U_\beta(x, y) = \frac{1}{\alpha} \binom{y-1}{x-1} B(x + \beta/\alpha, y - x + 1), \quad x \in \mathbb{N}_+, y \in \{x, x+1, \dots\} \quad (16.23.45)$$

If $\beta > 0$, the function $\beta U_\beta(x, \cdot)$ is the beta-negative binomial probability density function with parameters x , β/α , and 1:

$$\beta U_\beta(x, y) = \binom{y-1}{x-1} \frac{(\beta/\alpha)^{[x]} 1^{[y-x]}}{(1 + \beta/\alpha)^{[y]}}, \quad x \in \mathbb{N}, y \in \{x, x+1, \dots\} \quad (16.23.46)$$

Proof

The proof is very similar to the one above. Suppose that $\beta \geq 0$ and that $x, y \in \mathbb{N}_+$ with $y \geq x$. By definition

$$U_\beta(x, y) = \int_0^\infty e^{-\beta t} P_t(x, y) dt = \int_0^\infty e^{-\beta t} \binom{y-1}{x-1} e^{-\alpha t x} (1 - e^{-\alpha t})^{y-x} dt \quad (16.23.47)$$

Substitute $u = e^{-\alpha t}$ so that $du = -\alpha e^{-\alpha t} dt$ or equivalently $dt = -du/\alpha u$. After some algebra, the result is

$$U_\beta(x, y) = \frac{1}{\alpha} \binom{y-1}{x-1} \int_0^1 u^{x+\beta/\alpha-1} (1-u)^{y-x} du \quad (16.23.48)$$

By definition, the last integral is $B(x + \beta/\alpha, y - x + 1)$.

If we think of the Yule process in terms of particles that never die, but each particle gives birth to a new particle at rate α , then we can study the *age* of the particles at a given time. As usual, we can start with a single, new particle at time 0. So to set up the notation, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be the Yule branching chain with birth rate $\alpha \in (0, \infty)$, and assume that $X_0 = 1$. Let $\tau_0 = 0$ and for $n \in \mathbb{N}_+$, let τ_n denote the time of the n th transition (birth).

For $t \in [0, \infty)$, let A_t denote the total age of the particles at time t . Then

$$A_t = \sum_{n=0}^{X_t-1} (t - \tau_n), \quad t \in [0, \infty) \quad (16.23.49)$$

The random process $\mathbf{A} = \{A_t : t \in [0, \infty)\}$ is the *age process*.

Proof

Note that there have been $X_t - 1$ births in the interval $[0, t]$. For $n \in \{0, 1, \dots, X_t - 1\}$, the age at time t of the particle born at time τ_n is $t - \tau_n$.

Here is another expression for the age process.

Again, let $\mathbf{A} = \{A_t : t \in [0, \infty)\}$ be the age process for the Yule chain starting with a single particle. Then

$$A_t = \int_0^t X_s ds, \quad t \in [0, \infty) \quad (16.23.50)$$

Proof

Suppose that $X_t = k + 1$ where $k \in \mathbb{N}$, so that $\tau_k \leq t < \tau_{k+1}$. Note that $X_s = n$ for $\tau_{n-1} \leq s < \tau_n$ and $n \in \{1, 2, \dots, k\}$, while $X_s = k + 1$ for $\tau_k \leq s \leq t$. Hence

$$\int_0^t X_s ds = \sum_{n=1}^k n(\tau_n - \tau_{n-1}) + (k+1)(t - \tau_k) = (k+1)t - \sum_{n=0}^k \tau_n \quad (16.23.51)$$

From the previous result,

$$A_t = \sum_{n=0}^k (t - \tau_n) = (k+1)t - \sum_{n=0}^k \tau_n \quad (16.23.52)$$

With the last representation, we can easily find the expected total age at time t .

Again, let $\mathbf{A} = \{A_t : t \in [0, \infty)\}$ be the age process for the Yule chain starting with a single particle. Then

$$\mathbb{E}(A_t) = \frac{e^{\alpha t} - 1}{\alpha}, \quad t \in [0, \infty) \quad (16.23.53)$$

Proof

We can interchange the expected value and the integral by Fubini's theorem. So using the moment result [above](#),

$$\mathbb{E}(A_t) = \mathbb{E}\left(\int_0^t X_s ds\right) = \int_0^t \mathbb{E}(X_s) ds = \int_0^t e^{\alpha s} ds = \frac{e^{\alpha t} - 1}{\alpha} \quad (16.23.54)$$

The General Birth-Death Branching Chain

Next we consider the continuous-time branching chain in which each particle, at the end of its life, leaves either no children or two children. At each transition, the number of particles either increases by 1 or decreases by 1, and so such a branching chain is also a continuous-time birth-death chain. Specifically, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be a continuous-time branching chain with lifetime parameter $\alpha \in (0, \infty)$ and offspring probability density function f given by $f(0) = 1 - p$, $f(2) = p$, where $p \in [0, 1]$. When $p = 0$ we have the [pure death chain](#), and when $p = 1$ we have the [Yule process](#). We have already studied these, so the interesting case is when $p \in (0, 1)$ so that both extinction and explosion are possible.

The transition matrix of the jump chain and the generator matrix are given by

1. $Q(0, 0) = 1$, and $Q(x, x-1) = 1-p$, $Q(x, x+1) = p$ for $x \in \mathbb{N}_+$
2. $G(x, x) = -\alpha x$ for $x \in \mathbb{N}$, and $G(x, x-1) = \alpha(1-p)x$, $G(x, x+1) = \alpha p x$ for $x \in \mathbb{N}_+$

As mentioned earlier, \mathbf{X} is also a continuous-time birth-death chain on \mathbb{N} , with 0 absorbing. In state $x \in \mathbb{N}_+$, the birth rate is $\alpha p x$ and the death rate is $\alpha(1-p)x$. The moment functions are given next.

For $t \in [0, \infty)$,

1. $m_t = e^{\alpha(2p-1)t}$
2. If $p \neq \frac{1}{2}$,

$$v_t = \left[\frac{4p(1-p)}{2p-1} + (2p-1) \right] \left[e^{2\alpha(2p-1)t} - e^{\alpha(2p-1)t} \right] \quad (16.23.55)$$

If $p = \frac{1}{2}$, $v_t = 4\alpha p(1-p)t$.

Proof

These results follow from the general formulas above for m_t and v_t , since $\mu = 2p$ and $\sigma^2 = 4p(1-p)$.

The next result gives the generating function of the offspring distribution and the extinction probability.

For the birth-death branching chain,

1. $\Psi(r) = pr^2 + (1-p)$ for $r \in \mathbb{R}$.
2. $q = 1$ if $0 < p \leq \frac{1}{2}$ and $q = \frac{1-p}{p}$ if $\frac{1}{2} < p < 1$.

Proof




Figure 16.23.1 Graphs of $r \mapsto \Psi(r)$ and $r \mapsto r$ when $p = \frac{1}{3}$




Figure 16.23.2 Graphs of $r \mapsto \Psi(r)$ and $r \mapsto r$ when $p = \frac{2}{3}$

For $t \in [0, \infty)$, the generating function Φ_t is given by

$$\Phi_t(r) = \frac{pr - (1-p) + (1-p)(1-r)e^{\alpha(2p-1)t}}{pr - (1-p) + p(1-r)e^{\alpha(2p-1)t}}, \quad \text{if } p \neq 1/2$$

$$\Phi_t(r) = \frac{2r + (1-r)\alpha t}{2 + (1-r)\alpha t}, \quad \text{if } p = \frac{1}{2}$$

Solution

The integral equation for Φ_t is

$$\int_r^{\Phi_t(r)} \frac{du}{pu^2 + (1-p) - u} = \alpha t \quad (16.23.56)$$

The denominator in the integral factors into $(u - 1)[pu - (1 - p)]$. If $p \neq \frac{1}{2}$, use partial fractions, standard integration, and some algebra. If $p = \frac{1}{2}$ the factoring is $\frac{1}{2}(u - 1)^2$ and partial fractions is not necessary. Again, use standard integration and algebra.

This page titled [16.23: Continuous-Time Branching Chains](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

17: Martingales

Martingales, and their cousins *sub-martingales* and *super-martingales* are real-valued stochastic processes that are abstract generalizations of fair, favorable, and unfair gambling processes. The importance of martingales extends far beyond gambling, and indeed these random processes are among the most important in probability theory, with an incredible number and diversity of applications.

[17.1: Introduction to Martingales](#)

[17.2: Properties and Constructions](#)

[17.3: Stopping Times](#)

[17.4: Inequalities](#)

[17.5: Convergence](#)

[17.6: Backwards Martingales](#)

This page titled [17: Martingales](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

17.1: Introduction to Martingales

Basic Theory

Basic Assumptions

For our basic ingredients, we start with a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, having state space \mathbb{R} , and where the index set T (representing time) is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). So to review what all this means, Ω is the sample space, \mathcal{F} the σ -algebra of events, \mathbb{P} the probability measure on (Ω, \mathcal{F}) , and X_t is a random variable with values in \mathbb{R} for each $t \in T$. Next, we have a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, and we assume that \mathbf{X} is *adapted* to \mathfrak{F} . To review again, \mathfrak{F} is an increasing family of sub σ -algebras of \mathcal{F} , so that $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ for $s, t \in T$ with $s \leq t$, and X_t is measurable with respect to \mathcal{F}_t for $t \in T$. We think of \mathcal{F}_t as the collection of events up to time $t \in T$, thus encoding the information available at time t . Finally, we assume that $\mathbb{E}(|X_t|) < \infty$, so that the mean of X_t exists as a real number, for each $t \in T$.

There are two important special cases of the basic setup. The simplest case, of course, is when $\mathcal{F}_t = \sigma\{X_s : s \in T, s \leq t\}$ for $t \in T$, so that \mathfrak{F} is the *natural filtration* associated with \mathbf{X} . Another case that arises frequently is when we have a second stochastic process $\mathbf{Y} = \{Y_t : t \in T\}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ with values in a general measure space (S, \mathcal{S}) , and \mathfrak{F} is the natural filtration associated with \mathbf{Y} . So in this case, our main assumption is that X_t is measurable with respect to $\sigma\{Y_s : s \in T, s \leq t\}$ for $t \in T$.

The theory of martingales is beautiful, elegant, and mostly accessible in discrete time, when $T = \mathbb{N}$. But as with the theory of Markov processes, martingale theory is technically much more complicated in continuous time, when $T = [0, \infty)$. In this case, additional assumptions about the continuity of the sample paths $t \mapsto X_t$ and the filtration $t \mapsto \mathcal{F}_t$ are often necessary in order to have a nice theory. Specifically, we will assume that the process \mathbf{X} is right continuous and has left limits, and that the filtration \mathfrak{F} is right continuous and complete. These are the *standard assumptions* in continuous time.

Definitions

For the basic definitions that follow, you may need to review conditional expected value with respect to a σ -algebra.

The process \mathbf{X} is a *martingale* with respect to \mathfrak{F} if $\mathbb{E}(X_t | \mathcal{F}_s) = X_s$ for all $s, t \in T$ with $s \leq t$.

In the special case that \mathfrak{F} is the natural filtration associated with \mathbf{X} , we simply say that \mathbf{X} is a martingale, without reference to the filtration. In the special case that we have a second stochastic process $\mathbf{Y} = \{Y_t : t \in T\}$ and \mathfrak{F} is the natural filtration associated with \mathbf{Y} , we say that \mathbf{X} is a martingale with respect to \mathbf{Y} .

The term *martingale* originally referred to a portion of the harness of a horse, and was later used to describe gambling strategies, such as the one used in the Petersburg paradox, in which bets are doubled when a game is lost. To interpret the definitions above in terms of gambling, suppose that a gambler is at a casino, and that X_t represents her fortune at time $t \in T$ and \mathcal{F}_t the information available to her at time t . Suppose now that $s, t \in T$ with $s < t$ and that we think of s as the current time, so that t is a future time. If \mathbf{X} is a martingale with respect to \mathfrak{F} then the games are fair in the sense that the gambler's expected fortune at the future time t is the same as her current fortune at time s . To venture a bit from the casino, suppose that X_t is the price of a stock, or the value of a stock index, at time $t \in T$. If \mathbf{X} is a martingale, then the expected value at a future time, given all of our information, is the present value.



Figure 17.1.1: An English-style breastplate with a running martingale attachment. By Danielle M., CC BY 3.0, from Wikipedia

But as we will see, martingales are useful in probability far beyond the application to gambling and even far beyond financial applications generally. Indeed, martingales are of fundamental importance in modern probability theory. Here are two related definitions, with equality in the martingale condition replaced by inequalities.

Suppose again that the process \mathbf{X} and the filtration \mathfrak{F} satisfy the [basic assumptions](#) above.

1. \mathbf{X} is a *sub-martingale* with respect to \mathfrak{F} if $\mathbb{E}(X_t | \mathcal{F}_s) \geq X_s$ for all $s, t \in T$ with $s \leq t$.

2. \mathbf{X} is a *super-martingale* with respect to \mathfrak{F} if $\mathbb{E}(X_t | \mathcal{F}_s) \leq X_s$ for all $s, t \in T$ with $s \leq t$.

In the gambling setting, a sub-martingale models games that are favorable to the gambler on average, while a super-martingale models games that are unfavorable to the gambler on average. To venture again from the casino, suppose that X_t is the price of a stock, or the value of a stock index, at time $t \in T$. If \mathbf{X} is a sub-martingale, the expected value at a future time, given all of our information, is greater than the present value, and if \mathbf{X} is a super-martingale then the expected value at the future time is less than the present value. One hopes that a stock index is a sub-martingale.

Clearly \mathbf{X} is a martingale with respect to \mathfrak{F} if and only if it is both a sub-martingale and a super-martingale. Finally, recall that the conditional expected value of a random variable with respect to a σ -algebra is itself a random variable, and so the equations and inequalities in the definitions should be interpreted as holding with probability 1. In this section generally, statements involving random variables are assumed to hold with probability 1.

The conditions that define martingale, sub-martingale, and super-martingale make sense if the index set T is any totally ordered set. In some applications that we will consider later, $T = \{0, 1, \dots, n\}$ for fixed $n \in \mathbb{N}_+$. In the section on backwards martingales, $T = \{-n : n \in \mathbb{N}\}$ or $T = (-\infty, 0]$. In the case of discrete time when $T = \mathbb{N}$, we can simplify the definitions slightly.

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ satisfies the basic assumptions above.

1. \mathbf{X} is a martingale with respect to \mathfrak{F} if and only if $\mathbb{E}(X_{n+1} | \mathcal{F}_n) = X_n$ for all $n \in \mathbb{N}$.
2. \mathbf{X} is a sub-martingale with respect to \mathfrak{F} if and only if $\mathbb{E}(X_{n+1} | \mathcal{F}_n) \geq X_n$ for all $n \in \mathbb{N}$.
3. \mathbf{X} is a super-martingale with respect to \mathfrak{F} if and only if $\mathbb{E}(X_{n+1} | \mathcal{F}_n) \leq X_n$ for all $n \in \mathbb{N}$.

Proof

The conditions in the definitions clearly imply the conditions here, so we just need to show the opposite implications. Thus, assume that the condition in (a) holds and suppose that $k, n \in \mathbb{N}$ with $k < n$. Then $k \leq n-1$ so $\mathcal{F}_k \subseteq \mathcal{F}_{n-1}$ and hence

$$\mathbb{E}(X_n | \mathcal{F}_k) = \mathbb{E}[\mathbb{E}(X_n | \mathcal{F}_{n-1}) | \mathcal{F}_k] = \mathbb{E}(X_{n-1} | \mathcal{F}_k) \quad (17.1.1)$$

Repeating the argument, we get to

$$\mathbb{E}(X_n | \mathcal{F}_k) = \mathbb{E}(X_{k+1} | \mathcal{F}_k) = X_k \quad (17.1.2)$$

The proof for sub and super-martingales is analogous, with inequalities replacing the last equality.

The relations that define martingales, sub-martingales, and super-martingales hold for the ordinary (unconditional) expected values.

Suppose that $s, t \in T$ with $s \leq t$.

1. If \mathbf{X} is a martingale with respect to \mathfrak{F} then $\mathbb{E}(X_s) = \mathbb{E}(X_t)$.
2. If \mathbf{X} is a sub-martingale with respect to \mathfrak{F} then $\mathbb{E}(X_s) \leq \mathbb{E}(X_t)$.
3. If \mathbf{X} is a super-martingale with respect to \mathfrak{F} then $\mathbb{E}(X_s) \geq \mathbb{E}(X_t)$.

Proof

The results follow directly from the definitions, and the critical fact that $\mathbb{E}[\mathbb{E}(X_t | \mathcal{F}_s)] = \mathbb{E}(X_t)$ for $s, t \in T$.

So if \mathbf{X} is a martingale then \mathbf{X} has constant expected value, and this value is referred to as the mean of \mathbf{X} .

Examples

The goal for the remainder of this section is to give some classical examples of martingales, and by doing so, to show the wide variety of applications in which martingales occur. We will return to many of these examples in subsequent sections. Without further ado, we assume that all random variables are real-valued, unless otherwise specified, and that all expected values mentioned below exist in \mathbb{R} . Be sure to try the proofs yourself before expanding the ones in the text.

Constant Sequence

Our first example is rather trivial, but still worth noting.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and that X is a random variable that is measurable with respect to \mathcal{F}_0 and satisfies $\mathbb{E}(|X|) < \infty$. Let $X_t = X$ for $t \in T$. Then $\mathbf{X} = \{X_t : t \in T\}$ is a martingale with respect to \mathfrak{F} .

Proof

Since X is measurable with respect to \mathcal{F}_0 , it is measurable with respect to \mathcal{F}_t for all $t \in T$. Hence \mathbf{X} is adapted to \mathfrak{F} . If $s, t \in T$ with $s \leq t$, then $\mathbb{E}(X_t | \mathcal{F}_s) = \mathbb{E}(X | \mathcal{F}_s) = X = X_s$.

Partial Sums

For our next discussion, we start with one of the most basic martingales in discrete time, and the one with the simplest interpretation in terms of gambling. Suppose that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is a sequence of independent random variables with $\mathbb{E}(|V_k|) < \infty$ for $k \in \mathbb{N}$. Let

$$X_n = \sum_{k=0}^n V_k, \quad n \in \mathbb{N} \quad (17.1.3)$$

so that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is simply the *partial sum process* associated with \mathbf{V} .

For the partial sum process \mathbf{X} ,

1. If $\mathbb{E}(V_n) \geq 0$ for $n \in \mathbb{N}_+$ then \mathbf{X} is a sub-martingale.
2. If $\mathbb{E}(V_n) \leq 0$ for $n \in \mathbb{N}_+$ then \mathbf{X} is a super-martingale.
3. If $\mathbb{E}(V_n) = 0$ for $n \in \mathbb{N}_+$ then \mathbf{X} is a martingale.

Proof

Let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\} = \sigma\{V_0, V_1, \dots, V_n\}$ for $n \in \mathbb{N}$. Note first that

$$\mathbb{E}(|X_n|) \leq \sum_{k=0}^n \mathbb{E}(|V_k|) < \infty, \quad n \in \mathbb{N} \quad (17.1.4)$$

Next,

$$\mathbb{E}(X_{n+1} | \mathcal{F}_n) = \mathbb{E}(X_n + V_{n+1} | \mathcal{F}_n) = \mathbb{E}(X_n | \mathcal{F}_n) + \mathbb{E}(V_{n+1} | \mathcal{F}_n) = X_n + \mathbb{E}(V_{n+1}), \quad n \in \mathbb{N} \quad (17.1.5)$$

The last equality holds since X_n is measurable with respect to \mathcal{F}_n and V_{n+1} is independent of \mathcal{F}_n . The results now follow from the definitions.

In terms of gambling, if $X_0 = V_0$ is the gambler's initial fortune and V_i is the gambler's net winnings on the i th game, then X_n is the gamblers net fortune after n games for $n \in \mathbb{N}_+$. But partial sum processes associated with independent sequences are important far beyond gambling. In fact, much of classical probability deals with partial sums of independent and *identically distributed* variables. The entire chapter on Random Samples explores this setting.

Note that $\mathbb{E}(X_n) = \sum_{k=0}^n \mathbb{E}(V_k)$. Hence condition (a) is equivalent to $n \mapsto \mathbb{E}(X_n)$ increasing, condition (b) is equivalent to $n \mapsto \mathbb{E}(X_n)$ decreasing, and condition (c) is equivalent to $n \mapsto \mathbb{E}(X_n)$ constant. Here is another martingale associated with the partial sum process, known as the *second moment martingale*.

Suppose that $\mathbb{E}(V_k) = 0$ for $k \in \mathbb{N}_+$ and $\text{var}(V_k) < \infty$ for $k \in \mathbb{N}$. Let

$$Y_n = X_n^2 - \text{var}(X_n), \quad n \in \mathbb{N} \quad (17.1.6)$$

Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} .

Proof

Again, let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$ for $n \in \mathbb{N}$. Since the sequence \mathbf{V} is independent, note that

$$\text{var}(X_n) = \text{var}\left(\sum_{k=0}^n V_k\right) = \sum_{k=0}^n \text{var}(V_k) \quad (17.1.7)$$

Also, $\text{var}(V_k) = \mathbb{E}(V_k^2)$ since $\mathbb{E}(V_k) = 0$ for $k \in \mathbb{N}_+$. In particular, $\mathbb{E}(|Y_n|) < \infty$ for $n \in \mathbb{N}$. Next for $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{E}(Y_{n+1} | \mathcal{F}_n) &= \mathbb{E}[X_{n+1}^2 - \text{var}(X_{n+1}) | \mathcal{F}_n] = \mathbb{E}[(X_n + V_{n+1})^2 - \text{var}(X_{n+1}) | \mathcal{F}_n] \\ &= \mathbb{E}[X_n^2 + 2X_n V_{n+1} + V_{n+1}^2 - \text{var}(X_{n+1}) | \mathcal{F}_n] = X_n^2 + 2X_n \mathbb{E}(V_{n+1}) + \mathbb{E}(V_{n+1}^2) - \text{var}(X_{n+1}) \end{aligned}$$

since X_n is measurable with respect to \mathcal{F}_n and V_{n+1} is independent of \mathcal{F}_n . But $\mathbb{E}(V_{n+1}) = 0$ and $\mathbb{E}(V_{n+1}^2) - \text{var}(X_{n+1}) = -\text{var}(X_n)$. Hence we have $\mathbb{E}(Y_{n+1} | \mathcal{F}_n) = X_n^2 - \text{var}(X_n) = Y_n$ for $n \in \mathbb{N}$.

So under the assumptions in this theorem, both \mathbf{X} and \mathbf{Y} are martingales. We will generalize the results for partial sum processes below in the discussion on [processes with independent increments](#).

Martingale Difference Sequences

In the last discussion, we saw that the partial sum process associated with a sequence of independent, mean 0 variables is a martingale. Conversely, every martingale in discrete time can be written as a partial sum process of *uncorrelated* mean 0 variables. This representation gives some significant insight into the theory of martingales generally. Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a martingale with respect to the filtration $\mathcal{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$.

Let $V_0 = X_0$ and $V_n = X_n - X_{n-1}$ for $n \in \mathbb{N}_+$. The process $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is the *martingale difference sequence* associated with \mathbf{X} and

$$X_n = \sum_{k=0}^n V_k, \quad n \in \mathbb{N} \quad (17.1.8)$$

As promised, the martingale difference variables have mean 0, and in fact satisfy a stronger property.

Suppose that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is the martingale difference sequence associated with \mathbf{X} . Then

1. \mathbf{V} is adapted to \mathcal{F} .
2. $\mathbb{E}(V_n | \mathcal{F}_k) = 0$ for $k, n \in \mathbb{N}$ with $k < n$.
3. $\mathbb{E}(V_n) = 0$ for $n \in \mathbb{N}_+$.

Proof

1. Of course $V_0 = X_0$ is measurable with respect to \mathcal{F}_0 . For $n \in \mathbb{N}_+$, X_n and X_{n-1} , and hence V_n are measurable with respect to \mathcal{F}_n .
2. Let $k \in \mathbb{N}$. By the martingale and adapted properties,

$$\mathbb{E}(V_{k+1} | \mathcal{F}_k) = \mathbb{E}(X_{k+1} | \mathcal{F}_k) - \mathbb{E}(X_k | \mathcal{F}_k) = X_k - X_k = 0 \quad (17.1.9)$$

Next by the tower property,

$$\mathbb{E}(V_{k+2} | \mathcal{F}_k) = \mathbb{E}[\mathbb{E}(V_{k+2} | \mathcal{F}_{k+1}) | \mathcal{F}_k] = 0 \quad (17.1.10)$$

Continuing (or using induction) gives the general result.

3. Since \mathbf{X} is a martingale, it has constant mean, as noted above. Hence $\mathbb{E}(V_n) = \mathbb{E}(X_n) - \mathbb{E}(X_{n-1}) = 0$ for $n \in \mathbb{N}_+$. We could also use part (b).

Also as promised, if the martingale variables have finite variance, then the martingale difference variables are uncorrelated.

Suppose again that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is the martingale difference sequence associated with the martingale \mathbf{X} . Assume that $\text{var}(X_n) < \infty$ for $n \in \mathbb{N}$. Then \mathbf{V} is an uncorrelated sequence. Moreover,

$$\text{var}(X_n) = \sum_{k=0}^n \text{var}(V_k) = \text{var}(X_0) + \sum_{k=1}^n \mathbb{E}(V_k^2), \quad n \in \mathbb{N} \quad (17.1.11)$$

Proof

Let $k, n \in \mathbb{N}$ with $k < n$. To show that V_k and V_n are uncorrelated, we just need to show that $\mathbb{E}(V_k V_n) = 0$ (since $\mathbb{E}(V_n) = 0$). But by the previous result,

$$\mathbb{E}(V_k V_n) = \mathbb{E}[\mathbb{E}(V_k V_n | \mathcal{F}_k)] = \mathbb{E}[V_k \mathbb{E}(V_n | \mathcal{F}_k)] = 0 \quad (17.1.12)$$

Finally, the variance of a sum of uncorrelated variables is the sum of the variances. Since V_k has mean 0, $\text{var}(V_k) = \mathbb{E}(V_k^2)$ for $k \in \mathbb{N}_+$. Hence the formula for $\text{var}(X_n)$ holds.

We now know that a discrete-time martingale is the partial sum process associated with a sequence of uncorrelated variables. Hence we might hope that there are martingale versions of the fundamental theorems that hold for a partial sum process associated with an independent sequence. This turns out to be true, and is a basic reason for the importance of martingales.

Discrete-Time Random Walks

Suppose that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is a sequence of independent random variables with $\{V_n : n \in \mathbb{N}_+\}$ identically distributed. We assume that $\mathbb{E}(|V_n|) < \infty$ for $n \in \mathbb{N}$ and we let a denote the common mean of $\{V_n : n \in \mathbb{N}_+\}$. Let $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ be the partial sum process associated with \mathbf{V} so that

$$X_n = \sum_{i=0}^n V_i, \quad n \in \mathbb{N} \quad (17.1.13)$$

This setting is a special case of the more general partial sum process considered [above](#). The process \mathbf{X} is sometimes called a (discrete-time) *random walk*. The initial position $X_0 = V_0$ of the walker can have an arbitrary distribution, but then the *steps* that the walker takes are independent and identically distributed. In terms of gambling, $X_0 = V_0$ is the initial fortune of the gambler playing a sequence of independent and identical games. If V_i is the amount won (or lost) on game $i \in \mathbb{N}_+$, then X_n is the gambler's net fortune after n games.

For the random walk \mathbf{X} ,

1. \mathbf{X} is a martingale if $a = 0$.
2. \mathbf{X} is a sub-martingale if $a \geq 0$.
3. \mathbf{X} is a super-martingale if $a \leq 0$.

For the second moment martingale, suppose that V_n has common mean $a = 0$ and common variance $b^2 < \infty$ for $n \in \mathbb{N}_+$, and that $\text{var}(V_0) < \infty$.

Let $Y_n = X_n^2 - \text{var}(V_0) - b^2 n$ for $n \in \mathbb{N}$. Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} .

Proof

This follows from the corresponding result for a general partial sum process, above, since

$$\text{var}(X_n) = \sum_{k=0}^n \text{var}(V_k) = \text{var}(V_0) + b^2 n, \quad n \in \mathbb{N} \quad (17.1.14)$$

We will generalize the results for discrete-time random walks below, in the discussion on [processes with stationary, independent increments](#).

Partial Products

Our next discussion is similar to the one on partial sum processes [above](#), but with products instead of sums. So suppose that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is an independent sequence of nonnegative random variables with $\mathbb{E}(V_n) < \infty$ for $n \in \mathbb{N}$. Let

$$X_n = \prod_{i=0}^n V_i, \quad n \in \mathbb{N} \quad (17.1.15)$$

so that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is the *partial product process* associated with \mathbf{X} .

For the partial product process \mathbf{X} ,

1. If $\mathbb{E}(V_n) = 1$ for $n \in \mathbb{N}_+$ then \mathbf{X} is a martingale with respect to \mathbf{V}
2. If $\mathbb{E}(V_n) \geq 1$ for $n \in \mathbb{N}_+$ then \mathbf{X} is a sub-martingale with respect to \mathbf{V}
3. If $\mathbb{E}(V_n) \leq 1$ for $n \in \mathbb{N}_+$ then \mathbf{X} is a super-martingale with respect to \mathbf{V}

Proof

Let $\mathcal{F}_n = \sigma\{V_0, V_1, \dots, V_n\}$ for $n \in \mathbb{N}$. Since the variables are independent,

$$\mathbb{E}(X_n) = \prod_{i=0}^n \mathbb{E}(V_i) < \infty \quad (17.1.16)$$

Next,

$$\mathbb{E}(X_{n+1} | \mathcal{F}_n) = \mathbb{E}(X_n V_{n+1} | \mathcal{F}_n) = X_n \mathbb{E}(V_{n+1} | \mathcal{F}_n) = X_n \mathbb{E}(V_{n+1}) \quad n \in \mathbb{N} \quad (17.1.17)$$

since X_n is measurable with respect to \mathcal{F}_n and V_{n+1} is independent of \mathcal{F}_n . The results now follow from the definitions.

As with random walks, a special case of interest is when $\{V_n : n \in \mathbb{N}_+\}$ is an identically distributed sequence.

The Simple Random Walk

Suppose now that that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is a sequence of independent random variables with $\mathbb{P}(V_i = 1) = p$ and $\mathbb{P}(V_i = -1) = 1 - p$ for $i \in \mathbb{N}_+$, where $p \in (0, 1)$. Let $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ be the partial sum process associated with \mathbf{V} so that

$$X_n = \sum_{i=0}^n V_i, \quad n \in \mathbb{N} \quad (17.1.18)$$

Then \mathbf{X} is the simple random walk with parameter p , and of course, is a special case of the more general random walk studied [above](#). In terms of gambling, our gambler plays a sequence of independent and identical games, and on each game, wins €1 with probability p and loses €1 with probability $1 - p$. So if V_0 is the gambler's initial fortune, then X_n is her net fortune after n games.

For the simple random walk,

1. If $p > \frac{1}{2}$ then \mathbf{X} is a sub-martingale.
2. If $p < \frac{1}{2}$ then \mathbf{X} is a super-martingale.
3. If $p = \frac{1}{2}$ then \mathbf{X} is a martingale.

Proof

Note that $\mathbb{E}(V_n) = p - (1 - p) = 2p - 1$ for $n \in \mathbb{N}_+$, so the results follow from the theorem [above](#).

So case (a) corresponds to favorable games, case (b) to unfavorable games, and case (c) to fair games.

Open the simulation of the simple symmetric random. For various values of the number of trials n , run the simulation 1000 times and note the general behavior of the sample paths.

Here is the second moment martingale for the simple, symmetric random walk.

Consider the simple random walk with parameter $p = \frac{1}{2}$, and let $Y_n = X_n^2 - \text{var}(V_0) - n$ for $n \in \mathbb{N}$. Then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X}

Proof

Note that $\mathbb{E}(V_i) = 0$ and $\text{var}(V_i) = 1$ for each $i \in \mathbb{N}_+$, so the result follows from the general result above.

But there is another martingale that can be associated with the simple random walk, known as *De Moivre's martingale* and named for one of the early pioneers of probability theory, Abraham De Moivre.

For $n \in \mathbb{N}$ define

$$Z_n = \left(\frac{1-p}{p} \right)^{X_n} \quad (17.1.19)$$

Then $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} .

Proof

Note that

$$Z_n = \prod_{k=0}^n \left(\frac{1-p}{p} \right)^{V_k}, \quad n \in \mathbb{N} \quad (17.1.20)$$

and

$$\mathbb{E} \left[\left(\frac{1-p}{p} \right)^{V_k} \right] = \left(\frac{1-p}{p} \right)^1 p + \left(\frac{1-p}{p} \right)^{-1} (1-p) = 1, \quad k \in \mathbb{N}_+ \quad (17.1.21)$$

So the result follows from the theorem [above](#) on partial products.

The Beta-Bernoulli Process

Recall that the beta-Bernoulli process is constructed by randomizing the success parameter in a Bernoulli trials process with a beta distribution. Specifically we have a random variable P that has the beta distribution with parameters $a, b \in (0, \infty)$, and a sequence of indicator variables $\mathbf{X} = (X_1, X_2, \dots)$ such that given $P = p \in (0, 1)$, \mathbf{X} is a sequence of independent variables with $\mathbb{P}(X_i = 1) = p$ for $i \in \mathbb{N}_+$. As usual, we couch this in reliability terms, so that $X_i = 1$ means *success* on trial i and $X_i = 0$ means *failure*. In our study of this process, we showed that the finite-dimensional distributions are given by

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) = \frac{a^{[k]} b^{[n-k]}}{(a+b)^{[n]}}, \quad n \in \mathbb{N}_+, (x_1, x_2, \dots, x_n) \in \{0, 1\}^n \quad (17.1.22)$$

where we use the ascending power notation $r^{[j]} = r(r+1) \cdots (r+j-1)$ for $r \in \mathbb{R}$ and $j \in \mathbb{N}$. Next, let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ denote the partial sum process associated with \mathbf{X} , so that once again,

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (17.1.23)$$

Of course Y_n is the number of success in the first n trials and has the beta-binomial distribution defined by

$$\mathbb{P}(Y_n = k) = \binom{n}{k} \frac{a^{[k]} b^{[n-k]}}{(a+b)^{[n]}}, \quad k \in \{0, 1, \dots, n\} \quad (17.1.24)$$

Now let

$$Z_n = \frac{a + Y_n}{a + b + n}, \quad n \in \mathbb{N} \quad (17.1.25)$$

This variable also arises naturally. Let $\mathcal{F}_n = \sigma\{X_1, X_2, \dots, X_n\}$ for $n \in \mathbb{N}$. Then as shown in the section on the beta-Bernoulli process, $Z_n = \mathbb{E}(X_{n+1} | \mathcal{F}_n) = \mathbb{E}(P | \mathcal{F}_n)$. In statistical terms, the second equation means that Z_n is the *Bayesian estimator* of the unknown success probability p in a sequence of Bernoulli trials, when p is modeled by the random variable P .

$\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} .

Proof

Note that $0 \leq Z_n \leq 1$ so $\mathbb{E}(Z_n) < \infty$ for $n \in \mathbb{N}$. Next,

$$\mathbb{E}(Z_{n+1} | \mathcal{F}_n) = \mathbb{E}\left[\frac{a + Y_{n+1}}{a + b + n + 1} \mid \mathcal{F}_n\right] = \frac{\mathbb{E}[a + (Y_n + X_{n+1}) | \mathcal{F}_n]}{a + b + n + 1} = \frac{a + Y_n + \mathbb{E}(X_{n+1} | \mathcal{F}_n)}{a + b + n + 1} \quad (17.1.26)$$

As noted above, $\mathbb{E}(X_{n+1} | \mathcal{F}_n) = (a + Y_n)/(a + b + n)$. Substituting into the displayed equation above and doing a bit of algebra we have

$$\mathbb{E}(Z_{n+1} | \mathcal{F}_n) = \frac{(a + Y_n) + (a + Y_n)/(a + b + n)}{a + b + n + 1} = \frac{a + Y_n}{a + b + n} = Z_n \quad (17.1.27)$$

Open the beta-Binomial experiment. Run the simulation 1000 times for various values of the parameters, and compare the empirical probability density function with the true probability density function.

Pólya's Urn Process

Recall that in the simplest version of Pólya's urn process, we start with an urn containing a red and b green balls. At each discrete time step, we select a ball at random from the urn and then replace the ball and add c new balls of the same color to the urn. For the parameters, we need $a, b \in \mathbb{N}_+$ and $c \in \mathbb{N}$. For $i \in \mathbb{N}_+$, let X_i denote the color of the ball selected on the i th draw, where 1 means red and 0 means green. The process $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is a classical example of a sequence of exchangeable yet dependent variables. Let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ denote the partial sum process associated with \mathbf{X} , so that once again,

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (17.1.28)$$

Of course Y_n is the total number of red balls selected in the first n draws. Hence at time $n \in \mathbb{N}$, the total number of red balls in the urn is $a + cY_n$, while the total number of balls in the urn is $a + b + cn$ and so the *proportion* of red balls in the urn is

$$Z_n = \frac{a + cY_n}{a + b + cn} \quad (17.1.29)$$

$\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} .

Indirect proof

If $c = 0$ then $Z_n = a/(a+b)$ for $n \in \mathbb{N}$ so \mathbf{Z} is a constant martingale. If $c \in \mathbb{N}_+$ then \mathbf{Z} is equivalent to the beta-Bernoulli process with parameters a/c and b/c . Moreover,

$$Z_n = \frac{a + cY_n}{a + b + cn} = \frac{a/c + Y_n}{a/c + b/c + n}, \quad n \in \mathbb{N} \quad (17.1.30)$$

So \mathbf{Z} is a martingale by the previous theorem.

Direct Proof

Trivially, $0 \leq Z_n \leq 1$ so $\mathbb{E}(Z_n) < \infty$ for $n \in \mathbb{N}$. Let $\mathcal{F}_n = \sigma\{X_1, X_2, \dots, X_n\}$. For $n \in \mathbb{N}$,

$$\mathbb{E}(Z_{n+1} | \mathcal{F}_n) = \mathbb{E}\left[\frac{a + cY_{n+1}}{a + b + c(n+1)} \mid \mathcal{F}_n\right] = \frac{\mathbb{E}[a + c(Y_n + X_{n+1}) | \mathcal{F}_n]}{a + b + c(n+1)} = \frac{a + cY_n + c\mathbb{E}(X_{n+1} | \mathcal{F}_n)}{a + b + cn + c} \quad (17.1.31)$$

since Y_n is measurable with respect to \mathcal{F}_n . But the probability of selecting a red ball on draw $n+1$, given the history of the process up to time n , is simply the proportion of red balls in the urn at time n . That is,

$$\mathbb{E}(X_{n+1} | \mathcal{F}_n) = \mathbb{P}(X_{n+1} = 1 | \mathcal{F}_n) = Z_n = \frac{a + cY_n}{a + b + cn} \quad (17.1.32)$$

Substituting and simplifying gives $\mathbb{E}(Z_{n+1} | \mathcal{F}_n) = Z_n$.

Open the simulation of Pólya's Urn Experiment. Run the simulation 1000 times for various values of the parameters, and compare the empirical probability density function of the number of red ball selected to the true probability density function.

Processes with Independent Increments.

Our first example above concerned the [partial sum process](#) \mathbf{X} associated with a sequence of independent random variables \mathbf{V} . Such processes are the only ones in discrete time that have *independent increments*. That is, for $m, n \in \mathbb{N}$ with $m \leq n$, $X_n - X_m$ is independent of (X_0, X_1, \dots, X_m) . The [random walk process](#) has the additional property of *stationary increments*. That is, the distribution of $X_n - X_m$ is the same as the distribution of $X_{n-m} - X_0$ for $m, n \in \mathbb{N}$ with $m \leq n$. Let's consider processes in discrete or continuous time with these properties. Thus, suppose that $\mathbf{X} = \{X_t : t \in T\}$ satisfying the [basic assumptions](#) above relative to the filtration $\mathcal{F} = \{\mathcal{F}_s : s \in T\}$. Here are the two definitions.

The process \mathbf{X} has

1. *Independent increments* if $X_t - X_s$ is independent of \mathcal{F}_s for all $s, t \in T$ with $s \leq t$.
2. *Stationary increments* if $X_t - X_s$ has the same distribution as $X_{t-s} - X_0$ for all $s, t \in T$.

Processes with stationary and independent increments were studied in the Chapter on Markov processes. In continuous time (with the continuity assumptions we have imposed), such a process is known as a *Lévy process*, named for Paul Lévy, and also as a *continuous-time random walk*. For a process with independent increments (not necessarily stationary), the connection with martingales depends on the mean function m given by $m(t) = \mathbb{E}(X_t)$ for $t \in T$.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ has independent increments.

1. If m is increasing then \mathbf{X} is a sub-martingale.
2. If m is decreasing then \mathbf{X} is a super-martingale.
3. If m is constant then \mathbf{X} is a martingale

Proof

The proof is just like the one above for partial sum processes. Suppose that $s, t \in [0, \infty)$ with $s < t$. Then

$$\mathbb{E}(X_t | \mathcal{F}_s) = \mathbb{E}[X_s + (X_t - X_s) | \mathcal{F}_s] = \mathbb{E}(X_s | \mathcal{F}_s) + \mathbb{E}(X_t - X_s | \mathcal{F}_s) \quad (17.1.33)$$

But X_s is measurable with respect to \mathcal{F}_s and $X_t - X_s$ is independent of \mathcal{F}_s . So

$$\mathbb{E}(X_t | \mathcal{F}_s) = X_s + \mathbb{E}(X_t - X_s) = X_s + m(t) - m(s) \quad (17.1.34)$$

Compare this theorem with the corresponding theorem for the [partial sum process](#) above. Suppose now that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a stochastic process as above, with mean function m , and let $Y_t = X_t - m(t)$ for $t \in [0, \infty)$. The process $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is sometimes called the *compensated process* associated with \mathbf{X} and has mean function 0. If \mathbf{X} has independent increments, then clearly so does \mathbf{Y} . Hence the following result is a trivial corollary to our previous theorem.

Suppose that \mathbf{X} has independent increments. The compensated process \mathbf{Y} is a martingale.

Next we give the *second moment martingale* for a process with independent increments, generalizing the [second moment martingale](#) for a partial sum process.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ has independent increments with constant mean function and with $\text{var}(X_t) < \infty$ for $t \in T$. Then $\mathbf{Y} = \{Y_t : t \in T\}$ is a martingale where

$$Y_t = X_t^2 - \text{var}(X_t), \quad t \in T \quad (17.1.35)$$

Proof

The proof is essentially the same as for the partial sum process in discrete time. Suppose that $s, t \in T$ with $s < t$. Note that $\mathbb{E}(Y_t | \mathcal{F}_s) = \mathbb{E}(X_t^2 | \mathcal{F}_s) - \text{var}(X_t)$. Next,

$$X_t^2 = [(X_t - X_s) + X_s]^2 = (X_t - X_s)^2 + 2(X_t - X_s)X_s + X_s^2 \quad (17.1.36)$$

But $X_t - X_s$ is independent of \mathcal{F}_s , X_s is measurable with respect to \mathcal{F}_s , and $\mathbb{E}(X_t - X_s) = 0$ so

$$\mathbb{E}(X_t^2 | \mathcal{F}_s) = \mathbb{E}[(X_t - X_s)^2] + 2X_s \mathbb{E}(X_t - X_s) + X_s^2 = \mathbb{E}[(X_t - X_s)^2] + X_s^2 \quad (17.1.37)$$

But also by independence and since $X_t - X_s$ has mean 0,

$$\text{var}(X_t) = \text{var}[(X_t - X_s) + X_s] = \text{var}(X_s) + \text{var}(X_t - X_s)^2 = \text{var}(X_s) + \mathbb{E}[(X_t - X_s)^2] \quad (17.1.38)$$

Putting the pieces together gives

$$\mathbb{E}(Y_t | \mathcal{F}_s) = X_s^2 - \text{var}(X_s) = Y_s \quad (17.1.39)$$

Of course, since the mean function is constant, \mathbf{X} is also a martingale. For processes with independent *and* stationary increments (that is, random walks), the last two theorems simplify, because the mean and variance functions simplify.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ has stationary, independent increments, and let $a = \mathbb{E}(X_1 - X_0)$. Then

1. \mathbf{X} is a martingale if $a = 0$
2. \mathbf{X} is a sub-martingale if $a \geq 0$
3. \mathbf{X} is a super-martingale if $a \leq 0$

Proof

Recall that the mean function m is given by $m(t) = \mathbb{E}(X_0) + at$ for $t \in T$, so the result follows from the [corresponding result](#) for a process with independent increments.

Compare this result with the [corresponding one](#) above for discrete-time random walks. Our next result is the second moment martingale. Compare this with the [second moment martingale](#) for discrete-time random walks.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ has stationary, independent increments with $\mathbb{E}(X_0) = \mathbb{E}(X_1)$ and $b^2 = \mathbb{E}(X_1^2) < \infty$. Then $\mathbf{Y} = \{Y_t : t \in T\}$ is a martingale where

$$Y_t = X_t^2 - \text{var}(X_0) - b^2 t, \quad t \in T \quad (17.1.40)$$

Proof

Recall that if $\mathbb{E}(X_0) = \mathbb{E}(X_1)$ then \mathbf{X} has constant mean function. Also, $\text{var}(X_t) = \text{var}(X_0) + b^2 t$, so the result follows from the corresponding result for a process with independent increments.

In discrete time, as we have mentioned several times, all of these results reduce to the earlier results for [partial sum processes](#) and [random walks](#). In continuous time, the Poisson processes, named of course for Simeon Poisson, provides examples. The standard (homogeneous) Poisson counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ with constant rate $r \in (0, \infty)$ has stationary, independent increments and mean function given by $m(t) = rt$ for $t \in [0, \infty)$. More generally, suppose that $r : [0, \infty) \rightarrow (0, \infty)$ is piecewise continuous (and non-constant). The non-homogeneous Poisson counting process $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ with rate function r has independent increments and mean function given by

$$m(t) = \int_0^t r(s) ds, \quad t \in [0, \infty) \quad (17.1.41)$$

The increment $N_t - N_s$ has the Poisson distribution with parameter $m(t) - m(s)$ for $s, t \in [0, \infty)$ with $s < t$, so the process does not have stationary increments. In all cases, m is increasing, so the following results are corollaries of our general results:

Let $\mathbf{N} = \{N_t : t \in [0, \infty)\}$ be the Poisson counting process with rate function $r : [0, \infty) \rightarrow (0, \infty)$. Then

1. \mathbf{N} is a sub-martingale
2. The compensated process $\mathbf{X} = \{N_t - m(t) : t \in [0, \infty)\}$ is a martingale.

Open the simulation of the Poisson counting experiment. For various values of r and t , run the experiment 1000 times and compare the empirical probability density function of the number of arrivals with the true probability density function.

We will see further examples of processes with stationary, independent increments in continuous time (and so also examples of continuous-time martingales) in our study of Brownian motion.

Likelihood Ratio Tests

Suppose that (S, \mathcal{S}, μ) is a general measure space, and that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a sequence of independent, identically distributed random variables, taking values in S . In statistical terms, \mathbf{X} corresponds to *sampling* from the common distribution, which is usually not completely known. Indeed, the central problem in statistics is to draw inferences about the distribution from observations of \mathbf{X} . Suppose now that the underlying distribution either has probability density function g_0 or probability density function g_1 , with respect to μ . We assume that g_0 and g_1 are positive on S . Of course the common special cases of this setup are

- S is a measurable subset of \mathbb{R}^n for some $n \in \mathbb{N}_+$ and $\mu = \lambda_n$ is n -dimensional Lebesgue measure on S .
- S is a countable set and $\mu = \#$ is counting measure on S .

The likelihood ratio test is a hypothesis test, where the null and alternative hypotheses are

- H_0 : the probability density function is g_0 .
- H_1 : the probability density function is g_1 .

The test is based on the test statistic

$$L_n = \prod_{i=1}^n \frac{g_0(X_i)}{g_1(X_i)}, \quad n \in \mathbb{N} \quad (17.1.42)$$

known as the *likelihood ratio test statistic*. Small values of the test statistic are evidence in favor of the alternative hypothesis H_1 . Here is our result.

Under the alternative hypothesis H_1 , the process $\mathbf{L} = \{L_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} , known as the *likelihood ratio martingale*.

Proof

Let $\mathcal{F}_n = \sigma\{X_1, X_2, \dots, X_n\}$. For $n \in \mathbb{N}$,

$$\mathbb{E}(L_{n+1} \mid \mathcal{F}_n) = \mathbb{E}\left[L_n \frac{g_0(X_{n+1})}{g_1(X_{n+1})} \mid \mathcal{F}_n\right] = L_n \mathbb{E}\left[\frac{g_0(X_{n+1})}{g_1(X_{n+1})}\right] \quad (17.1.43)$$

Since L_n is measurable with respect to \mathcal{F}_n and $g_0(X_{n+1})/g_1(X_{n+1})$ is independent of \mathcal{F}_n . But under H_1 , and using the change of variables formula for expected value, we have

$$\mathbb{E}\left[\frac{g_0(X_{n+1})}{g_1(X_{n+1})}\right] = \int_S \frac{g_0(x)}{g_1(x)} g_1(x) d\mu(x) = \int_S g_0(x) d\mu(x) = 1 \quad (17.1.44)$$

This result also follows essentially from the theorem [above](#) on partial products. The sequence $\mathbf{Z} = (Z_1, Z_2, \dots)$ given by $Z_i = g_0(X_i)/g_1(X_i)$ for $i \in \mathbb{N}_+$ is independent and identically distributed, and as just shown, has mean 1 under H_1 .

Branching Processes

In the simplest model of a *branching process*, we have a system of *particles* each of which can die out or split into new particles of the same type. The fundamental assumption is that the particles act independently, each with the same offspring distribution on \mathbb{N} . We will let f denote the (discrete) probability density function of the number of offspring of a particle, m the mean of the distribution, and ϕ the probability generating function of the distribution. Thus, if U is the number of children of a particle, then $f(n) = \mathbb{P}(U = n)$ for $n \in \mathbb{N}$, $m = \mathbb{E}(U)$, and $\phi(t) = \mathbb{E}(t^U)$ defined at least for $t \in (-1, 1]$.

Our interest is in *generational time* rather than *absolute time*: the original particles are in generation 0, and recursively, the children a particle in generation n belong to generation $n+1$. Thus, the stochastic process of interest is $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ where X_n is the

number of particles in the n th generation for $n \in \mathbb{N}$. The process \mathbf{X} is a Markov chain and was studied in the section on discrete-time branching chains. In particular, one of the fundamental problems is to compute the probability q of extinction starting with a single particle:

$$q = \mathbb{P}(X_n = 0 \text{ for some } n \in \mathbb{N} \mid X_0 = 1) \quad (17.1.45)$$

Then, since the particles act independently, the probability of extinction starting with $x \in \mathbb{N}$ particles is simply q^x . We will assume that $f(0) > 0$ and $f(0) + f(1) < 1$. This is the interesting case, since it means that a particle has a positive probability of dying without children and a positive probability of producing more than 1 child. The fundamental result, you may recall, is that q is the smallest fixed point of ϕ (so that $\phi(q) = q$) in the interval $[0, 1]$. Here are two martingales associated with the branching process:

Each of the following is a martingale with respect to \mathbf{X} .

1. $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ where $Y_n = X_n/m^n$ for $n \in \mathbb{N}$.
2. $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ where $Z_n = q^{X_n}$ for $n \in \mathbb{N}$.

Proof

Let $\mathcal{F}_n = \sigma\{X_0, X_1, \dots, X_n\}$. For $n \in \mathbb{N}$, note that X_{n+1} can be written in the form

$$X_{n+1} = \sum_{i=1}^{X_n} U_i \quad (17.1.46)$$

where $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent variables, each with PDF f (and hence mean μ and PGF ϕ), and with \mathbf{U} independent of \mathcal{F}_n . Think of U_i as the number of children of the i th particle in generation n .

1. For $n \in \mathbb{N}$,

$$\mathbb{E}(Y_{n+1} \mid \mathcal{F}_n) = \mathbb{E}\left(\frac{X_{n+1}}{m^{n+1}} \mid \mathcal{F}_n\right) = \frac{1}{m^{n+1}} \mathbb{E}\left(\sum_{i=1}^{X_n} U_i \mid \mathcal{F}_n\right) = \frac{1}{m^{n+1}} m X_n = \frac{X_n}{m^n} = Y_n \quad (17.1.47)$$

2. For $n \in \mathbb{N}$

$$\mathbb{E}(Z_{n+1} \mid \mathcal{F}_n) = \mathbb{E}(q^{X_{n+1}} \mid \mathcal{F}_n) = \mathbb{E}\left(q^{\sum_{i=1}^{X_n} U_i} \mid \mathcal{F}_n\right) = [\phi(q)]^{X_n} = q^{X_n} = Z_n \quad (17.1.48)$$

Doob's Martingale

Our next example is one of the simplest, but most important. Indeed, as we will see later in the section on convergence, this type of martingale is almost universal in the sense that every uniformly integrable martingale is of this type. The process is constructed by conditioning a fixed random variable on the σ -algebras in a given filtration, and thus *accumulating information* about the random variable.

Suppose that $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ is a filtration on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and that X is a real-valued random variable with $\mathbb{E}(|X|) < \infty$. Define $X_t = \mathbb{E}(X \mid \mathcal{F}_t)$ for $t \in T$. Then $\mathbf{X} = \{X_t : t \in T\}$ is a martingale with respect to \mathfrak{F} .

Proof

For $t \in T$, recall that $|X_t| = |\mathbb{E}(X \mid \mathcal{F}_t)| \leq \mathbb{E}(|X| \mid \mathcal{F}_t)$. Taking expected values gives $\mathbb{E}(|X_t|) \leq \mathbb{E}(|X|) < \infty$. Suppose that $s, t \in T$ with $s < t$. Using the tower property of conditional expected value,

$$\mathbb{E}(X_t \mid \mathcal{F}_s) = \mathbb{E}[\mathbb{E}(X \mid \mathcal{F}_t) \mid \mathcal{F}_s] = \mathbb{E}(X \mid \mathcal{F}_s) = X_s \quad (17.1.49)$$

The martingale in the last theorem is known as *Doob's martingale* and is named for Joseph Doob who did much of the pioneering work on martingales. It's also known as the *Lévy martingale*, named for Paul Lévy.

Doob's martingale arises naturally in the statistical context of Bayesian estimation. Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of independent random variables whose common distribution depends on an unknown real-valued parameter θ , with values in a parameter space $A \subseteq \mathbb{R}$. For each $n \in \mathbb{N}_+$, let $\mathcal{F}_n = \sigma\{X_1, X_2, \dots, X_n\}$ so that $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}_+\}$ is the natural filtration associated with \mathbf{X} . In Bayesian estimation, we model the unknown parameter θ with a random variable Θ taking values in A and having a specified *prior distribution*. The *Bayesian estimator* of θ based on the sample $\mathbf{X}_n = (X_1, X_2, \dots, X_n)$ is

$$U_n = \mathbb{E}(\Theta \mid \mathcal{F}_n), \quad n \in \mathbb{N}_+ \quad (17.1.50)$$

So it follows that the sequence of Bayesian estimators $\mathbf{U} = (U_n : n \in \mathbb{N}_+)$ is a Doob martingale. The estimation referred to in the discussion of the [beta-Bernoulli](#) process above is a special case.

Density Functions

For this example, you may need to review general measures and density functions in the chapter on Distributions. We start with our probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ in discrete time. Suppose now that μ is a finite measure on the sample space (Ω, \mathcal{F}) . For each $n \in \mathbb{N}$, the restriction of μ to \mathcal{F}_n is a measure on the measurable space (Ω, \mathcal{F}_n) , and similarly the restriction of \mathbb{P} to \mathcal{F}_n is a probability measure on (Ω, \mathcal{F}_n) . To save notation and terminology, we will refer to these as μ and \mathbb{P} on \mathcal{F}_n , respectively. Suppose now that μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F}_n for each $n \in \mathbb{N}$. Recall that this means that if $A \in \mathcal{F}_n$ and $\mathbb{P}(A) = 0$ then $\mu(A) = 0$ for every $B \in \mathcal{F}_n$ with $B \subseteq A$. By the Radon-Nikodym theorem, μ has a density function $X_n : \Omega \rightarrow \mathbb{R}$ with respect to \mathbb{P} on \mathcal{F}_n for each $n \in \mathbb{N}_+$. The density function of a measure with respect to a positive measure is known as a *Radon-Nikodym derivative*. The theorem and the derivative are named for Johann Radon and Otto Nikodym. Here is our main result.

$\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathfrak{F} .

Proof

Let $n \in \mathbb{N}$. By definition, X_n is measurable with respect to \mathcal{F}_n . Also, $\mathbb{E}(|X_n|) = \|\mu\|$ (the total variation of μ) for each $n \in \mathbb{N}$. Since μ is a finite measure, $\|\mu\| < \infty$. By definition,

$$\mu(A) = \int_A X_n d\mathbb{P} = \mathbb{E}(X_n; A), \quad A \in \mathcal{F}_n \quad (17.1.51)$$

On the other hand, if $A \in \mathcal{F}_n$ then $A \in \mathcal{F}_{n+1}$ and so $\mu(A) = \mathbb{E}(X_{n+1}; A)$. So to summarize, X_n is \mathcal{F}_n -measurable and $\mathbb{E}(X_{n+1}; A) = \mathbb{E}(X_n; A)$ for all $A \in \mathcal{F}_n$. By definition, this means that $\mathbb{E}(X_{n+1} | \mathcal{F}_n) = X_n$, and so \mathbf{X} is a martingale with respect to \mathfrak{F} .

Note that μ may not be absolutely continuous with respect to \mathbb{P} on \mathcal{F} or even on $\mathcal{F}_\infty = \sigma(\bigcup_{n=0}^\infty \mathcal{F}_n)$. On the other hand, if μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F}_∞ then μ has a density function X with respect to \mathbb{P} on \mathcal{F}_∞ . So a natural question in this case is the relationship between the martingale \mathbf{X} and the random variable X . You may have already guessed the answer, but at any rate it will be given in the section on convergence.

This page titled [17.1: Introduction to Martingales](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

17.2: Properties and Constructions

Basic Theory

Preliminaries

As in the Introduction, we start with a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, having state space \mathbb{R} , and where the index set T (representing time) is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). Next, we have a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, and we assume that \mathbf{X} is *adapted* to \mathfrak{F} . So \mathfrak{F} is an increasing family of sub σ -algebras of \mathcal{F} and X_t is measurable with respect to \mathcal{F}_t for $t \in T$. We think of \mathcal{F}_t as the collection of events up to time $t \in T$. We assume that $\mathbb{E}(|X_t|) < \infty$, so that the mean of X_t exists as a real number, for each $t \in T$. Finally, in continuous time where $T = [0, \infty)$, we make the standard assumptions that \mathbf{X} is right continuous and has left limits, and that the filtration \mathfrak{F} is right continuous and complete. Please recall the following from the Introduction:

Definitions

1. \mathbf{X} is a *martingale* with respect to \mathfrak{F} if $\mathbb{E}(X_t | \mathcal{F}_s) = X_s$ for all $s, t \in T$ with $s \leq t$.
2. \mathbf{X} is a *sub-martingale* with respect to \mathfrak{F} if $\mathbb{E}(X_t | \mathcal{F}_s) \geq X_s$ for all $s, t \in T$ with $s \leq t$.
3. \mathbf{X} is a *super-martingale* with respect to \mathfrak{F} if $\mathbb{E}(X_t | \mathcal{F}_s) \leq X_s$ for all $s, t \in T$ with $s \leq t$.

Our goal in this section is to give a number of basic properties of martingales and to give ways of constructing martingales from other types of processes. The deeper, fundamental theorems will be studied in the following sections.

Basic Properties

Our first result is that the martingale property is preserved under a coarser filtration.

Suppose that the process \mathbf{X} and the filtration \mathfrak{F} satisfy the basic assumptions [above](#) and that \mathfrak{G} is a filtration coarser than \mathfrak{F} so that $\mathcal{G}_t \subseteq \mathcal{F}_t$ for $t \in T$. If \mathbf{X} is a martingale (sub-martingale, super-martingale) with respect to \mathfrak{F} and \mathbf{X} is adapted to \mathfrak{G} then \mathbf{X} is a martingale (sub-martingale, super-martingale) with respect to \mathfrak{G} .

Proof

Suppose that $s, t \in T$ with $s \leq t$. The proof uses the tower and increasing properties of conditional expected value, and the fact that \mathbf{X} is adapted to \mathfrak{G}

1. If \mathbf{X} is a martingale with respect to \mathfrak{F} then

$$\mathbb{E}(X_t | \mathcal{G}_s) = \mathbb{E}[\mathbb{E}(X_t | \mathcal{F}_s) | \mathcal{G}_s] = \mathbb{E}(X_s | \mathcal{G}_s) = X_s \quad (17.2.1)$$

2. If \mathbf{X} is a sub-martingale with respect to \mathfrak{F} then

$$\mathbb{E}(X_t | \mathcal{G}_s) = \mathbb{E}[\mathbb{E}(X_t | \mathcal{F}_s) | \mathcal{G}_s] \geq \mathbb{E}(X_s | \mathcal{G}_s) = X_s \quad (17.2.2)$$

3. If \mathbf{X} is a super-martingale with respect to \mathfrak{F} then

$$\mathbb{E}(X_t | \mathcal{G}_s) = \mathbb{E}[\mathbb{E}(X_t | \mathcal{F}_s) | \mathcal{G}_s] \leq \mathbb{E}(X_s | \mathcal{G}_s) = X_s \quad (17.2.3)$$

In particular, if \mathbf{X} is a martingale (sub-martingale, super-martingale) with respect to *some* filtration, then it is a martingale (sub-martingale, super-martingale) with respect to its own natural filtration.

The relations that define martingales, sub-martingales, and super-martingales hold for the ordinary (unconditional) expected values. We had this result in the last section, but it's worth repeating.

Suppose that $s, t \in T$ with $s \leq t$.

1. If \mathbf{X} is a martingale with respect to \mathfrak{F} then $\mathbb{E}(X_s) = \mathbb{E}(X_t)$.
2. If \mathbf{X} is a sub-martingale with respect to \mathfrak{F} then $\mathbb{E}(X_s) \leq \mathbb{E}(X_t)$.
3. If \mathbf{X} is a super-martingale with respect to \mathfrak{F} then $\mathbb{E}(X_s) \geq \mathbb{E}(X_t)$.

Proof

The results follow directly from the definitions, and the critical fact that $\mathbb{E}[\mathbb{E}(X_t | \mathcal{F}_s)] = \mathbb{E}(X_t)$ for $s, t \in T$.

So if \mathbf{X} is a martingale then \mathbf{X} has constant expected value, and this value is referred to as the *mean* of \mathbf{X} . The martingale properties are preserved under sums of the stochastic processes.

For the processes $\mathbf{X} = \{X_t : t \in T\}$ and $\mathbf{Y} = \{Y_t : t \in T\}$, let $\mathbf{X} + \mathbf{Y} = \{X_t + Y_t : t \in T\}$. If \mathbf{X} and \mathbf{Y} are martingales (sub-martingales, super-martingales) with respect to \mathcal{F} then $\mathbf{X} + \mathbf{Y}$ is a martingale (sub-martingale, super-martingale) with respect to \mathcal{F} .

Proof

The results follow easily from basic properties of expected value and conditional expected value. First note that $\mathbb{E}(|X_t + Y_t|) \leq \mathbb{E}(|X_t|) + \mathbb{E}(|Y_t|) < \infty$ for $t \in T$. Next $\mathbb{E}(X_t + Y_t | \mathcal{F}_s) = \mathbb{E}(X_t | \mathcal{F}_s) + \mathbb{E}(Y_t | \mathcal{F}_s)$ for $s, t \in T$ with $s \leq t$.

The sub-martingale and super-martingale properties are preserved under multiplication by a positive constant and are reversed under multiplication by a negative constant.

For the process $\mathbf{X} = \{X_t : t \in T\}$ and the constant $c \in \mathbb{R}$, let $c\mathbf{X} = \{cX_t : t \in T\}$.

1. If \mathbf{X} is a martingale with respect to \mathcal{F} then $c\mathbf{X}$ is also a martingale with respect to \mathcal{F}
2. If \mathbf{X} is a sub-martingale with respect to \mathcal{F} , then $c\mathbf{X}$ is a sub-martingale if $c > 0$, a super-martingale if $c < 0$, and a martingale if $c = 0$.
3. If \mathbf{X} is a super-martingale with respect to \mathcal{F} , then $c\mathbf{X}$ is a super-martingale if $c > 0$, a sub-martingale if $c < 0$, and a martingale if $c = 0$.

Proof

The results follow easily from basic properties of expected value and conditional expected value. First note that $\mathbb{E}(|cX_t|) = |c|\mathbb{E}(|X_t|) < \infty$ for $t \in T$. Next $\mathbb{E}(cX_t | \mathcal{F}_s) = c\mathbb{E}(X_t | \mathcal{F}_s)$ for $s, t \in T$ with $s \leq t$.

Property (a), together with the previous additive property, means that the collection of martingales with respect to a fixed filtration \mathcal{F} forms a vector space. Here is a class of transformations that turns martingales into sub-martingales.

Suppose that \mathbf{X} takes values in an interval $S \subseteq \mathbb{R}$ and that $g : S \rightarrow \mathbb{R}$ is convex with $\mathbb{E}[|g(X_t)|] < \infty$ for $t \in T$. If either of the following conditions holds then $g(\mathbf{X}) = \{g(X_t) : t \in T\}$ is a sub-martingale with respect to \mathcal{F} :

1. \mathbf{X} is a martingale.
2. \mathbf{X} is a sub-martingale and g is also increasing.

Proof




Figure 17.2.1: A convex function and several supporting lines

Here is the most important special case of the previous result:

Suppose again that \mathbf{X} is a martingale with respect to \mathcal{F} . Let $k \in [1, \infty)$ and suppose that $\mathbb{E}(|X_t|^k) < \infty$ for $t \in T$. Then the process $|\mathbf{X}|^k = \{|X_t|^k : t \in T\}$ is a sub-martingale with respect to \mathcal{F}

Proof




Figure 17.2.2: The graphs of $x \mapsto |x|$, $x \mapsto |x|^2$ and $x \mapsto |x|^3$ on the interval $[-2, 2]$

In particular, if \mathbf{X} is a martingale relative to \mathfrak{F} then $|\mathbf{X}| = \{ |X_t| : t \in T \}$ is a sub-martingale relative to \mathfrak{F} . Here is a related result that we will need later. First recall that the *positive and negative parts* of $x \in \mathbb{R}$ are $x^+ = \max\{x, 0\}$ and $x^- = \max\{-x, 0\}$, so that $x^+ \geq 0$, $x^- \geq 0$, $x = x^+ - x^-$, and $|x| = x^+ + x^-$.




Figure 17.2.3: The graph of $x \mapsto x^+$ on the interval $[-5, 5]$

If $\mathbf{X} = \{X_t : t \in T\}$ is a sub-martingale relative to $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ then $\mathbf{X}^+ = \{X_t^+ : t \in T\}$ is also a sub-martingale relative to \mathfrak{F} .

Proof

As shown in the graph above, the function $x \mapsto x^+$ is increasing and convex on \mathbb{R} .

Our last result of this discussion is that if we sample a continuous-time martingale at an increasing sequence of time points, we get a discrete-time martingale.

Suppose again that the process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ and the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ satisfy the basic assumptions above. Suppose also that $\{t_n : n \in \mathbb{N}\} \subset [0, \infty)$ is a strictly increasing sequence of time points with $t_0 = 0$, and define $Y_n = X_{t_n}$ and $\mathcal{G}_n = \mathcal{F}_{t_n}$ for $n \in \mathbb{N}$. If \mathbf{X} is a martingale (sub-martingale, super-martingale) with respect to \mathfrak{F} then $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale (sub-martingale, super-martingale) with respect to \mathcal{G} .

Proof

Since the time points are increasing, it's clear that \mathcal{G} is a discrete-time filtration. Next, $\mathbb{E}(|Y_n|) = \mathbb{E}(|X_{t_n}|) < \infty$. Finally, suppose that \mathbf{X} is a martingale and $n, k \in \mathbb{N}$ with $k < n$. Then $t_k < t_n$ so

$$\mathbb{E}(Y_n | \mathcal{G}_k) = \mathbb{E}(X_{t_n} | \mathcal{F}_{t_k}) = X_{t_k} = Y_k \quad (17.2.4)$$

Hence \mathbf{Y} is also a martingale. The proofs for sub and super-martingales are similar, but with inequalities replacing the second equality.

This result is often useful for extending proofs of theorems in discrete time to continuous time.

The Martingale Transform

Our next discussion, in discrete time, shows how to build a new martingale from an existing martingale and an predictable process. This construction turns out to be very useful, and has an interesting gambling interpretation. To review the definition, recall that $\{Y_n : n \in \mathbb{N}_+\}$ is *predictable* relative to the filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ if Y_n is measurable with respect to \mathcal{F}_{n-1} for $n \in \mathbb{N}_+$. Think of Y_n as the bet that a gambler makes on game $n \in \mathbb{N}_+$. The gambler can base the bet on all of the information she has at that point, including the outcomes of the previous $n - 1$ games. That is, she can base the bet on the information encoded in \mathcal{F}_{n-1} .

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is adapted to the filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ and that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}_+\}$ is predictable relative to \mathfrak{F} . The *transform* of \mathbf{X} by \mathbf{Y} is the process $\mathbf{Y} \cdot \mathbf{X}$ defined by

$$(\mathbf{Y} \cdot \mathbf{X})_n = X_0 + \sum_{k=1}^n Y_k (X_k - X_{k-1}), \quad n \in \mathbb{N} \quad (17.2.5)$$

The motivating example behind the transform, in terms of a gambler making a sequence of bets, is given in an [example below](#). Note that $\mathbf{Y} \cdot \mathbf{X}$ is also adapted to \mathfrak{F} . Note also that the transform depends on \mathbf{X} only through X_0 and $\{X_n - X_{n-1} : n \in \mathbb{N}_+\}$. If \mathbf{X} is a martingale, this sequence is the martingale difference sequence studied in Introduction.

Suppose $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is adapted to the filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ and that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a bounded process, predictable relative to \mathfrak{F} .

1. If \mathbf{X} is a martingale relative to \mathfrak{F} then $\mathbf{Y} \cdot \mathbf{X}$ is also a martingale relative to \mathfrak{F} .
2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} and \mathbf{Y} is nonnegative, then $\mathbf{Y} \cdot \mathbf{X}$ is also a sub-martingale relative to \mathfrak{F} .
3. If \mathbf{X} is a super-martingale relative to \mathfrak{F} and \mathbf{Y} is nonnegative, then $\mathbf{Y} \cdot \mathbf{X}$ is also a super-martingale relative to \mathfrak{F} .

Proof

Suppose that $|Y_n| \leq c$ for $n \in \mathbb{N}$ where $c \in (0, \infty)$. Then

$$\mathbb{E}(|(\mathbf{Y} \cdot \mathbf{X})_n|) \leq \mathbb{E}(|X_0|) + c \sum_{k=1}^n [\mathbb{E}(|X_k|) + \mathbb{E}(|X_{k-1}|)] < \infty, \quad n \in \mathbb{N} \quad (17.2.6)$$

Next, for $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{E}[(\mathbf{Y} \cdot \mathbf{X})_{n+1} \mid \mathcal{F}_n] &= \mathbb{E}[(\mathbf{Y} \cdot \mathbf{X})_n + Y_{n+1}(X_{n+1} - X_n) \mid \mathcal{F}_n] = (\mathbf{Y} \cdot \mathbf{X})_n + Y_{n+1} \mathbb{E}(X_{n+1} - X_n \mid \mathcal{F}_n) \\ &= (\mathbf{Y} \cdot \mathbf{X})_n + Y_{n+1} [\mathbb{E}(X_{n+1} \mid \mathcal{F}_n) - X_n] \end{aligned}$$

since $(\mathbf{Y} \cdot \mathbf{X})_n$, Y_{n+1} and X_n are \mathcal{F}_n -measurable. The results now follow from the definitions of martingale, sub-martingale, and super-martingale.

This construction is known as a *martingale transform*, and is a discrete version of the stochastic integral that we will study in the chapter on Brownian motion. The result also holds if instead of \mathbf{Y} being bounded, we have \mathbf{X} bounded and $\mathbb{E}(|Y_n|) < \infty$ for $n \in \mathbb{N}_+$.

The Doob Decomposition

The next result, in discrete time, shows how to decompose a basic stochastic process into a martingale and a predictable process. The result is known as the *Doob decomposition theorem* and is named for Joseph Doob who developed much of the modern theory of martingales.

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ satisfies the [basic assumptions](#) above relative to the filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$. Then $X_n = Y_n + Z_n$ for $n \in \mathbb{N}$ where $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale relative to \mathfrak{F} and $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is predictable relative to \mathfrak{F} . The decomposition is unique.

1. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then \mathbf{Z} is increasing.
2. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then \mathbf{Z} is decreasing.

Proof

Recall that the basic assumptions mean that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is adapted to \mathfrak{F} and $\mathbb{E}(|X_n|) < \infty$ for $n \in \mathbb{N}$. Define $Z_0 = 0$ and

$$Z_n = \sum_{k=1}^n [\mathbb{E}(X_k \mid \mathcal{F}_{k-1}) - X_{k-1}], \quad n \in \mathbb{N}_+ \quad (17.2.7)$$

Then Z_n is measurable with respect to \mathcal{F}_{n-1} for $n \in \mathbb{N}_+$ so \mathbf{Z} is predictable with respect to \mathfrak{F} . Now define

$$Y_n = X_n - Z_n = X_n - \sum_{k=1}^n [\mathbb{E}(X_k \mid \mathcal{F}_{k-1}) - X_{k-1}], \quad n \in \mathbb{N} \quad (17.2.8)$$

Then $\mathbb{E}(|Y_n|) < \infty$ and trivially $X_n = Y_n + Z_n$ for $n \in \mathbb{N}$. Next,

$$\begin{aligned}\mathbb{E}(Y_{n+1} | \mathcal{F}_n) &= \mathbb{E}(X_{n+1} | \mathcal{F}_n) - Z_{n+1} = \mathbb{E}(X_{n+1} | \mathcal{F}_n) - \sum_{k=1}^{n+1} [\mathbb{E}(X_k | \mathcal{F}_{k-1}) - X_{k-1}] \\ &= X_n - \sum_{k=1}^n [\mathbb{E}(X_k | \mathcal{F}_{k-1}) - X_{k-1}] = Y_n, \quad n \in \mathbb{N}\end{aligned}$$

Hence \mathbf{Y} is a martingale. Conversely, suppose that \mathbf{X} has the decomposition in terms of \mathbf{Y} and \mathbf{Z} given in the theorem. Since \mathbf{Y} is a martingale and \mathbf{Z} is predictable,

$$\begin{aligned}\mathbb{E}(X_n - X_{n-1} | \mathcal{F}_{n-1}) &= \mathbb{E}(Y_n | \mathcal{F}_{n-1}) - \mathbb{E}(Y_{n-1} | \mathcal{F}_{n-1}) + \mathbb{E}(Z_n | \mathcal{F}_{n-1}) - \mathbb{E}(Z_{n-1} | \mathcal{F}_{n-1}) \\ &= Y_{n-1} - Y_{n-1} + Z_n - Z_{n-1} = Z_n - Z_{n-1}, \quad n \in \mathbb{N}_+\end{aligned}$$

Also $Z_0 = 0$ so \mathbf{X} uniquely determines \mathbf{Z} . But $Y_n = X_n - Z_n$ for $n \in \mathbb{N}$, so \mathbf{X} uniquely determines \mathbf{Y} also.

1. If \mathbf{X} is a sub-martingale then $\mathbb{E}(X_n | \mathcal{F}_{n-1}) - X_{n-1} \geq 0$ for $n \in \mathbb{N}_+$ so \mathbf{Z} is increasing.
2. If \mathbf{X} is a super-martingale then $\mathbb{E}(X_n | \mathcal{F}_{n-1}) - X_{n-1} \leq 0$ for $n \in \mathbb{N}_+$ so \mathbf{Z} is decreasing.

A decomposition of this form is more complicated in continuous time, in part because the definition of a predictable process is more subtle and complex. The decomposition theorem holds in continuous time, with our basic assumptions and the additional assumption that the collection of random variables $\{X_\tau : \tau \text{ is a finite-valued stopping time}\}$ is uniformly integrable. The result is known as the *Doob-Meyer decomposition theorem*, named additionally for Paul Meyer.

Markov Processes

As you might guess, there are important connections between Markov processes and martingales. Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a (homogeneous) Markov process with state space (S, \mathcal{S}) , relative to the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$. Let $\mathbf{P} = \{P_t : t \in T\}$ denote the collection of transition kernels of \mathbf{X} , so that

$$P_t(x, A) = \mathbb{P}(X_t \in A | X_0 = x), \quad x \in S, A \in \mathcal{S} \quad (17.2.9)$$

Recall that (like all probability kernels), P_t operates (on the right) on (measurable) functions $h : S \rightarrow \mathbb{R}$ by the rule

$$P_t h(x) = \int_S P_t(x, dy) h(y) = \mathbb{E}[h(X_t) | X_0 = x], \quad x \in S \quad (17.2.10)$$

assuming as usual that the expected value exists. Here is the critical definition that we will need.

Suppose that $h : S \rightarrow \mathbb{R}$ and that $\mathbb{E}[|h(X_t)|] < \infty$ for $t \in T$.

1. h is *harmonic* for \mathbf{X} if $P_t h = h$ for $t \in T$.
2. h is *sub-harmonic* for \mathbf{X} if $P_t h \geq h$ for $t \in T$.
3. h is *super-harmonic* for \mathbf{X} if $P_t h \leq h$ for $t \in T$.

The following theorem gives the fundamental connection between the two types of stochastic processes. Given the similarity in the terminology, the result may not be a surprise.

Suppose that $h : S \rightarrow \mathbb{R}$ and $\mathbb{E}[|h(X_t)|] < \infty$ for $t \in T$. Define $h(\mathbf{X}) = \{h(X_t) : t \in T\}$.

1. h is harmonic for \mathbf{X} if and only if $h(\mathbf{X})$ is a martingale with respect to \mathfrak{F} .
2. h is sub-harmonic for \mathbf{X} if and only if $h(\mathbf{X})$ is a sub-martingale with respect to \mathfrak{F} .
3. h is super-harmonic for \mathbf{X} if and only if $h(\mathbf{X})$ is a super-martingale with respect to \mathfrak{F} .

Proof

Suppose that $s, t \in T$ with $s \leq t$. Then by the Markov property,

$$\mathbb{E}[h(X_t) | \mathcal{F}_s] = \mathbb{E}[h(X_t) | X_s] = P_{t-s} h(X_s) \quad (17.2.11)$$

So if h is harmonic, $\mathbb{E}[h(X_t) | \mathcal{F}_s] = h(X_s)$ so $\{h(X_t) : t \in T\}$ is a martingale. Conversely, if $\{h(X_t) : t \in T\}$ is a martingale, then $P_{t-s} h(X_s) = h(X_s)$. Letting $s = 0$ and $X_0 = x$ gives $P_t h(x) = h(x)$ so h is harmonic. The proofs for sub and super-martingales are similar, with inequalities replacing the equalities.

Several of the examples given in the Introduction can be re-interpreted in the context of harmonic functions of Markov chains. We explore some of these below.

Examples

Let \mathcal{R} denote the usual set of Borel measurable subsets of \mathbb{R} , and for $A \in \mathcal{R}$ and $x \in \mathbb{R}$ let $A - x = \{y - x : y \in A\}$. Let I denote the identity function on \mathbb{R} , so that $I(x) = x$ for $x \in \mathbb{R}$. We will need this notation in a couple of our applications below.

Random Walks

Suppose that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is a sequence of independent, real-valued random variables, with $\{V_n : n \in \mathbb{N}_+\}$ identically distributed and having common probability measure Q on $(\mathbb{R}, \mathcal{R})$ and mean $a \in \mathbb{R}$. Recall from the Introduction that the partial sum process $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ associated with \mathbf{V} is given by

$$X_n = \sum_{i=0}^n V_i, \quad n \in \mathbb{N} \quad (17.2.12)$$

and that \mathbf{X} is a (discrete-time) *random walk*. But \mathbf{X} is also a discrete-time Markov process with one-step transition kernel P given by $P(x, A) = Q(A - x)$ for $x \in \mathbb{R}$ and $A \in \mathcal{R}$.

The identity function I is

1. Harmonic for \mathbf{X} if $a = 0$.
2. Sub-harmonic for \mathbf{X} if $a \geq 0$.
3. Super-harmonic for \mathbf{X} if $a \leq 0$.

Proof

Note that

$$PI(x) = \mathbb{E}(X_1 \mid X_0 = x) = x + \mathbb{E}(X_1 - X_0 \mid X_0 = x) = x + \mathbb{E}(V_1 \mid X_0 = x) = I(x) + a \quad (17.2.13)$$

Since V_1 and $X_0 = V_0$ are independent. The results now follow from the definitions.

It now follows from our [theorem above](#) that \mathbf{X} is a martingale if $a = 0$, a sub-martingale if $a > 0$, and a super-martingale if $a < 0$. We showed these results directly from the definitions in the Introduction.

The Simple Random Walk

Suppose now that that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is a sequence of independent random variables with $\mathbb{P}(V_i = 1) = p$ and $\mathbb{P}(V_i = -1) = 1 - p$ for $i \in \mathbb{N}_+$, where $p \in (0, 1)$. Let $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ be the partial sum process associated with \mathbf{V} so that

$$X_n = \sum_{i=0}^n V_i, \quad n \in \mathbb{N} \quad (17.2.14)$$

Then \mathbf{X} is the simple random walk with parameter p . In terms of gambling, our gambler plays a sequence of independent and identical games, and on each game, wins €1 with probability p and loses €1 with probability $1 - p$. So if X_0 is the gambler's initial fortune, then X_n is her net fortune after n games. In the Introduction we showed that \mathbf{X} is a martingale if $p = \frac{1}{2}$, a super-martingale if $p < \frac{1}{2}$, and a sub-martingale if $p > \frac{1}{2}$. But suppose now that instead of making constant unit bets, the gambler makes bets that depend on the outcomes of previous games. This leads to a [martingale transform](#) as studied above.

Suppose that the gambler bets Y_n on game $n \in \mathbb{N}_+$ (at even stakes), where $Y_n \in [0, \infty)$ depends on $(V_0, V_1, V_2, \dots, V_{n-1})$ and satisfies $\mathbb{E}(Y_n) < \infty$. So the process $\mathbf{Y} = \{Y_n : n \in \mathbb{N}_+\}$ is predictable with respect to \mathbf{X} , and the gambler's net winnings after n games is

$$(\mathbf{Y} \cdot \mathbf{X})_n = V_0 + \sum_{k=1}^n Y_k V_k = X_0 + \sum_{k=1}^n Y_k (X_k - X_{k-1}) \quad (17.2.15)$$

1. $\mathbf{Y} \cdot \mathbf{X}$ is a sub-martingale if $p > \frac{1}{2}$.
2. $\mathbf{Y} \cdot \mathbf{X}$ is a super-martingale if $p < \frac{1}{2}$.

3. $\mathbf{Y} \cdot \mathbf{X}$ is a martingale if $p = \frac{1}{2}$.

Proof

These result follow immediately the theorem for [martingale transforms](#) above.

The simple random walk \mathbf{X} is also a discrete-time Markov chain on \mathbb{Z} with one-step transition matrix P given by $P(x, x+1) = p$, $P(x, x-1) = 1-p$.

The function h given by $h(x) = \left(\frac{1-p}{p}\right)^x$ for $x \in \mathbb{Z}$ is harmonic for \mathbf{X} .

Proof

For $x \in \mathbb{Z}$,

$$\begin{aligned} Ph(x) &= ph(x+1) + (1-p)h(x-1) = p\left(\frac{1-p}{p}\right)^{x+1} + (1-p)\left(\frac{1-p}{p}\right)^{x-1} \\ &= \frac{(1-p)^{x+1}}{p^x} + \frac{(1-p)^x}{p^{x-1}} = \left(\frac{1-p}{p}\right)^x [(1-p) + p] = h(x) \end{aligned}$$

It now follows from our [theorem above](#) that the process $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ given by $Z_n = \left(\frac{1-p}{p}\right)^{X_n}$ for $n \in \mathbb{N}$ is a martingale. We showed this directly from the definition in the Introduction. As you may recall, this is *De Moivre's martingale* and named for Abraham De Moivre.

Branching Processes

Recall the discussion of the simple *branching process* from the Introduction. The fundamental assumption is that the particles act independently, each with the same offspring distribution on \mathbb{N} . As before, we will let f denote the (discrete) probability density function of the number of offspring of a particle, m the mean of the distribution, and ϕ the probability generating function of the distribution. We assume that $f(0) > 0$ and $f(0) + f(1) < 1$ so that a particle has a positive probability of dying without children and a positive probability of producing more than 1 child. Recall that q denotes the probability of extinction, starting with a single particle.

The stochastic process of interest is $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ where X_n is the number of particles in the n th generation for $n \in \mathbb{N}$. Recall that \mathbf{X} is a discrete-time Markov chain on \mathbb{N} with one-step transition matrix P given by $P(x, y) = f^{*x}(y)$ for $x, y \in \mathbb{N}$ where f^{*x} denotes the convolution power of order x of f .

The function h given by $h(x) = q^x$ for $x \in \mathbb{N}$ is harmonic for \mathbf{X} .

Proof

For $x \in \mathbb{N}$,

$$Ph(x) = \sum_{y \in \mathbb{N}} P(x, y)h(y) = \sum_{y \in \mathbb{N}} f^{*x}(y)q^y \quad (17.2.16)$$

The last expression is the probability generating function of f^{*x} evaluated at q . But this PGF is simply ϕ^x and q is a fixed point of ϕ so we have

$$Ph(x) = [\phi(q)]^x = q^x = h(x) \quad (17.2.17)$$

It now follows from our [theorem above](#) that the process $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is a martingale where $Z_n = q^{X_n}$ for $n \in \mathbb{N}$. We showed this directly from the definition in the Introduction. We also showed that the process $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale where $Y_n = X_n/m^n$ for $n \in \mathbb{N}$. But we can't write $Y_n = h(X_n)$ for a function h defined on the state space, so we can't interpret this martingale in terms of a harmonic function.

General Random Walks

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a stochastic process satisfying the [basic assumptions](#) above relative to the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$. Recall from the Introduction that the term *increment* refers to a difference of the form $X_{s+t} - X_s$ for $s, t \in T$.

The process \mathbf{X} has *independent increments* if this increment is always independent of \mathcal{F}_s , and has *stationary increments* this increment always has the same distribution as $X_t - X_0$. In discrete time, a process with stationary, independent increments is simply a [random walk](#) as discussed above. In continuous time, a process with stationary, independent increments (and with the continuity assumptions we have imposed) is called a *continuous-time random walk*, and also a *Lévy process*, named for Paul Lévy.

So suppose that \mathbf{X} has stationary, independent increments. For $t \in T$ let Q_t denote the probability distribution of $X_t - X_0$ on $(\mathbb{R}, \mathcal{R})$, so that Q_t is also the probability distribution aof $X_{s+t} - X_s$ for every $s, t \in T$. From our previous study, we know that \mathbf{X} is a Markov processes with transition kernel P_t at time $t \in T$ given by

$$P_t(x, A) = Q_t(A - x); \quad x \in \mathbb{R}, A \in \mathcal{R} \quad (17.2.18)$$

We also know that $\mathbb{E}(X_t - X_0) = at$ for $t \in T$ where $a = \mathbb{E}(X_1 - X_0)$ (assuming of course that the last expected value exists in \mathbb{R}).

The identity function I is .

1. Harmonic for \mathbf{X} if $a = 0$.
2. Sub-harmonic for \mathbf{X} if $a \geq 0$.
3. Super-harmonic for \mathbf{X} if $a \leq 0$.

Proof

Note that

$$P_t I(x) = \mathbb{E}(X_t \mid X_0 = x) = x + \mathbb{E}(X_t - X_0 \mid X_0 = x) = I(x) + at \quad (17.2.19)$$

since $X_t - X_0$ is independent of X_0 . The results now follow from the definitions.

It now follows that \mathbf{X} is a martingale if $a = 0$, a sub-martingale if $a \geq 0$, and a super-martingale if $a \leq 0$. We showed this directly in the Introduction. Recall that in continuous time, the Poisson counting process has stationary, independent increments, as does standard Brownian motion

This page titled [17.2: Properties and Constructions](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

17.3: Stopping Times

Basic Theory

As in the Introduction, we start with a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, having state space \mathbb{R} , and where the index set T (representing time) is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). Next, we have a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, and we assume that \mathbf{X} is adapted to \mathfrak{F} . So \mathfrak{F} is an increasing family of sub σ -algebras of \mathcal{F} and X_t is measurable with respect to \mathcal{F}_t for $t \in T$. We think of \mathcal{F}_t as the collection of events up to time $t \in T$. We assume that $\mathbb{E}(|X_t|) < \infty$, so that the mean of X_t exists as a real number, for each $t \in T$. Finally, in continuous time where $T = [0, \infty)$, we make the standard assumption that \mathbf{X} is right continuous and has left limits, and that the filtration \mathfrak{F} is right continuous and complete.

Our general goal in this section is to see if some of the important martingale properties are preserved if the deterministic time $t \in T$ is replaced by a (random) stopping time. Recall that a random time τ with values in $T \cup \{\infty\}$ is a *stopping time* relative to \mathfrak{F} if $\{\tau \leq t\} \in \mathcal{F}_t$ for $t \in T$. So a stopping time is a random time that does not require that we see into the future. That is, we can tell if $\tau \leq t$ from the information available at time t . Next recall that the σ -algebra associated with the stopping time τ is

$$\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \in T\} \quad (17.3.1)$$

So \mathcal{F}_τ is the collection of events up to the random time τ just as \mathcal{F}_t is the collection of events up to the deterministic time $t \in T$. In terms of a gambler playing a sequence of games, the time that the gambler decides to stop playing must be a stopping time, and in fact this interpretation is the origin of the name. That is, the time when the gambler decides to stop playing can only depend on the information that the gambler has up to that point in time.

Optional Stopping

The basic martingale equation $\mathbb{E}(X_t | \mathcal{F}_s) = X_s$ for $s, t \in T$ with $s \leq t$ can be generalized by replacing both s and t by bounded stopping times. The result is known as the *Doob's optional stopping theorem* and is named again for Joseph Doob. Suppose that $\mathbf{X} = \{X_t : t \in T\}$ satisfies the basic assumptions above with respect to the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$

Suppose that ρ and τ are bounded stopping times relative to \mathfrak{F} with $\rho \leq \tau$.

1. If \mathbf{X} is a martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau | \mathcal{F}_\rho) = X_\rho$.
2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau | \mathcal{F}_\rho) \geq X_\rho$.
3. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau | \mathcal{F}_\rho) \leq X_\rho$.

Proof in discrete time

1. Suppose that $\tau \leq k$ where $k \in \mathbb{N}_+$ and let $A \in \mathcal{F}_\tau$. For $j \in \mathbb{N}$ with $j \leq k$, $A \cap \{\tau = j\} \in \mathcal{F}_j$. Hence by the martingale property,

$$\mathbb{E}(X_k; A \cap \{\tau = j\}) = \mathbb{E}(X_j; A \cap \{\tau = j\}) = \mathbb{E}(X_\tau; A \cap \{\tau = j\}) \quad (17.3.2)$$

Since k is an upper bound on τ , the events $A \cap \{\tau = j\}$ for $j = 0, 1, \dots, k$ partition A , so summing the displayed equation over j gives $\mathbb{E}(X_k; A) = \mathbb{E}(X_\tau; A)$. By definition of conditional expectation, $\mathbb{E}(X_k | \mathcal{F}_\tau) = X_\tau$. But since k is also an upper bound for ρ we also have $\mathbb{E}(X_k | \mathcal{F}_\rho) = X_\rho$. Finally using the tower property we have

$$X_\rho = \mathbb{E}(X_k | \mathcal{F}_\rho) = \mathbb{E}[\mathbb{E}(X_k | \mathcal{F}_\tau) | \mathcal{F}_\rho] = \mathbb{E}[\mathbb{E}(X_\tau | \mathcal{F}_\tau) | \mathcal{F}_\rho] = \mathbb{E}(X_\tau | \mathcal{F}_\rho) \quad (17.3.3)$$

2. If \mathbf{X} is a sub-martingale, then by the Doob decomposition theorem, $X_n = Y_n + Z_n$ for $n \in \mathbb{N}$ where $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is a martingale relative to \mathfrak{F} and $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is increasing and is predictable relative to \mathfrak{F} . So

$$\mathbb{E}(X_\tau | \mathcal{F}_\rho) = \mathbb{E}(Y_\tau | \mathcal{F}_\rho) + \mathbb{E}(Z_\tau | \mathcal{F}_\rho) \quad (17.3.4)$$

But $\mathbb{E}(Y_\tau | \mathcal{F}_\rho) = Y_\rho$ by part (a) and since \mathbf{Z} is increasing, $\mathbb{E}(Z_\tau | \mathcal{F}_\rho) \geq \mathbb{E}(Z_\rho | \mathcal{F}_\rho) = Z_\rho$. Hence $\mathbb{E}(X_\tau | \mathcal{F}_\rho) \geq X_\rho$.

3. The proof when \mathbf{X} is a super-martingale is just like (b), except that the process \mathbf{Z} is decreasing.

Proof in continuous time

Suppose that \mathbf{X} is a martingale. We need to show that $\mathbb{E}(X_\tau; A) = \mathbb{E}(X_\rho; A)$ for every $A \in \mathcal{F}_\rho$. Let $\rho_n = \lceil 2^n \rho \rceil / 2^n$ and $\tau_n = \lceil 2^n \tau \rceil / 2^n$ for $n \in \mathbb{N}$. The stopping times ρ_n and τ_n take values in a countable set T_n for each $n \in \mathbb{N}$, and $\rho_n \downarrow \rho$ and

$\tau_n \downarrow \tau$ as $n \rightarrow \infty$. The process $\{X_t : t \in T_n\}$ is a discrete-time martingale for each $n \in \mathbb{N}$. By the right continuity of \mathbf{X} ,

$$X_{\rho_n} \rightarrow X_\rho, \quad X_{\tau_n} \rightarrow X_\tau \text{ as } n \rightarrow \infty \quad (17.3.5)$$

Suppose next that $\tau \leq c$ where $c \in (0, \infty)$ so that $\rho \leq c$ also. Then $\rho_n \leq c+1$ and $\tau_n \leq c+1$ for $n \in \mathbb{N}$ so the discrete stopping times are uniformly bounded. From the discrete version of the theorem, $X_{\rho_n} = \mathbb{E}(X_{c+1} | \mathcal{F}_{\rho_n})$ and $X_{\tau_n} = \mathbb{E}(X_{c+1} | \mathcal{F}_{\tau_n})$ for $n \in \mathbb{N}$. It then follows that the sequences $\{X_{\rho_n} : n \in \mathbb{N}\}$ and $\{X_{\tau_n} : n \in \mathbb{N}\}$ are uniformly integrable and hence $X_{\rho_n} \rightarrow X_\rho$ and $X_{\tau_n} \rightarrow X_\tau$ as $n \rightarrow \infty$ in mean as well as with probability 1. Now let $A \in \mathcal{F}_\rho$. Since $\rho \leq \rho_n$, $\mathcal{F}_\rho \subseteq \mathcal{F}_{\rho_n}$ and so $A \in \mathcal{F}_{\rho_n}$ for each $n \in \mathbb{N}$. By the theorem in discrete time,

$$\mathbb{E}(X_{\tau_n}; A) = \mathbb{E}(X_{\rho_n}; A), \quad n \in \mathbb{N} \quad (17.3.6)$$

Letting $n \rightarrow \infty$ gives $\mathbb{E}(X_\tau; A) = \mathbb{E}(X_\rho; A)$. The proofs in parts (b) and (c) are as in the discrete time.

The assumption that the stopping times are bounded is critical. A [counterexample](#) when this assumption does not hold is given below. Here are a couple of simple corollaries:

Suppose again that ρ and τ are bounded stopping times relative to \mathfrak{F} with $\rho \leq \tau$.

1. If \mathbf{X} is a martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau) = \mathbb{E}(X_\rho)$.
2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau) \geq \mathbb{E}(X_\rho)$.
3. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau) \leq \mathbb{E}(X_\rho)$.

Proof

Recall that $\mathbb{E}(X_\tau) = \mathbb{E}[\mathbb{E}(X_\tau | \mathcal{F}_\rho)]$, so the results are immediate from the optional stopping theorem.

Suppose that τ is a bounded stopping time relative to \mathfrak{F} .

1. If \mathbf{X} is a martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau) = \mathbb{E}(X_0)$.
2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau) \geq \mathbb{E}(X_0)$.
3. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then $\mathbb{E}(X_\tau) \leq \mathbb{E}(X_0)$.

The Stopped Martingale

For our next discussion, we first need to recall how to stop a stochastic process at a stopping time.

Suppose that \mathbf{X} satisfies the assumptions above and that τ is a stopping time relative to the filtration \mathfrak{F} . The *stopped process* $X^\tau = \{X_t^\tau : t \in [0, \infty)\}$ is defined by

$$X_t^\tau = X_{t \wedge \tau}, \quad t \in [0, \infty) \quad (17.3.7)$$

Details

In continuous time, our standard assumptions ensure that \mathbf{X}^τ is a valid stochastic process and is adapted to \mathfrak{F} . That is, X_t^τ is measurable with respect to \mathcal{F}_t for each $t \in [0, \infty)$. Moreover, \mathbf{X}^τ is also right continuous and has left limits.

So $X_t^\tau = X_t$ if $t < \tau$ and $X_t^\tau = X_\tau$ if $t \geq \tau$. In particular, note that $X_0^\tau = X_0$. If X_t is the fortune of a gambler at time $t \in T$, then X_t^τ is the revised fortune at time t when τ is the stopping time of the gambler. Our next result, known as the *elementary stopping theorem*, is that a martingale stopped at a stopping time is still a martingale.

Suppose again that \mathbf{X} satisfies the assumptions above, and that τ is a stopping time relative to \mathfrak{F} .

1. If \mathbf{X} is a martingale relative to \mathfrak{F} then so is \mathbf{X}^τ .
2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then so is \mathbf{X}^τ .
3. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then so is \mathbf{X}^τ .

General proof

If $s, t \in T$ with $s \leq t$ then $\tau \wedge s$ and $\tau \wedge t$ are bounded stopping times with $\tau \wedge s \leq \tau \wedge t$. So the results follows immediately from the [optional stopping theorem](#) above.

Special proof in discrete time

In discrete time, there is a simple direct proof using the martingale transform. So suppose that $T = \mathbb{N}$ and define the process $\mathbf{Y} = \{Y_n : n \in \mathbb{N}_+\}$ by

$$Y_n = \mathbf{1}(\tau \geq n) = 1 - \mathbf{1}(\tau \leq n-1), \quad n \in \mathbb{N}_+ \quad (17.3.8)$$

By definition of a stopping time, $\{\tau \leq n-1\} \in \mathcal{F}_{n-1}$ for $n \in \mathbb{N}_+$, so the process \mathbf{Y} is predictable. Of course, \mathbf{Y} is a bounded, nonnegative process also. The transform of \mathbf{X} by \mathbf{Y} is

$$(\mathbf{Y} \cdot \mathbf{X})_n = X_0 + \sum_{k=1}^n Y_k (X_k - X_{k-1}) = X_0 + \sum_{k=1}^n \mathbf{1}(\tau \geq k) (X_k - X_{k-1}), \quad n \in \mathbb{N}_+ \quad (17.3.9)$$

But note that $X_k^\tau - X_{k-1}^\tau = X_k - X_{k-1}$ if $\tau \geq k$ and $X_k^\tau - X_{k-1}^\tau = X_\tau - X_\tau = 0$ if $\tau < k$. That is, $X_k^\tau - X_{k-1}^\tau = \mathbf{1}(\tau \geq k) (X_k - X_{k-1})$. Hence

$$(\mathbf{Y} \cdot \mathbf{X})_n = X_0 + \sum_{k=1}^n (X_k^\tau - X_{k-1}^\tau) = X_0 + X_n^\tau - X_0^\tau = X_n^\tau, \quad n \in \mathbb{N}_+ \quad (17.3.10)$$

But if \mathbf{X} is a martingale (sub-martingale) (super-martingale), then so is the transform $\mathbf{Y} \cdot \mathbf{X} = \mathbf{X}_\tau$.

The elementary stopping theorem is bad news for the gambler playing a sequence of games. If the games are fair or unfavorable, then no stopping time, regardless of how cleverly designed, can help the gambler. Since a stopped martingale is still a martingale, the the mean property holds.

Suppose again that \mathbf{X} satisfies the assumptions above, and that τ is a stopping time relative to \mathfrak{F} . Let $t \in T$.

1. If \mathbf{X} is a martingale relative to \mathfrak{F} then $\mathbb{E}(X_{t \wedge \tau}) = \mathbb{E}(X_0)$
2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then $\mathbb{E}(X_{t \wedge \tau}) \geq \mathbb{E}(X_0)$
3. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then $\mathbb{E}(X_{t \wedge \tau}) \leq \mathbb{E}(X_0)$

Optional Stopping in Discrete Time

A simple corollary of the [optional stopping theorem](#) is that if \mathbf{X} is a martingale and τ a bounded stopping time, then $\mathbb{E}(X_\tau) = \mathbb{E}(X_0)$ (with the appropriate inequalities if \mathbf{X} is a sub-martingale or a super-martingale). Our next discussion centers on other conditions which give these results in discrete time. Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ satisfies the basic assumptions above with respect to the filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$, and that τ is a stopping time relative to \mathfrak{F} .

Suppose that $|X_n|$ is bounded uniformly in $n \in \mathbb{N}$ and that τ is finite.

4. If \mathbf{X} is a martingale then $\mathbb{E}(X_\tau) = \mathbb{E}(X_0)$.
5. If \mathbf{X} is a sub-martingale then $\mathbb{E}(X_\tau) \geq \mathbb{E}(X_0)$.
6. If \mathbf{X} is a super-martingale then $\mathbb{E}(X_\tau) \leq \mathbb{E}(X_0)$.

Proof

Assume that \mathbf{X} is a super-martingale. The proof for a sub-martingale are similar, and then the results follow immediately for a martingale. The main tool is the [mean property](#) above for the stopped super-martingale:

$$\mathbb{E}(X_{\tau \wedge n}) \leq \mathbb{E}(X_0), \quad n \in \mathbb{N} \quad (17.3.11)$$

Since $\tau < \infty$ with probability 1, $\tau \wedge n \rightarrow \tau$ as $n \rightarrow \infty$, also with probability 1. Since $|X_n|$ is bounded in $n \in T$, it follows from the bounded convergence theorem that $\mathbb{E}(X_{\tau \wedge n}) \rightarrow \mathbb{E}(X_\tau)$ as $n \rightarrow \infty$. Letting $n \rightarrow \infty$ in the displayed equation gives $\mathbb{E}(X_\tau) \leq \mathbb{E}(X_0)$.

Suppose that $|X_{n+1} - X_n|$ is bounded uniformly in $n \in \mathbb{N}$ and that $\mathbb{E}(\tau) < \infty$.

4. If \mathbf{X} is a martingale then $\mathbb{E}(X_\tau) = \mathbb{E}(X_0)$.
5. If \mathbf{X} is a sub-martingale then $\mathbb{E}(X_\tau) \geq \mathbb{E}(X_0)$.
6. If \mathbf{X} is a super-martingale then $\mathbb{E}(X_\tau) \leq \mathbb{E}(X_0)$.

Proof

Assume that \mathbf{X} is a super-martingale. The proofs for a sub-martingale are similar, and then the results follow immediately for a martingale. The main tool once again is the [mean property](#) above for the stopped super-martingale:

$$\mathbb{E}(X_{\tau \wedge n}) \leq \mathbb{E}(X_0), \quad n \in \mathbb{N} \quad (17.3.12)$$

Suppose that $|X_{n+1} - X_n| \leq c$ where $c \in (0, \infty)$. Then

$$|X_{\tau \wedge n} - X_0| = \left| \sum_{k=1}^{\tau \wedge n} (X_k - X_{k-1}) \right| \leq \sum_{k=1}^{\tau \wedge n} |X_k - X_{k-1}| \leq c(\tau \wedge n) \leq c\tau \quad (17.3.13)$$

Hence $|X_{\tau \wedge n}| \leq c\tau + |X_0|$. Since $\mathbb{E}(\tau) < \infty$ we know that $\tau < \infty$ with probability 1, so as before, $\tau \wedge n \rightarrow \tau$ as $n \rightarrow \infty$. Also $\mathbb{E}(c\tau + |X_0|) < \infty$ so by the dominated convergence theorem, $\mathbb{E}(X_{\tau \wedge n}) \rightarrow \mathbb{E}(X_\tau)$ as $n \rightarrow \infty$. So again letting $n \rightarrow \infty$ in the displayed equation gives $\mathbb{E}(X_\tau) \leq \mathbb{E}(X_0)$.

Let's return to our original interpretation of a martingale \mathbf{X} representing the fortune of a gambler playing fair games. The gambler could choose to quit at a random time τ , but τ would have to be a stopping time, based on the gambler's information encoded in the filtration \mathfrak{F} . Under the conditions of the theorem, no such scheme can help the gambler in terms of expected value.

Examples and Applications

The Simple Random Walk

Suppose that $\mathbf{V} = (V_1, V_2, \dots)$ is a sequence of independent, identically distributed random variables with $\mathbb{P}(V_i = 1) = p$ and $\mathbb{P}(V_i = -1) = 1 - p$ for $i \in \mathbb{N}_+$, where $p \in (0, 1)$. Let $\mathbf{X} = (X_0, X_1, X_2, \dots)$ be the partial sum process associated with \mathbf{V} so that

$$X_n = \sum_{i=1}^n V_i, \quad n \in \mathbb{N} \quad (17.3.14)$$

Then \mathbf{X} is the simple random walk with parameter p . In terms of gambling, our gambler plays a sequence of independent and identical games, and on each game, wins €1 with probability p and loses €1 with probability $1 - p$. So X_n is the gambler's total net winnings after n games. We showed in the Introduction that \mathbf{X} is a martingale if $p = \frac{1}{2}$ (the fair case), a sub-martingale if $p > \frac{1}{2}$ (the favorable case), and a super-martingale if $p < \frac{1}{2}$ (the unfair case). Now, for $c \in \mathbb{Z}$, let

$$\tau_c = \inf\{n \in \mathbb{N} : X_n = c\} \quad (17.3.15)$$

where as usual, $\inf(\emptyset) = \infty$. So τ_c is the first time that the gambler's fortune reaches c . What if the gambler simply continues playing until her net winnings is some specified positive number (say €1 000 000)? Is that a workable strategy?

Suppose that $p = \frac{1}{2}$ and that $c \in \mathbb{N}_+$.

1. $\mathbb{P}(\tau_c < \infty) = 1$
2. $\mathbb{E}(X_{\tau_c}) = c \neq 0 = \mathbb{E}(X_0)$
3. $\mathbb{E}(\tau_c) = \infty$

Proof

Parts (a) and (c) hold since \mathbf{X} is a null recurrent Markov chain. Part (b) follows from (a) since trivially $X_{\tau_c} = c$ if $\tau_c < \infty$.

Note that part (b) does not contradict the [optional stopping theorem](#) because of part (c). The strategy of waiting until the net winnings reaches a specified goal c is unsustainable. Suppose now that the gambler plays until the net winnings either falls to a specified negative number (a loss that she can tolerate) or reaches a specified positive number (a goal she hopes to reach).

Suppose again that $p = \frac{1}{2}$. For $a, b \in \mathbb{N}_+$, let $\tau = \tau_{-a} \wedge \tau_b$. Then

1. $\mathbb{E}(\tau) < \infty$
2. $\mathbb{E}(X_\tau) = 0$
3. $\mathbb{P}(\tau_{-a} < \tau_b) = b/(a+b)$

Proof

1. We will let X_0 have an arbitrary value in the set $\{-a, -a+1, \dots, b-1, b\}$, so that we can use Markov chain techniques. Let $m(x) = \mathbb{E}(\tau \mid X_0 = x)$ for x in this set. Conditioning on the first state and using the Markov property we have

$$m(x) = 1 + \frac{1}{2}m(x-1) + \frac{1}{2}m(x+1), \quad x \in \{-a+1, \dots, b-1\} \quad (17.3.16)$$

with boundary conditions $m(-a) = m(b) = 0$. The linear recurrence relation can be solved explicitly, but all that we care about is the fact that the solution is finite.

2. The [optional sampling theorem](#) applies, so $\mathbb{E}(X_\tau) = \mathbb{E}(X_0) = 0$.
3. Let $q = \mathbb{P}(\tau_{-a} < \tau_b)$ so that $1 - q = \mathbb{P}(\tau_b < \tau_{-a})$. By definition, $X_\tau = -a$ if $\tau_{-a} < \tau_b$ and $X_\tau = b$ if $\tau_b < \tau_{-a}$. So from (b), $q(-a) + (1 - q)b = 0$ and therefore $q = b/(a+b)$.

So gambling until the net winnings either falls to $-a$ or reaches b is a workable strategy, but alas has expected value 0. Here's another example that shows that the first version of the optional sampling theorem can fail if the stopping times are not bounded.

Suppose again that $p = \frac{1}{2}$. Let $a, b \in \mathbb{N}_+$ with $a < b$. Then $\tau_a < \tau_b < \infty$ but

$$b = \mathbb{E}(X_{\tau_b} \mid \mathcal{F}_{\tau_a}) \neq X_{\tau_a} = a \quad (17.3.17)$$

Proof

Since $X_0 = 0$, the process \mathbf{X} must reach a before reaching b . As before, $\tau_b < \infty$ but $\mathbb{E}(\tau_b) = \infty$ since \mathbf{X} is a null recurrent Markov chain.

This result does not contradict the [optional stopping theorem](#) since the stopping times are not bounded.

Wald's Equation

Wald's equation, named for Abraham Wald is a formula for the expected value of the sum of a random number of independent, identically distributed random variables. We have considered this before, in our discussion of conditional expected value and our discussion of random samples, but martingale theory leads to a particularly simple and elegant proof.

Suppose that $\mathbf{X} = (X_n : n \in \mathbb{N}_+)$ is a sequence of independent, identically distributed variables with common mean $\mu \in \mathbb{R}$. If N is a stopping time for \mathbf{X} with $\mathbb{E}(N) < \infty$ then

$$\mathbb{E}\left(\sum_{k=1}^N X_k\right) = \mathbb{E}(N)\mu \quad (17.3.18)$$

Proof

Let \mathcal{F} denote the natural filtration associated with \mathbf{X} . Let $c = \mathbb{E}(|X_n|)$, so that by assumption, $c < \infty$. Finally, let

$$Y_n = \sum_{k=1}^n (X_k - \mu) \quad n \in \mathbb{N}_+ \quad (17.3.19)$$

Then $\mathbf{Y} = (Y_n : n \in \mathbb{N}_+)$ is a martingale relative to \mathcal{F} , with mean 0. Note that

$$\mathbb{E}(|Y_{n+1} - Y_n|) = \mathbb{E}(|X_{n+1} - \mu|) \leq c + |\mu|, \quad n \in \mathbb{N}_+ \quad (17.3.20)$$

Hence a discrete version of the [optional stopping theorem](#) applies and we have $\mathbb{E}(Y_N) = 0$. Therefore

$$0 = \mathbb{E}(Y_N) = \mathbb{E}\left[\sum_{k=1}^N (X_k - \mu)\right] = \mathbb{E}\left(\sum_{k=1}^N X_k - N\mu\right) = \mathbb{E}\left(\sum_{k=1}^N X_k\right) - \mathbb{E}(N)\mu \quad (17.3.21)$$

Patterns in Multinomial Trials

Patterns in multinomial trials were studied in the chapter on Renewal Processes. As is often the case, martingales provide a more elegant solution. Suppose that $\mathbf{L} = (L_1, L_2, \dots)$ is a sequence of independent, identically distributed random variables taking values in a finite set S , so that \mathbf{L} is a sequence of multinomial trials. Let f denote the common probability density function so that for a generic trial variable L , we have $f(a) = \mathbb{P}(L = a)$ for $a \in S$. We assume that all outcomes in S are actually possible, so $f(a) > 0$ for $a \in S$.

In this discussion, we interpret S as an *alphabet*, and we write the sequence of variables in concatenation form, $\mathbf{L} = L_1 L_2 \cdots$ rather than standard sequence form. Thus the sequence is an infinite string of letters from our alphabet S . We are interested in the first occurrence of a particular finite substring of letters (that is, a “word” or “pattern”) in the infinite sequence. The following definition will simplify the notation.

If $\mathbf{a} = a_1 a_2 \cdots a_k$ is a word of length $k \in \mathbb{N}_+$ from the alphabet S , define

$$f(\mathbf{a}) = \prod_{i=1}^k f(a_i) \quad (17.3.22)$$

so $f(\mathbf{a})$ is the probability of k consecutive trials producing word \mathbf{a} .

So, fix a word $\mathbf{a} = a_1 a_2 \cdots a_k$ of length $k \in \mathbb{N}_+$ from the alphabet S , and consider the number of trials $N_{\mathbf{a}}$ until \mathbf{a} is completed. Our goal is compute $\nu(\mathbf{a}) = \mathbb{E}(N_{\mathbf{a}})$. We do this by casting the problem in terms of a sequence of gamblers playing fair games and then using the [optional stopping theorem](#) above. So suppose that if a gambler bets $c \in (0, \infty)$ on a letter $a \in S$ on a trial, then the gambler wins $c/f(a)$ if a occurs on that trial and wins 0 otherwise. The expected value of this bet is

$$f(a) \frac{c}{f(a)} - c = 0 \quad (17.3.23)$$

and so the bet is fair. Consider now a gambler with an initial fortune 1. When she starts playing, she bets 1 on a_1 . If she wins, she bet her entire fortune $1/f(a_1)$ on the next trial on a_2 . She continues in this way: as long as she wins, she bets her entire fortune on the next trial on the next letter of the word, until either she loses or completes the word \mathbf{a} . Finally, we consider a sequence of independent gamblers playing this strategy, with gambler i starting on trial i for each $i \in \mathbb{N}_+$.

For a finite word \mathbf{a} from the alphabet S , $\nu(\mathbf{a})$ is the total winnings by all of the players at time $N_{\mathbf{a}}$.

Proof

Let X_n denote the total fortunes of all of the gamblers after trial $n \in \mathbb{N}_+$. Since all of the bets are fair, $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is a martingale with mean 0. We will show that the conditions in the discrete version of the [optional sampling theorem](#) hold. First, consider disjoint blocks of trials of length k , that is

$$((L_1, L_2, \dots, L_k), (L_{k+1}, L_{k+2}, \dots, L_{2k}), \dots) \quad (17.3.24)$$

Let $M_{\mathbf{a}}$ denote the index of the first such block that forms the letter \mathbf{a} . This variable has the geometric distribution on \mathbb{N}_+ with success parameter $f(\mathbf{a})$ and so in particular, $\mathbb{E}(M_{\mathbf{a}}) = 1/f(\mathbf{a})$. But clearly $N_{\mathbf{a}} \leq kM_{\mathbf{a}}$ so $\nu(\mathbf{a}) < k/f(\mathbf{a}) < \infty$. Next note that all of the gamblers have stopped playing by time N , so clearly $|X_{n+1} - X_n| \leq 1/f(a)$ for $n \in \mathbb{N}_+$. So the optional stopping theorem applies, and hence $\mathbb{E}(X_{N_{\mathbf{a}}}) = 0$. But note that $\nu(\mathbf{a})$ can also be interpreted as the expected amount of money invested by the gamblers (1 unit at each time until the game ends at time $N_{\mathbf{a}}$), and hence this must also be the total winnings at time $N_{\mathbf{a}}$ (which is deterministic).

Given \mathbf{a} , we can compute the total winnings precisely. By definition, trials $N - k + 1, \dots, N$ form the word \mathbf{a} for the first time. Hence for $i \leq N - k$, gambler i loses at some point. Also by definition, gambler $N - k + 1$ wins all of her bets, completes word \mathbf{a} and so collects $1/f(\mathbf{a})$. The complicating factor is that gamblers $N - k + 2, \dots, N$ may or may not have won all of their bets at the point when the game ends. The following exercise illustrates this.

Suppose that \mathbf{L} is a sequence of Bernoulli trials (so $S = \{0, 1\}$) with success probability $p \in (0, 1)$. For each of the following strings, find the expected number of trials needed to complete the string.

1. 001
2. 010

Solution

Let $q = 1 - p$.

1. For the word 001, gambler $N - 2$ wins $\frac{1}{q^2 p}$ on her three bets. Gambler $N - 2$ makes two bets, winning the first but losing the second. Gambler N loses her first (and only) bet. Hence $\nu(001) = \frac{1}{q^2 p}$

- For the word 010, gambler $N - 2$ wins $\frac{1}{q^2 p}$ on her three bets as before. Gambler $N - 1$ loses his first bet. Gambler N wins $1/q$ on his first (and only) bet. So $\nu(010) = \frac{1}{q^2 p} + \frac{1}{q}$

The difference between the two words is that the word in (b) has a *prefix* (a proper string at the beginning of the word) that is also a *suffix* (a proper string at the end of the word). Word \mathbf{a} has no such prefix. Thus we are led naturally to the following dichotomy:

Suppose that \mathbf{a} is a finite word from the alphabet S . If no proper prefix of \mathbf{a} is also a suffix, then \mathbf{a} is *simple*. Otherwise, \mathbf{a} is *compound*.

Here is the main result, which of course is the same as when the problem was solved using renewal theory.

Suppose that \mathbf{a} is a finite word in the alphabet S .

- If \mathbf{a} is simple then $\nu(\mathbf{a}) = 1/f(\mathbf{a})$.
- If \mathbf{a} is compound, then $\nu(\mathbf{a}) = 1/f(\mathbf{a}) + \nu(\mathbf{b})$ where \mathbf{b} is the longest word that is both a prefix and a suffix of \mathbf{a} .

Proof

The ingredients are in place from our previous discussion. Suppose that \mathbf{a} has length $k \in \mathbb{N}_+$.

- If \mathbf{a} is simple, only player $N - k + 1$ wins, and she wins $1/f(\mathbf{a})$.
- Suppose \mathbf{a} is compound and \mathbf{b} is the largest proper prefix-suffix. player $N - k + 1$ wins $1/f(\mathbf{a})$ as always. The winnings of players $N - k + 2, \dots, N$ are the same as the winnings of a new sequence of gamblers playing a new sequence of trials with the goal of reaching word \mathbf{b} .

For a compound word, we can use (b) to reduce the computation to simple words.

Consider Bernoulli trials with success probability $p \in (0, 1)$. Find the expected number of trials until each of the following strings is completed.

- 1011011
- $11 \cdots 1$ (k times)

Solutions

Again, let $q = 1 - p$.

- $\nu(1011011) = \frac{1}{p^5 q^2} + \nu(1011) = \frac{1}{p^5 q^2} + \frac{1}{p^3 q} + \nu(1) = \frac{1}{p^5 q^2} + \frac{1}{p^3 q} + \frac{1}{p}$
- Let $\mathbf{1}_j$ denote a string of j 1s for $j \in \mathbb{N}_+$. If $k \geq 2$ then $\nu(\mathbf{1}_k) = 1/p^k + \nu(\mathbf{1}_{k-1})$. Hence

$$\nu(\mathbf{1}_k) = \sum_{j=1}^k \frac{1}{p^j} \quad (17.3.25)$$

Recall that an *ace-six flat die* is a six-sided die for which faces 1 and 6 have probability $\frac{1}{4}$ each while faces 2, 3, 4, and 5 have probability $\frac{1}{8}$ each. Ace-six flat dice are sometimes used by gamblers to cheat.

Suppose that an ace-six flat die is thrown repeatedly. Find the expected number of throws until the pattern 6165616 occurs.

Solution

From our main theorem,

$$\begin{aligned} \nu(6165616) &= \frac{1}{f(6165616)} + \nu(616) = \frac{1}{f(6165616)} + \frac{1}{f(616)} + \nu(6) \\ &= \frac{1}{f(6165616)} + \frac{1}{f(616)} + \frac{1}{f(6)} = \frac{1}{(1/4)^6(1/8)} + \frac{1}{(1/4)^3} + \frac{1}{1/4} = 32\,836 \end{aligned}$$

Suppose that a monkey types randomly on a keyboard that has the 26 lower-case letter keys and the space key (so 27 keys). Find the expected number of keystrokes until the monkey produces each of the following phrases:

1. it was the best of times
2. to be or not to be

Solution

1. $27^{24} \approx 2.258 \times 10^{34}$
2. $27^5 + 27^{18} \approx 5.815 \times 10^{25}$

The Secretary Problem

The secretary problem was considered in the chapter on Finite Sampling Models. In this discussion we will solve a variation of the problem using martingales. Suppose that there are $n \in \mathbb{N}_+$ candidates for a job, or perhaps potential marriage partners. The candidates arrive sequentially in random order and are interviewed. We measure the quality of each candidate by a number in the interval $[0, 1]$. Our goal is to select the very best candidate, but once a candidate is rejected, she cannot be recalled. Mathematically, our assumptions are that the sequence of candidate variables $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is independent and that each is uniformly distributed on the interval $[0, 1]$ (and so has the *standard uniform distribution*). Our goal is to select a stopping time τ with respect to \mathbf{X} that maximizes $\mathbb{E}(X_\tau)$, the expected value of the chosen candidate. The following sequence will play a critical role as a sequence of thresholds.

Define the sequence $\mathbf{a} = (a_k : k \in \mathbb{N})$ by $a_0 = 0$ and $a_{k+1} = \frac{1}{2}(1 + a_k^2)$ for $k \in \mathbb{N}$. Then

1. $a_k < 1$ for $k \in \mathbb{N}$.
2. $a_k < a_{k+1}$ for $k \in \mathbb{N}$.
3. $a_k \rightarrow 1$ as $k \rightarrow \infty$.
4. If X is uniformly distributed on $[0, 1]$ then $\mathbb{E}(X \vee a_k) = a_{k+1}$ for $k \in \mathbb{N}$.

Proof

1. Note that $a_1 = \frac{1}{2} < 1$. Suppose that $a_k < 1$ for some $k \in \mathbb{N}_+$. Then $a_{k+1} = \frac{1}{2}(1 + a_k^2) < \frac{1}{2}(1 + 1) = 1$.
2. Note that $0 = a_0 < a_1 = \frac{1}{2}$. Suppose that $a_k > a_{k-1}$ for some $k \in \mathbb{N}_+$. Then $a_{k+1} = \frac{1}{2}(1 + a_k^2) > \frac{1}{2}(1 + a_{k-1}^2) = a_k$.
3. Since the sequence is increasing and bounded above, $a_\infty = \lim_{k \rightarrow \infty} a_k$ exists. Taking limits in the recursion relation gives $a_\infty = \frac{1}{2}(1 + a_\infty^2)$ or equivalently $(a_\infty - 1)^2 = 0$.
4. For $k \in \mathbb{N}$,

$$\mathbb{E}(X \vee a_k) = \int_0^1 (x \vee a_k) dx = \int_0^{a_k} a_k dx + \int_{a_k}^1 x dx = \frac{1}{2}(1 + a_k^2) = a_{k+1} \quad (17.3.26)$$

Since $a_0 = 0$, all of the terms of the sequence are in $[0, 1]$ by (a). Approximations of the first 10 terms are

$$(0, 0.5, 0.625, 0.695, 0.742, 0.775, 0.800, 0.820, 0.836, 0.850, 0.861, \dots) \quad (17.3.27)$$

Property (d) gives some indication of why the sequence is important for the secretary problem. At any rate, the next theorem gives the solution. To simplify the notation, let $\mathbb{N}_n = \{0, 1, \dots, n\}$ and $\mathbb{N}_n^+ = \{1, 2, \dots, n\}$.

The stopping time $\tau = \inf \{k \in \mathbb{N}_n^+ : X_k > a_{n-k}\}$ is optimal for the secretary problem with n candidates. The optimal value is $\mathbb{E}(X_\tau) = a_n$.

Proof

Let $\mathcal{F} = \{\mathcal{F}_k : k \in \mathbb{N}_n^+\}$ be the natural filtration of \mathbf{X} , and suppose that ρ is a stopping time for \mathcal{F} . Define $\mathbf{Y} = \{Y_k : k \in \mathbb{N}_n\}$ by $Y_0 = 0$ and $Y_k = X_{\rho \wedge k} \vee a_{n-k}$ for $k \in \mathbb{N}_n^+$. We will show that \mathbf{Y} is a super-martingale with respect to \mathcal{F} . First, on the event $\rho \leq k-1$,

$$\mathbb{E}(Y_k | \mathcal{F}_{k-1}) = \mathbb{E}[(X_\rho \vee a_{n-k}) | \mathcal{F}_{k-1}] = X_\rho \vee a_{n-k} \leq X_\rho \vee a_{n-k+1} = Y_{k-1} \quad (17.3.28)$$

where we have used the fact that $X_\rho \mathbf{1}(\rho \leq k-1)$ is measurable with respect to \mathcal{F}_{k-1} and the fact that the sequence \mathbf{a} is increasing. On the event $\rho > k-1$,

$$\mathbb{E}(Y_k | \mathcal{F}_{k-1}) = \mathbb{E}(X_k \vee a_{n-k} | \mathcal{F}_{k-1}) = \mathbb{E}(X_k \vee a_{n-k}) = a_{n-k+1} \leq Y_{k-1} \quad (17.3.29)$$

where we have used the fact that X_k and \mathcal{F}_{k-1} are independent, and part (d) of the [previous result](#). Since \mathbf{Y} is a super-martingale and ρ is bounded, the optional stopping theorem applies and we have

$$\mathbb{E}(X_\rho) \leq \mathbb{E}(X_\rho \vee a_{n-\rho}) = \mathbb{E}(Y_\rho) \leq \mathbb{E}(Y_0) = a_n \quad (17.3.30)$$

so a_n is an upper bound on the expected value of the candidate chosen by the stopping time ρ .

Next, we will show that in the special case that $\rho = \tau$, the process \mathbf{Y} is a martingale. On the event $\tau \leq k-1$ we have $\mathbb{E}(Y_k | \mathcal{F}_{k-1}) = X_\tau \vee a_{n-k}$ as before. But by definition, $X_\tau \geq a_{n-\tau} \geq a_{n-k+1} \geq a_{n-k}$ so on this event,

$$\mathbb{E}(Y_k | \mathcal{F}_{k-1}) = X_\tau = X_\tau \vee a_{n-k+1} = Y_{k-1} \quad (17.3.31)$$

On the event $\tau > k-1$ we have $\mathbb{E}(Y_k | \mathcal{F}_{k-1}) = a_{n-k+1}$ as before. But on this event, $Y_{k-1} = a_{n-k+1}$. Now since \mathbf{Y} is a martingale and τ is bounded, the optional stopping theorem applies and we have

$$\mathbb{E}(X_\tau) = \mathbb{E}(X_\tau \vee a_{n-\tau}) = \mathbb{E}(Y_\tau) = \mathbb{E}(Y_0) = a_n \quad (17.3.32)$$

Here is a specific example:

For $n = 5$, the decision rule is as follows:

1. Select candidate 1 if $X_1 > 0.742$; otherwise,
2. select candidate 2 if $X_2 > 0.695$; otherwise,
3. select candidate 3 if $X_3 > 0.625$; otherwise,
4. select candidate 4 if $X_4 > 0.5$; otherwise,
5. select candidate 5.

The expected value of our chosen candidate is 0.775.

In our original version of the secretary problem, we could only observe the *relative ranks* of the candidates, and our goal was to maximize the probability of picking the best candidate. With $n = 5$, the optimal strategy is to let the first two candidates go by and then pick the first candidate after that is better than all previous candidates, if she exists. If she does not exist, of course, we must select candidate 5. The probability of picking the best candidate is 0.433

This page titled 17.3: Stopping Times is shared under a CC BY 2.0 license and was authored, remixed, and/or curated by Kyle Siegrist (Random Services) via source content that was edited to the style and standards of the LibreTexts platform.

17.4: Inequalities

Basic Theory

In this section, we will study a number of interesting inequalities associated with martingales and their sub-martingale and super-martingale cousins. These turn out to be very important for both theoretical reasons and for applications. You may need to review infimums and supremums.

Basic Assumptions

As in the Introduction, we start with a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, having state space \mathbb{R} , and where the index set T (representing time) is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). Next, we have a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, and we assume that \mathbf{X} is *adapted* to \mathfrak{F} . So \mathfrak{F} is an increasing family of sub σ -algebras of \mathcal{F} and X_t is measurable with respect to \mathcal{F}_t for $t \in T$. We think of \mathcal{F}_t as the collection of events up to time $t \in T$. We assume that $\mathbb{E}(|X_t|) < \infty$, so that the mean of X_t exists as a real number, for each $t \in T$. Finally, in continuous time where $T = [0, \infty)$, we make the standard assumptions that $t \mapsto X_t$ is right continuous and has left limits, and that the filtration \mathfrak{F} is right continuous and complete.

Maximal Inequalities

For motivation, let's review a modified version of Markov's inequality, named for Andrei Markov.

If X is a real-valued random variable then

$$\mathbb{P}(X \geq x) \leq \frac{1}{x} \mathbb{E}(X; X \geq x), \quad x \in (0, \infty) \quad (17.4.1)$$

Proof

The modified version has essentially the same elegant proof as the original. Clearly

$$x \mathbf{1}(X \geq x) \leq X \mathbf{1}(X \geq x), \quad x \in (0, \infty) \quad (17.4.2)$$

Taking expected values through the inequality gives $x \mathbb{P}(X \geq x) \leq \mathbb{E}(X; X \geq x)$. Dividing both sides by x gives the result (and it is at this point that we need $x > 0$).

So Markov's inequality gives an upper bound on the probability that X exceeds a given positive value x , in terms of a moment of X . Now let's return to our stochastic process $\mathbf{X} = \{X_t : t \in T\}$. To simplify the notation, let $T_t = \{s \in T : s \leq t\}$ for $t \in T$. Here is the main definition:

For the process \mathbf{X} , define the corresponding *maximal process* $\mathbf{U} = \{U_t : t \in T\}$ by

$$U_t = \sup\{X_s : s \in T_t\}, \quad t \in T \quad (17.4.3)$$

Clearly, the maximal process is increasing, so that if $s, t \in T$ with $s \leq t$ then $U_s \leq U_t$. A trivial application of Markov's inequality above would give

$$\mathbb{P}(U_t \geq x) \leq \frac{1}{x} \mathbb{E}(U_t; U_t \geq x), \quad x > 0 \quad (17.4.4)$$

But when \mathbf{X} is a sub-martingale, the following theorem gives a much stronger result by replacing the first occurrence of U_t on the right with X_t . The theorem is known as *Doob's sub-martingale maximal inequality* (or more simply as *Doob's inequality*), named once again for Joseph Doob who did much of the pioneering work on martingales. A sub-martingale has an *increasing property* of sorts in the sense that if $s, t \in T$ with $s \leq t$ then $\mathbb{E}(X_t | \mathcal{F}_s) \geq X_s$, so it's perhaps not entirely surprising that such a bound is possible.

Suppose that \mathbf{X} is a sub-martingale. For $t \in T$, let $U_t = \sup\{X_s : s \in T_t\}$. Then

$$\mathbb{P}(U_t \geq x) \leq \frac{1}{x} \mathbb{E}(X_t; U_t \geq x), \quad x \in (0, \infty) \quad (17.4.5)$$

Proof in the discrete time

So $T = \mathbb{N}$ and the maximal process is given by $U_n = \max\{X_k : k \in \mathbb{N}_n\}$ for $n \in \mathbb{N}$. Let $x \in (0, \infty)$, and define $\tau_x = \min\{k \in \mathbb{N} : X_k \geq x\}$ where as usual, $\min(\emptyset) = \infty$. The random time τ_x is a stopping time relative to \mathfrak{F} . Moreover, the processes $\{U_n : n \in \mathbb{N}\}$ and $\{\tau_x : x \in (0, \infty)\}$ are inverses in the sense that for $n \in \mathbb{N}$ and $x \in (0, \infty)$,

$$U_n \geq x \text{ if and only if } \tau_x \leq n \quad (17.4.6)$$

We have seen this type of duality before—in the Poisson process and more generally in renewal processes. Let $n \in \mathbb{N}$. First note that

$$\mathbb{E}(X_{\tau_x \wedge n}) = \mathbb{E}(X_{\tau_x \wedge n}; \tau_x \leq n) + \mathbb{E}(X_{\tau_x \wedge n}; \tau_x > n) \quad (17.4.7)$$

If $\tau_x \leq n$ then $X_{\tau_x \wedge n} = X_{\tau_x} \geq x$. On the other hand if $\tau_x > n$ then $X_{\tau_x \wedge n} = X_n$. So we have

$$\mathbb{E}(X_{\tau_x \wedge n}) \geq x\mathbb{P}(\tau_x \leq n) + \mathbb{E}(X_n; \tau_x > n) = x\mathbb{P}(U_t \geq x) + \mathbb{E}(X_n; \tau_x > n) \quad (17.4.8)$$

Similarly,

$$\mathbb{E}(X_n) = \mathbb{E}(X_n; \tau_x \leq n) + \mathbb{E}(X_n; \tau_x > n) = \mathbb{E}(X_n; U_t \geq x) + \mathbb{E}(X_n; \tau_x > n) \quad (17.4.9)$$

But by the optional stopping theorem, $\mathbb{E}(X_{\tau_x \wedge n}) \leq \mathbb{E}(X_n)$. Hence we have

$$x\mathbb{P}(U_t \geq x) + \mathbb{E}(X_n; \tau_x > n) \leq \mathbb{E}(X_n; U_t \geq x) + \mathbb{E}(X_n; \tau_x > n) \quad (17.4.10)$$

Subtracting the common term and then dividing both sides by x gives the result

Proof in continuous time

For $k \in \mathbb{N}$, let $\mathbb{D}_k^+ = \{j/2^k : j \in \mathbb{N}\}$ denote the set of nonnegative dyadic rationals (or *binary rationals*) of rank k or less. For $t \in [0, \infty)$ let $T_t^k = (\mathbb{D}_k^+ \cap [0, t]) \cup \{t\}$, so that T_t^k is the finite set of such dyadic rationals that are less than t , with t added to the set. Note that T_t^k has an ordered enumeration, so $\mathbf{X}^k = \{X_s : s \in T_t^k\}$ is a discrete-time sub-martingale for each $k \in \mathbb{N}$. Let $U_t^k = \sup\{X_s : s \in T_t^k\}$ for $k \in \mathbb{N}$. Note that $T_t^j \subset T_t^k \subset [0, t]$ for $t \in [0, \infty)$ and for $j, k \in \mathbb{N}$ with $j < k$ and therefore $U_t^j \leq U_t^k \leq U_t$. It follows that for $x \in (0, \infty)$,

$$\{U_t^j \geq x\} \subseteq \{U_t^k \geq x\} \subseteq \{U_t \geq x\} \quad (17.4.11)$$

The set \mathbb{D}^+ of all nonnegative dyadic rationals is dense in $[0, \infty)$ and so since \mathbf{X} is right continuous and has left limits, it follows that if $U_t \geq x$ then $U_t^k \geq x$ for some $k \in \mathbb{N}$. That is, we have

$$\{U_t \geq x\} = \bigcup_{k=0}^{\infty} \{U_t^k \geq x\} \quad (17.4.12)$$

The maximal inequality applies to the discrete-time sub-martingale \mathbf{X}^k and so

$$P(U_t^k \geq x) \leq \frac{1}{x} \mathbb{E}(X_t; U_t^k \geq x) \quad (17.4.13)$$

for each $k \in \mathbb{N}$. By the monotone convergence theorem, the left side converges to $\mathbb{P}(U_t \geq x)$ as $k \rightarrow \infty$ and the right side converges to $\mathbb{E}(X; U_t \geq x)$ as $k \rightarrow \infty$.

There are a number of simple corollaries of the maximal inequality. For the first, recall that the *positive part* of $x \in \mathbb{R}$ is $x^+ = x \vee 0$, so that $x^+ = x$ if $x > 0$ and $x^+ = 0$ if $x \leq 0$.

Suppose that \mathbf{X} is a sub-martingale. For $t \in T$, let $V_t = \sup\{X_s^+ : s \in T_t\}$. Then

$$\mathbb{P}(V_t \geq x) \leq \frac{1}{x} \mathbb{E}(X_t^+; V_t \geq x), \quad x \in (0, \infty) \quad (17.4.14)$$

Proof

Recall that since \mathbf{X} is a sub-martingale and $x \mapsto x^+$ is increasing and convex, $\mathbf{X}^+ = \{X_t^+ : t \in T\}$ is also a sub-martingale. Hence the result follows from the general maximal inequality for sub-martingales.

As a further simple corollary, note that

$$\mathbb{P}(V_t \geq x) \leq \frac{1}{x} \mathbb{E}(X_t^+), \quad x \in (0, \infty) \quad (17.4.15)$$

This is sometimes how the maximal inequality is given in the literature.

Suppose that \mathbf{X} is a martingale. For $t \in T$, let $W_t = \sup\{|X_s| : s \in T_t\}$. Then

$$\mathbb{P}(W_t \geq x) \leq \frac{1}{x} \mathbb{E}(|X_t|; W_t \geq x), \quad x \in (0, \infty) \quad (17.4.16)$$

Proof

Recall that since \mathbf{X} is a martingale, and $x \mapsto |x|$ is convex, $|\mathbf{X}| = \{|X_t| : t \in T\}$ is a sub-martingale. Hence the result follows from the general maximal inequality for sub-martingales.

Once again, a further simple corollary is

$$\mathbb{P}(W_t \geq x) \leq \frac{1}{x} \mathbb{E}(|X_t|), \quad x \in (0, \infty) \quad (17.4.17)$$

Next recall that for $k \in (1, \infty)$, the k -norm of a real-valued random variable X is $\|X\|_k = [\mathbb{E}(|X|^k)]^{1/k}$, and the vector space \mathcal{L}_k consists of all real-valued random variables for which this norm is finite. The following theorem is the *norm version* of the Doob's maximal inequality.

Suppose again that \mathbf{X} is a martingale. For $t \in T$, let $W_t = \sup\{|X_s| : s \in T_t\}$. Then for $k > 1$,

$$\|W_t\|_k \leq \frac{k}{k-1} \|X_t\|_k \quad (17.4.18)$$

Proof

Fix $t \in T$. If $\mathbb{E}(|X_t|^k) = \infty$, the inequality trivial holds, so assume that $\mathbb{E}(|X_t|^k) < \infty$, and thus that $X_t \in \mathcal{L}_k$. The proof relies fundamentally on Hölder's inequality, and for that inequality to work, we need to truncate the variable W_t and consider instead the bounded random variable $W_t \wedge c$ where $c \in (0, \infty)$. First we need to show that

$$\mathbb{P}(W_t \wedge c \geq x) \leq \frac{1}{x} \mathbb{E}(|X_t|; W_t \wedge c \geq x), \quad x \in (0, \infty) \quad (17.4.19)$$

If $c < x$, both sides are 0. If $c \geq x$, $\{W_t \wedge c \geq x\} = \{W_t \geq x\}$ and so from the [maximal inequality](#) above,

$$\mathbb{P}(W_t \wedge c \geq x) = \mathbb{P}(W_t \geq x) \leq \frac{1}{x} \mathbb{E}(|X_t|; W_t \geq x) = \mathbb{E}(|X_t|; W_t \wedge c \geq x) \quad (17.4.20)$$

Next recall that

$$\|W_t \wedge c\|_k^k = \mathbb{E}[(W_t \wedge c)^k] = \int_0^\infty kx^{k-1} \mathbb{P}(W_t \wedge c \geq x) dx \quad (17.4.21)$$

Applying the inequality gives

$$\mathbb{E}[(W_t \wedge c)^k] \leq \int_0^\infty kx^{k-2} \mathbb{E}[|X_t|; W_t \wedge c \geq x] dx \quad (17.4.22)$$

By Fubini's theorem we can interchange the expected value and the integral which gives

$$\mathbb{E}[(W_t \wedge c)^k] \leq \mathbb{E} \left[\int_0^{W_t \wedge c} kx^{k-2} |X_t| dx \right] = \frac{k}{k-1} \mathbb{E}[|X_t| (W_t \wedge c)^{k-1}] \quad (17.4.23)$$

But $X_t \in \mathcal{L}_k$ and $(W_t \wedge c)^{k-1} \in \mathcal{L}_j$ where $j = k/(k-1)$ is the exponent conjugate to k . So an application of Hölder's inequality gives

$$\|W_t \wedge c\|_k^k \leq \frac{k}{k-1} \|X_t\|_k \| (W_t \wedge c)^{k-1} \|_j = \frac{k}{k-1} \|X_t\|_k \|W_t \wedge c\|_k^{k-1} \quad (17.4.24)$$

where we have used the simple fact that $\|(W_t \wedge c)^{k-1}\|_j = \|W_t \wedge c\|_k^{k-1}$. Dividing by this factor gives

$$\|W_t \wedge c\|_k \leq \frac{k}{k-1} \|X_t\|_k \quad (17.4.25)$$

Finally, $\|W_t \wedge c\|_k \uparrow \|W_t\|_k$ as $c \rightarrow \infty$ by the monotone convergence theorem. So letting $c \rightarrow \infty$ in the last displayed equation gives

$$\|W_t\|_k \leq \frac{k}{k-1} \|X_t\|_k \quad (17.4.26)$$

Once again, $\mathbf{W} = \{W_t : t \in T\}$ is the maximal process associated with $|\mathbf{X}| = \{|X_t| : t \in T\}$. As noted in the proof, $j = k/(k-1)$ is the exponent *conjugate* to k , satisfying $1/j + 1/k = 1$. So this version of the maximal inequality states that the k norm of the maximum of the martingale \mathbf{X} on T_t is bounded by j times the k norm of X_t , where j and k are conjugate exponents. Stated just in terms of expected value, rather than norms, the \mathcal{L}_k maximal inequality is

$$\mathbb{E}(|W_t|^k) \leq \left(\frac{k}{k-1}\right)^k \mathbb{E}(|X_t|^k) \quad (17.4.27)$$

Our final result in this discussion is a variation of the maximal inequality for super-martingales.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a nonnegative super-martingale, and let $U_\infty = \sup\{X_t : t \in T\}$. Then

$$\mathbb{P}(U_\infty \geq x) \leq \frac{1}{x} \mathbb{E}(X_0), \quad x \in (0, \infty) \quad (17.4.28)$$

Proof

Let $Y_t = -X_t$ for $t \in T$. Since \mathbf{X} is a super-martingale, \mathbf{Y} is a sub-martingale. And since \mathbf{X} is nonnegative, $Y_t^+ = X_t$ for $t \in T$. Let $U_t = \sup\{X_s : s \in T_t\} = \sup\{Y_s^+ : s \in T_t\}$ for $t \in T$. By the maximal inequality for sub-martingales, and since \mathbf{X} is a super-martingale we have for $t \in T$,

$$\mathbb{P}(U_t \geq x) \leq \frac{1}{x} \mathbb{E}(Y_t^+) = \frac{1}{x} \mathbb{E}(X_t) \leq \frac{1}{x} \mathbb{E}(X_0), \quad x \in (0, \infty) \quad (17.4.29)$$

Next note that $U_t \uparrow U_\infty$ as $t \rightarrow \infty$. Let $x \in (0, \infty)$ and $\epsilon \in (0, x)$. If $U_\infty \geq x$ then $U_t \geq x - \epsilon$ for sufficiently large $t \in T$. Hence

$$\{U_\infty \geq x\} \subseteq \bigcup_{k=1}^{\infty} \{U_k \geq x - \epsilon\} \quad (17.4.30)$$

Using the continuity theorem for increasing events, and our result above we have

$$\mathbb{P}(U_\infty \geq x) \leq \lim_{k \rightarrow \infty} \mathbb{P}(U_k \geq x - \epsilon) \leq \frac{1}{x - \epsilon} \mathbb{E}(X_0) \quad (17.4.31)$$

Since this holds for all $\epsilon \in (0, x)$, it follows that $\mathbb{P}(U_\infty \geq x) \leq \frac{1}{x} \mathbb{E}(X_0)$.

The Up-Crossing Inequality

The up-crossing inequality gives a bound on how much a sub-martingale (or super-martingale) can oscillate, and is the main tool in the martingale convergence theorems that will be studied in the next section. It should come as no surprise by now that the inequality is due to Joseph Doob. We start with the discrete-time case.

Suppose that $\mathbf{x} = (x_n : n \in \mathbb{N})$ is a sequence of real numbers, and that $a, b \in \mathbb{R}$ with $a < b$. Define $t_0(\mathbf{x}) = 0$ and then recursively define

$$\begin{aligned} s_{k+1}(\mathbf{x}) &= \inf\{n \in \mathbb{N} : n \geq t_k(\mathbf{x}), x_n \leq a\}, \quad k \in \mathbb{N} \\ t_{k+1}(\mathbf{x}) &= \inf\{n \in \mathbb{N} : n \geq s_{k+1}(\mathbf{x}), x_n \geq b\}, \quad k \in \mathbb{N} \end{aligned}$$

1. The number of up-crossings of the interval $[a, b]$ by the sequence \mathbf{x} up to time $n \in \mathbb{N}$ is

$$u_n(a, b, \mathbf{x}) = \sup\{k \in \mathbb{N} : t_k(\mathbf{x}) \leq n\} \quad (17.4.32)$$

2. The total number of up-crossings of the interval $[a, b]$ by the sequence \mathbf{x} is

$$u_{\infty}(a, b, \mathbf{x}) = \sup\{k \in \mathbb{N} : t_k(\mathbf{x}) < \infty\} \quad (17.4.33)$$

Details

As usual, we define $\inf(\emptyset) = \infty$. Note that if $t_k(\mathbf{x}) < \infty$ for $k \in \mathbb{N}_+$, then $(x_n : n = s_k(\mathbf{x}), \dots, t_k(\mathbf{x}))$ is the k th up-crossing of the interval $[a, b]$ by the sequence \mathbf{x} .

So informally, as the name suggests, $u_n(a, b, \mathbf{x})$ is the number of times that the sequence (x_0, x_1, \dots, x_n) goes from a value below a to one above b , and $u(a, b, \mathbf{x})$ is the number of times the entire sequence \mathbf{x} goes from a value below a to one above b . Here are a few of simple properties:

Suppose again that $\mathbf{x} = (x_n : n \in \mathbb{N})$ is a sequence of real numbers and that $a, b \in \mathbb{R}$ with $a < b$.

1. $u_n(a, b, \mathbf{x})$ is increasing in $n \in \mathbb{N}$.
2. $u_n(a, b, \mathbf{x}) \rightarrow u(a, b, \mathbf{x})$ as $n \rightarrow \infty$.
3. If $c, d \in \mathbb{R}$ with $a < c < d < b$ then $u_n(c, d, \mathbf{x}) \geq u_n(a, b, \mathbf{x})$ for $n \in \mathbb{N}$, and $u(c, d, \mathbf{x}) \geq u(a, b, \mathbf{x})$.

Proof

1. Note that $\{k \in \mathbb{N} : t_k(\mathbf{x}) \leq n\} \subseteq \{k \in \mathbb{N} : t_k(\mathbf{x}) \leq n+1\}$.
2. Note that $\bigcup_{n=0}^{\infty} \{k \in \mathbb{N} : t_k(\mathbf{x}) \leq n\} = \{k \in \mathbb{N} : t_k(\mathbf{x}) < \infty\}$.
3. Every up-crossing of $[a, b]$ is also an up-crossing of $[c, d]$.

The importance of the definitions is found in the following theorem. Recall that $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, \infty\}$ is the set of extended real numbers, and \mathbb{Q} is the set of rational real numbers.

Suppose again that $\mathbf{x} = (x_n : n \in \mathbb{N})$ is a sequence of real numbers. Then $\lim_{n \rightarrow \infty} x_n$ exists in \mathbb{R}^* is and only if $u_{\infty}(a, b, \mathbf{x}) < \infty$ for every $a, b \in \mathbb{Q}$ with $a < b$.

Proof

We prove the contrapositive. Note that the following statements are equivalent:

1. $\lim_{n \rightarrow \infty} x_n$ does not exist in \mathbb{R}^* .
2. $\liminf_{n \rightarrow \infty} x_n < \limsup_{n \rightarrow \infty} x_n$.
3. There exists $a, b \in \mathbb{Q}$ with $a < b$ and with $x_n \leq a$ for infinitely many $n \in \mathbb{N}$ and $x_n \geq b$ for infinitely many $n \in \mathbb{N}$.
4. There exists $a, b \in \mathbb{Q}$ with $a < b$ and $u_{\infty}(a, b, \mathbf{x}) = \infty$.

Clearly the theorem is true with \mathbb{Q} replaced with \mathbb{R} , but the countability of \mathbb{Q} will be important in the martingale convergence theorem. As a simple corollary, if \mathbf{x} is bounded and $u_{\infty}(a, b, \mathbf{x}) < \infty$ for every $a, b \in \mathbb{Q}$ with $a < b$, then \mathbf{x} converges in \mathbb{R} . The up-crossing inequality for a discrete-time martingale \mathbf{X} gives an upper bound on the expected number of up-crossings of \mathbf{X} up to time $n \in \mathbb{N}$ in terms of a moment of X_n .

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ satisfies the basic assumptions with respect to the filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$, and let $a, b \in \mathbb{R}$ with $a < b$. Let $U_n = u_n(a, b, \mathbf{X})$, the random number of up-crossings of $[a, b]$ by \mathbf{X} up to time $n \in \mathbb{N}$.

1. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then

$$\mathbb{E}(U_n) \leq \frac{1}{b-a} \mathbb{E}[(X_n - a)^-] \leq \frac{1}{b-a} [\mathbb{E}(X_n^-) + |a|] \leq \frac{1}{b-a} [\mathbb{E}(|X_n|) + |a|], \quad n \in \mathbb{N} \quad (17.4.34)$$

2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then

$$\mathbb{E}(U_n) \leq \frac{1}{b-a} \mathbb{E}[(X_n - a)^+] \leq \frac{1}{b-a} [\mathbb{E}(X_n^+) + |a|] \leq \frac{1}{b-a} [\mathbb{E}(|X_n|) + |a|], \quad n \in \mathbb{N} \quad (17.4.35)$$

Proof

In the context of the up-crossing definition above, let $\sigma_k = s_k(\mathbf{X})$ and $\tau_k = t_k(\mathbf{X})$. These are the random times that define the up-crossings of \mathbf{X} . Let $Y_k = X_{\tau_k \wedge n} - X_{\sigma_k \wedge n}$ and then define $Z_n = \sum_{k=1}^n Y_k$. To understand the sum, let's take cases for the k th term Y_k :

- If $\tau_k \leq n$ then $Y_k = X_{\tau_k} - X_{\sigma_k} \geq b - a$. By definition, the first U_n terms are of this form.
- If $\sigma_k \leq n < \tau_k$ then $Y_k = X_n - X_{\sigma_k} \geq X_n - a$. There is at most one such term, with index $k = U_n + 1$.
- If $\sigma_k > n$ then $Y_k = X_n - X_n = 0$.

Hence $Z_n \geq (b-a)U_n + (X_n - a)\mathbf{1}(\sigma_{U_n+1} \leq n)$ and so $(b-a)U_n \leq Z_n - (X_n - a)\mathbf{1}(\sigma_{U_n+1} \leq n)$. Next note that $\sigma_k \wedge n$ and $\tau_k \wedge n$ are bounded stopping times and of course $\sigma_k \wedge n \leq \tau_k \wedge n$.

1. If \mathbf{X} is a super-martingale, it follows from the optional stopping theorem that

$$\mathbb{E}(Y_k) = \mathbb{E}(X_{\tau_k \wedge n}) - \mathbb{E}(X_{\sigma_k \wedge n}) \leq 0 \quad (17.4.36)$$

and therefore $\mathbb{E}(Z_n) \leq 0$. Finally, $-(X_n - a)\mathbf{1}(\sigma_{U_n+1} \leq n) \leq (X_n - a)^-$. Taking expected values gives

$$(b-a)\mathbb{E}(U_n) \leq \mathbb{E}(Z_n) + \mathbb{E}[(X_n - a)^-] \leq \mathbb{E}[(X_n - a)^-] \quad (17.4.37)$$

The remaining parts of the inequality follow since $(x-a)^- \leq x^- + |a| \leq |x| + |a|$ for $x \in \mathbb{R}$.

Additional details

The process $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ in the proof can be viewed as a transform of $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ by a predictable process. Specifically, for $n \in \mathbb{N}_+$, let $I_n = 1$ if $\sigma_k < n \leq \tau_k$ for some $k \in \mathbb{N}$, and let $I_n = 0$ otherwise. Since σ_k and τ_k are stopping times, note that $\{I_n = 1\} \in \mathcal{F}_{n-1}$ for $n \in \mathbb{N}_+$. Hence the process $\mathbf{I} = \{I_n : n \in \mathbb{N}_+\}$ is predictable with respect to \mathcal{F} . Moreover, the transform of \mathbf{X} by \mathbf{I} is

$$(\mathbf{I} \cdot \mathbf{X})_n = \sum_{j=1}^n I_j(X_j - X_{j-1}) = \sum_{k=1}^n (X_{\tau_k \wedge n} - X_{\sigma_k \wedge n}) = Z_n, \quad n \in \mathbb{N} \quad (17.4.38)$$

Since \mathbf{I} is a nonnegative process, if \mathbf{X} is a martingale (sub-martingale, super-martingale), then $\mathbf{I} \cdot \mathbf{X}$ is also a martingale (sub-martingale, super-martingale).

Of course if \mathbf{X} is a martingale with respect to \mathcal{F} then both inequalities apply. In continuous time, as usual, the concepts are more complicated and technical.

Suppose that $\mathbf{x} : [0, \infty) \rightarrow \mathbb{R}$ and that $a, b \in \mathbb{R}$ with $a < b$.

1. If $I \subset [0, \infty)$ is finite, define $t_0^I(\mathbf{x}) = 0$ and then recursively define

$$\begin{aligned} s_{k+1}^I(\mathbf{x}) &= \inf \{t \in I : t \geq t_k^I(\mathbf{x}), x_t \leq a\}, \quad k \in \mathbb{N} \\ t_{k+1}^I(\mathbf{x}) &= \inf \{t \in I : t \geq s_{k+1}^I(\mathbf{x}), x_t \geq b\}, \quad k \in \mathbb{N} \end{aligned}$$

The number of up-crossings of the interval $[a, b]$ by the function \mathbf{x} restricted to I is

$$u_I(a, b, \mathbf{x}) = \sup \{k \in \mathbb{N} : t_k^I(\mathbf{x}) < \infty\} \quad (17.4.39)$$

2. If $I \subseteq [0, \infty)$ is infinite, the number of up-crossings of the interval $[a, b]$ by \mathbf{x} restricted to I is

$$u_I(a, b, \mathbf{x}) = \sup \{u_J(a, b, \mathbf{x}) : J \text{ is finite and } J \subset I\} \quad (17.4.40)$$

To simplify the notation, we will let $u_t(a, b, \mathbf{x}) = u_{[0, t]}(a, b, \mathbf{x})$, the number of up-crossings of $[a, b]$ by \mathbf{x} on $[0, t]$, and $u_\infty(a, b, \mathbf{x}) = u_{[0, \infty)}(a, b, \mathbf{x})$, the total number of up-crossings of $[a, b]$ by \mathbf{x} . In continuous time, the definition of up-crossings is built out of finite subsets of $[0, \infty)$ for measurability concerns, which arise when we replace the deterministic function \mathbf{x} with a stochastic process \mathbf{X} . Here are the simple properties that are analogous to our previous ones.

Suppose again that $\mathbf{x} : [0, \infty) \rightarrow \mathbb{R}$ and that $a, b \in \mathbb{R}$ with $a < b$.

1. If $I, J \subseteq [0, \infty)$ with $I \subseteq J$, then $u_I(a, b, \mathbf{x}) \leq u_J(a, b, \mathbf{x})$.
2. If $(I_n : n \in \mathbb{N})$ is an increasing sequence of sets in $[0, \infty)$ and $J = \bigcup_{n=0}^\infty I_n$ then $u_{I_n}(a, b, \mathbf{x}) \rightarrow u_J(a, b, \mathbf{x})$ as $n \rightarrow \infty$.
3. If $c, d \in \mathbb{R}$ with $a < c < d < b$ and $I \subset [0, \infty)$ then $u_I(c, d, \mathbf{x}) \geq u_I(a, b, \mathbf{x})$.

Proof

1. The result follows easily from the definitions if I is finite (and J either finite or infinite). If I is infinite (and hence so is J), note that

$$\{u_K(a, b, \mathbf{x}) : K \text{ is finite and } K \subseteq I\} \subseteq \{u_K(a, b, \mathbf{x}) : K \text{ is finite and } K \subseteq J\} \quad (17.4.41)$$

2. Since I_n is increasing in $n \in \mathbb{N}$ (in the subset partial order), note that if $K \subset [0, \infty)$ is finite, then $K \subseteq J$ if and only if $K \subseteq I_n$ for some $n \in \mathbb{N}$.
3. Every up-crossing of $[a, b]$ is an up-crossing of $[c, d]$.

The following result is the reason for studying up-crossings in the first place. Note that the definition built from finite set is sufficient.

Suppose that $\mathbf{x} : [0, \infty) \rightarrow \mathbb{R}$. Then $\lim_{t \rightarrow \infty} x_t$ exists in \mathbb{R}^* if and only if $u_\infty(a, b, \mathbf{x}) < \infty$ for every $a, b \in \mathbb{Q}$ with $a < b$.

Proof

As in the discrete-time case, we prove the contrapositive. The proof is almost the same: The following statements are equivalent:

1. $\lim_{t \rightarrow \infty} x_t$ does not exist in \mathbb{R}^* .
2. $\liminf_{t \rightarrow \infty} x_t < \limsup_{t \rightarrow \infty} x_t$.
3. There exists $a, b \in \mathbb{Q}$ with $a < b$ and there exists $s_n, t_n \in [0, \infty)$ with $x_{s_n} \leq a$ for $n \in \mathbb{N}$ and $x_{t_n} \geq b$ for $n \in \mathbb{N}$.
4. There exists $a, b \in \mathbb{Q}$ with $a < b$ and $u_\infty(a, b, \mathbf{x}) = \infty$.

Finally, here is the up-crossing inequality for martingales in continuous time. Once again, the inequality gives a bound on the expected number of up-crossings.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ satisfies the basic assumptions with respect to the filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$, and let $a, b \in \mathbb{R}$ with $a < b$. Let $U_t = u_t(a, b, \mathbf{X})$, the random number of up-crossings of $[a, b]$ by \mathbf{X} up to time $t \in [0, \infty)$.

1. If \mathbf{X} is a super-martingale relative to \mathfrak{F} then

$$\mathbb{E}(U_t) \leq \frac{1}{b-a} \mathbb{E}[(X_t - a)^-] \leq \frac{1}{b-a} [\mathbb{E}(X_t^-) + |a|] \leq \frac{1}{b-a} [\mathbb{E}(|X_t|) + |a|], \quad t \in [0, \infty) \quad (17.4.42)$$

2. If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then

$$\mathbb{E}(U_t) \leq \frac{1}{b-a} \mathbb{E}[(X_t - a)^+] \leq \frac{1}{b-a} [\mathbb{E}(X_t^+) + |a|] \leq \frac{1}{b-a} [\mathbb{E}(|X_t|) + |a|], \quad t \in [0, \infty) \quad (17.4.43)$$

Proof

Suppose that \mathbf{X} is a sub-martingale; the proof for a super-martingale is analogous. Fix $t \in [0, \infty)$ and $a, b \in \mathbb{R}$ with $a < b$. For $I \subseteq [0, \infty)$ let $U_I = u_I(a, b, \mathbf{X})$, the number of up-crossings of $[a, b]$ by \mathbf{X} restricted to I . Suppose that I is finite and that $t \in I$ is the maximum of I . Since \mathbf{X} restricted to I is also a sub-martingale, the discrete-time up-crossing theorem applies and so

$$\mathbb{E}(U_I) \leq \frac{1}{b-a} \mathbb{E}[(X_t - a)^+] \quad (17.4.44)$$

Since $U_t = \sup\{U_I : I \text{ is finite and } I \subset [0, t]\}$, there exists finite I_n for $n \in \mathbb{N}$ with $U_{I_n} \uparrow U_t$ as $n \rightarrow \infty$. In particular, U_t is measurable. By property (a) in the theorem [above](#), there exists such a sequence with I_n increasing in n and $t \in I_n$ for each $n \in \mathbb{N}$. By the monotone convergence theorem, $\mathbb{E}(U_{I_n}) \rightarrow \mathbb{E}(U_t)$ as $n \rightarrow \infty$. So by the displayed equation above,

$$\mathbb{E}(U_t) \leq \frac{1}{b-a} \mathbb{E}[(X_t - a)^+] \quad (17.4.45)$$

Examples and Applications

Kolmogorov's Inequality

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is a sequence of independent variables with $\mathbb{E}(X_n) = 0$ and $\text{var}(X_n) = \mathbb{E}(X_n^2) < \infty$ for $n \in \mathbb{N}_+$. Let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ be the partial sum process associated with \mathbf{X} , so that

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (17.4.46)$$

From the Introduction we know that \mathbf{Y} is a martingale. A simple application of the maximal inequality gives the following result, which is known as *Kolmogorov's inequality*, named for Andrei Kolmogorov.

For $n \in \mathbb{N}$, let $U_n = \max\{|Y_i| : i \in \mathbb{N}_n\}$. Then

$$\mathbb{P}(U_n \geq x) \leq \frac{1}{x^2} \text{var}(Y_n) = \frac{1}{x^2} \sum_{i=1}^n \mathbb{E}(X_i^2), \quad x \in (0, \infty) \quad (17.4.47)$$

Proof

As noted above, \mathbf{Y} is a martingale. Since the function $x \mapsto x^2$ on \mathbb{R} is convex, $\mathbf{Y}^2 = \{Y_n^2 : n \in \mathbb{N}\}$ is a sub-martingale. Let $V_n = \max\{Y_i^2 : i \in \mathbb{N}_n\}$ for $n \in \mathbb{N}$, and let $x \in (0, \infty)$. Applying the maximal inequality for sub-martingales we have

$$\mathbb{P}(U_n \geq x) = \mathbb{P}(V_n \geq x^2) \leq \frac{1}{x^2} \mathbb{E}(Y_n^2) = \frac{1}{x^2} \text{var}(Y_n) \quad (17.4.48)$$

Finally, since \mathbf{X} is an independent sequence,

$$\text{var}(Y_n) = \sum_{i=1}^n \text{var}(X_i) = \sum_{i=1}^n \mathbb{E}(X_i^2) \quad (17.4.49)$$

Red and Black

In the game of red and black, a gambler plays a sequence of Bernoulli games with success parameter $p \in (0, 1)$ at even stakes. The gambler starts with an initial fortune x and plays until either she is ruined or reaches a specified target fortune a , where $x, a \in (0, \infty)$ with $x < a$. When $p \leq \frac{1}{2}$, so that the games are fair or unfair, an optimal strategy is *bold play*: on each game, the gambler bets her entire fortune or just what is needed to reach the target, whichever is smaller. In the section on bold play we showed that when $p = \frac{1}{2}$, so that the games are fair, the probability of winning (that is, reaching the target a starting with x) is x/a . We can use the maximal inequality for super-martingales to show that indeed, one cannot do better.

To set up the notation and review various concepts, let X_0 denote the gambler's initial fortune and let X_n denote the outcome of game $n \in \mathbb{N}_+$, where 1 denotes a win and -1 a loss. So $\{X_n : n \in \mathbb{N}\}$ is a sequence of independent variables with $\mathbb{P}(X_n = 1) = p$ and $\mathbb{P}(X_n = -1) = 1 - p$ for $n \in \mathbb{N}_+$. (The initial fortune X_0 has an unspecified distribution on $(0, \infty)$.) The gambler is at a casino after all, so of course $p \leq \frac{1}{2}$. Let

$$Y_n = \sum_{i=0}^n X_i, \quad n \in \mathbb{N} \quad (17.4.50)$$

so that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is the partial sum process associated with $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$. Recall that \mathbf{Y} is also known as the *simple random walk* with parameter p , and since $p \leq \frac{1}{2}$, is a super-martingale. The process $\{X_n : n \in \mathbb{N}_+\}$ is the *difference sequence* associated with \mathbf{Y} . Next let Z_n denote the amount that the gambler bets on game $n \in \mathbb{N}_+$. The process $\mathbf{Z} = \{Z_n : n \in \mathbb{N}_+\}$ is *predictable* with respect to $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$, so that Z_n is measurable with respect to $\sigma\{X_0, X_1, \dots, X_{n-1}\}$ for $n \in \mathbb{N}_+$. So the gambler's fortune after n games is

$$W_n = X_0 + \sum_{i=1}^n Z_i X_i = X_0 + \sum_{i=1}^n Z_i (Y_i - Y_{i-1}) \quad (17.4.51)$$

Recall that $\mathbf{W} = \{W_n : n \in \mathbb{N}\}$ is the *transform* of \mathbf{Z} with \mathbf{Y} , denoted $\mathbf{W} = \mathbf{Z} \cdot \mathbf{Y}$. The gambler is not allowed to go into debt and so we must have $Z_n \leq W_{n-1}$ for $n \in \mathbb{N}_+$: the gambler's bet on game n cannot exceed her fortune after game $n - 1$. What's the probability that the gambler can ever reach or exceed the target a starting with fortune $x < a$?

Let $U_\infty = \sup\{W_n : n \in \mathbb{N}\}$. Suppose that $x, a \in (0, \infty)$ with $x < a$ and that $X_0 = x$. Then

$$\mathbb{P}(U_\infty \geq a) \leq \frac{x}{a} \quad (17.4.52)$$

Proof

Since Y is a super-martingale and Z is nonnegative, the transform $W = Z \cdot Y$ is also a super-martingale. By the inequality for nonnegative super-martingales [above](#):

$$\mathbb{P}(U_\infty \geq a) \leq \frac{1}{a} \mathbb{E}(W_0) = \frac{x}{a} \quad (17.4.53)$$

Note that the only assumptions made on the gambler's sequence of bets Z is that the sequence is predictable, so that the gambler cannot see into the future, and that gambler cannot go into debt. Under these basic assumptions, no strategy can do any better than bold play. However, there *are* strategies that do as well as bold play; these are variations on bold play.

Open the simulation of the red and black game. Select bold play and $p = \frac{1}{2}$. Play the game with various values of initial and target fortunes.

This page titled [17.4: Inequalities](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

17.5: Convergence

Basic Theory

Basic Assumptions

As in the Introduction, we start with a stochastic process $\mathbf{X} = \{X_t : t \in T\}$ on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, having state space \mathbb{R} , and where the index set T (representing time) is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). Next, we have a filtration $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$, and we assume that \mathbf{X} is *adapted* to \mathfrak{F} . So \mathfrak{F} is an increasing family of sub σ -algebras of \mathcal{F} and X_t is measurable with respect to \mathcal{F}_t for $t \in T$. We think of \mathcal{F}_t as the collection of events up to time $t \in T$. We assume that $\mathbb{E}(|X_t|) < \infty$, so that the mean of X_t exists as a real number, for each $t \in T$. Finally, in continuous time where $T = [0, \infty)$, we need the additional assumptions that $t \mapsto X_t$ is right continuous and has left limits, and that the filtration \mathfrak{F} is standard (that is, right continuous and complete). Recall also that $\mathcal{F}_\infty = \sigma(\bigcup_{t \in T} \mathcal{F}_t)$, and this is the σ -algebra that encodes our information over all time.

The Martingale Convergence Theorems

If \mathbf{X} is a sub-martingale relative to \mathfrak{F} then \mathbf{X} has an *increasing property* of sorts: $E(X_t | \mathcal{F}_s) \geq X_s$ for $s, t \in T$ with $s \leq t$. Similarly, if \mathbf{X} is a super-martingale relative to \mathfrak{F} then \mathbf{X} has a *decreasing property* of sorts, since the last inequality is reversed. Thus, there is hope that if this increasing or decreasing property is coupled with an appropriate boundedness property, then the sub-martingale or super-martingale might converge, in some sense, as $t \rightarrow \infty$. This is indeed the case, and is the subject of this section. The martingale convergence theorems, first formulated by Joseph Doob, are among the most important results in the theory of martingales. The *first martingale convergence theorem* states that if the expected absolute value is bounded in the time, then the martingale process converges with probability 1.

Suppose that $\mathbf{X} = \{X_t : t \in T\}$ is a sub-martingale or a super-martingale with respect to $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ and that $\mathbb{E}(|X_t|)$ is bounded in $t \in T$. Then there exists a random variable X_∞ that is measurable with respect to \mathcal{F}_∞ such that $\mathbb{E}(|X_\infty|) < \infty$ and $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$ with probability 1.

Proof

The proof is simple using the up-crossing inequality. Let $T_t = \{s \in T : s \leq t\}$ for $t \in T$. For $a, b \in \mathbb{R}$ with $a < b$, let $U_t(a, b)$ denote the number of up-crossings of the interval $[a, b]$ by the process \mathbf{X} on T_t , and let $U_\infty(a, b)$ denote the number of up-crossings of $[a, b]$ by \mathbf{X} on T . Recall that $U_t \uparrow U_\infty$ as $t \rightarrow \infty$. Suppose that $\mathbb{E}(|X_t|) < c$ for $t \in T$, where $c \in (0, \infty)$. By the up-crossing inequality,

$$\mathbb{E}[U_t(a, b)] \leq \frac{1}{b-a} [|a| + \mathbb{E}(|X_t|)] \leq \frac{|a| + c}{b-a}, \quad n \in \mathbb{N} \quad (17.5.1)$$

By the monotone convergence theorem, it follows that

$$\mathbb{E}[U_\infty(a, b)] < \frac{|a| + c}{b-a} < \infty \quad (17.5.2)$$

Hence $\mathbb{P}[U_\infty(a, b) < \infty] = 1$. Therefore with probability 1, $U_\infty(a, b) < \infty$ for every $a, b \in \mathbb{Q}$ with $a < b$. By our characterization of convergence in terms of up-crossings, it follows that there exists a random variable X_∞ with values in $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, \infty\}$ such that with probability 1, $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$. Note that X is measurable with respect to \mathcal{F}_∞ . By Fatou's lemma,

$$\mathbb{E}(|X_\infty|) \leq \liminf_{t \rightarrow \infty} \mathbb{E}(|X_t|) < \infty \quad (17.5.3)$$

Hence $\mathbb{P}(X_\infty \in \mathbb{R}) = 1$.

The boundedness condition means that \mathbf{X} is bounded (in norm) as a subset of the vector space \mathcal{L}_1 . Here is a very simple, but useful corollary:

If $\mathbf{X} = \{X_t : t \in T\}$ is a nonnegative super-martingale with respect to $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ then there exists a random variable X_∞ , measurable with respect to \mathcal{F}_∞ , such that $X_t \rightarrow X_\infty$ with probability 1.

Proof

Since \mathbf{X} is a nonnegative super-martingale, $\mathbb{E}(|X_t|) = \mathbb{E}(X_t) \leq \mathbb{E}(X_0)$ for $t \in T$. Hence the previous martingale convergence theorem applies.

Of course, the corollary applies to a nonnegative martingale as a special case. For the second *martingale convergence theorem* you will need to review uniformly integrable variables. Recall also that for $k \in [1, \infty)$, the k -norm of a random variable X is

$$\|X\|_k = \left[\mathbb{E}(|X|^k) \right]^{1/k} \quad (17.5.4)$$

and \mathcal{L}_k is the normed vector space of all real-valued random variables for which this norm is finite. *Convergence in mean* refers to convergence in \mathcal{L}_1 and more generally, *convergence in k th mean* refers to convergence in \mathcal{L}_k .

Suppose that \mathbf{X} is a uniformly integrable and is a sub-martingale or super-martingale with respect to \mathfrak{F} . Then there exists a random variable X_∞ , measurable with respect to \mathcal{F}_∞ such that $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$ with probability 1 and in mean. Moreover, if \mathbf{X} is a martingale with respect to \mathfrak{F} then $X_t = \mathbb{E}(X_\infty | \mathcal{F}_t)$ for $t \in T$.

Proof

Since $\mathbf{X} = \{X_t : t \in T\}$ is uniformly integrable, $\mathbb{E}(|X_t|)$ is bounded in $t \in T$. Hence the by the first martingale convergence theorem, there exists X_∞ that is measurable with respect to \mathcal{F}_∞ such that $\mathbb{E}(|X_\infty|) < \infty$ and $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$ with probability 1. By the uniform integrability theorem, the convergence is also in mean, so that $\mathbb{E}(|X_t - X_\infty|) \rightarrow 0$ as $t \rightarrow \infty$. Suppose now that \mathbf{X} is a martingale with respect to \mathfrak{F} . For fixed $s \in T$ we know that $\mathbb{E}(X_t | \mathcal{F}_s) \rightarrow \mathbb{E}(X_\infty | \mathcal{F}_s)$ as $t \rightarrow \infty$ (with probability 1). But $\mathbb{E}(X_t | \mathcal{F}_s) = X_s$ for $t \geq s$ so it follows that $X_s = \mathbb{E}(X_\infty | \mathcal{F}_s)$.

As a simple corollary, recall that if $\|X_t\|_k$ is bounded in $t \in T$ for some $k \in (1, \infty)$ then \mathbf{X} is uniformly integrable, and hence the second martingale convergence theorem applies. But we can do better.

Suppose again that $\mathbf{X} = \{X_t : t \in T\}$ is a sub-martingale or super-martingale with respect to $\mathfrak{F} = \{\mathcal{F}_t : t \in T\}$ and that $\|X_t\|_k$ is bounded in $t \in T$ for some $k \in (1, \infty)$. Then there exists a random variable $X_\infty \in \mathcal{L}_k$ such that $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$ in \mathcal{L}_k .

Proof

Suppose that $\|X_t\|_k \leq c$ for $t \in T$ where $c \in (0, \infty)$. Since $\|X\|_1 \leq \|X\|_k$, we have $\mathbb{E}(|X_t|)$ bounded in $t \in T$ so the first martingale convergence theorem applies. Hence there exists X_∞ , measurable with respect to \mathcal{F}_∞ , such that $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$ with probability 1. Equivalently, with probability 1,

$$|X_t - X_\infty|^k \rightarrow 0 \text{ as } t \rightarrow \infty \quad (17.5.5)$$

Next, for $t \in T$, let $T_t = \{s \in T : s \leq t\}$ define $W_t = \sup\{|X_s| : s \in T_t\}$. by the norm version of the maximal inequality,

$$\|W_t\|_k \leq \frac{k}{k-1} \|X_t\|_k \leq \frac{kc}{k-1}, \quad t \in T \quad (17.5.6)$$

If we let $W_\infty = \sup\{|X_s| : s \in T\}$, then by the montone convergence theorem

$$\|W_\infty\|_k = \lim_{t \rightarrow \infty} \|W_t\|_k \leq \frac{ck}{k-1} \quad (17.5.7)$$

So $W_\infty \in \mathcal{L}_k$. But $|X_\infty| \leq W_\infty$ so $X_\infty \in \mathcal{L}_k$ also. Moreover, $|X_t - X_\infty|^k \leq 2^k W_\infty^k$, so applying the dominated convergence theorem to the first displayed equation above, we have $\mathbb{E}(|X_t - X_\infty|^k) \rightarrow 0$ as $t \rightarrow \infty$.

Example and Applications

In this subsection, we consider a number of applications of the martingale convergence theorems. One indication of the importance of martingale theory is the fact that many of the classical theorems of probability have simple and elegant proofs when formulated in terms of martingales.

Simple Random Walk

Suppose now that $\mathbf{V} = \{V_n : n \in \mathbb{N}\}$ is a sequence of independent random variables with $\mathbb{P}(V_i = 1) = p$ and $\mathbb{P}(V_i = -1) = 1 - p$ for $i \in \mathbb{N}_+$, where $p \in (0, 1)$. Let $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ be the partial sum process associated with \mathbf{V} so that

$$X_n = \sum_{i=0}^n V_i, \quad n \in \mathbb{N} \quad (17.5.8)$$

Recall that \mathbf{X} is the simple random walk with parameter p . From our study of Markov chains, we know that $p > \frac{1}{2}$ then $X_n \rightarrow \infty$ as $n \rightarrow \infty$ and if $p < \frac{1}{2}$ then $X_n \rightarrow -\infty$ as $n \rightarrow \infty$. The chain is transient in these two cases. If $p = \frac{1}{2}$, the chain is (null) recurrent and so visits every state in \mathbb{N} infinitely often. In this case X_n does not converge as $n \rightarrow \infty$. But of course $\mathbb{E}(X_n) = n(2p - 1)$ for $n \in \mathbb{N}$, so the martingale convergence theorems do not apply.

Doob's Martingale

Recall that if X is a random variable with $\mathbb{E}(|X|) < \infty$ and we define $X_t = \mathbb{E}(X | \mathcal{F}_t)$ for $t \in T$, then $\mathbf{X} = \{X_t : t \in T\}$ is a martingale relative to \mathfrak{F} and is known as a Doob martingale, named for you know whom. So the second martingale convergence theorem states that every uniformly integrable martingale is a Doob martingale. Moreover, we know that the Doob martingale \mathbf{X} constructed from X and \mathfrak{F} is uniformly integrable, so the second martingale convergence theorem applies. The last remaining question is the relationship between X and the limiting random variable X_∞ . The answer may come as no surprise.

Let $\mathbf{X} = \{X_t : t \in T\}$ be the Doob martingale constructed from X and \mathfrak{F} . Then $X_t \rightarrow X_\infty$ as $t \rightarrow \infty$ with probability 1 and in mean, where

$$X_\infty = \mathbb{E}(X | \mathcal{F}_\infty) \quad (17.5.9)$$

Of course if $\mathcal{F}_\infty = \mathcal{F}$, which is quite possible, then $X_\infty = X$. At the other extreme, if $\mathcal{F}_t = \{\emptyset, \Omega\}$, the trivial σ -algebra for all $t \in T$, then $X_\infty = \mathbb{E}(X)$, a constant.

Kolmogorov Zero-One Law

Suppose that $\mathbf{X} = (X_n : n \in \mathbb{N}_+)$ is a sequence of random variables with values in a general state space (S, \mathcal{S}) . Let $\mathcal{G}_n = \sigma\{X_k : k \geq n\}$ for $n \in \mathbb{N}_+$, and let $\mathcal{G}_\infty = \bigcap_{n=1}^\infty \mathcal{G}_n$. So \mathcal{G}_∞ is the *tail σ -algebra* of \mathbf{X} , the collection of events that depend only on the terms of the sequence with arbitrarily large indices. For example, if the sequence is real-valued (or more generally takes values in a metric space), then the event that X_n has a limit as $n \rightarrow \infty$ is a tail event. If $B \in \mathcal{S}$, then the event that $X_n \in B$ for infinitely many $n \in \mathbb{N}_+$ is another tail event. The *Kolmogorov zero-one law*, named for Andrei Kolmogorov, states that if \mathbf{X} is an independent sequence, then the tail events are essentially deterministic.

Suppose that \mathbf{X} is a sequence of independent random variables. If $A \in \mathcal{G}_\infty$ then $\mathbb{P}(A) = 0$ or $\mathbb{P}(A) = 1$.

Proof

Let $\mathcal{F}_n = \sigma\{X_k : k \leq n\}$ for $n \in \mathbb{N}_+$ so that $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}_+\}$ is the natural filtration associated with \mathbf{X} . As with our notation above, let $\mathcal{F}_\infty = \sigma\left(\bigcup_{n \in \mathbb{N}_+} \mathcal{F}_n\right)$. Now let $A \in \mathcal{G}_\infty$ be a tail event. Then $\{\mathbb{E}(\mathbf{1}_A | \mathcal{F}_n) : n \in \mathbb{N}_+\}$ is the Doob martingale associated with the indicator variable $\mathbf{1}_A$ and \mathfrak{F} . By our results above, $\mathbb{E}(\mathbf{1}_A | \mathcal{F}_n) \rightarrow \mathbb{E}(\mathbf{1}_A | \mathcal{F}_\infty)$ as $n \rightarrow \infty$ with probability 1. But $A \in \mathcal{F}_\infty$ so $\mathbb{E}(\mathbf{1}_A | \mathcal{F}_\infty) = \mathbf{1}_A$. On the other hand, $A \in \mathcal{G}_{n+1}$ and the σ -algebras \mathcal{G}_{n+1} and \mathcal{F}_n are independent. Therefore $\mathbb{E}(\mathbf{1}_A | \mathcal{F}_n) = \mathbb{P}(A)$ for each $n \in \mathbb{N}_+$. Thus $\mathbb{P}(A) = \mathbf{1}_A$.

Tail events and the Kolmogorov zero-one law were studied earlier in the section on measure in the chapter on probability spaces. A random variable that is measurable with respect to \mathcal{G}_∞ is a *tail random variable*. From the Kolmogorov zero-one law, a real-valued tail random variable for an independent sequence must be a constant (with probability 1).

Branching Processes

Recall the discussion of the simple *branching process* from the Introduction. The fundamental assumption is that the particles act independently, each with the same offspring distribution on \mathbb{N} . As before, we will let f denote the (discrete) probability density function of the number of offspring of a particle, m the mean of the distribution, and q the probability of extinction starting with a single particle. We assume that $f(0) > 0$ and $f(0) + f(1) < 1$ so that a particle has a positive probability of dying without children and a positive probability of producing more than 1 child.

The stochastic process of interest is $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ where X_n is the number of particles in the n th generation for $n \in \mathbb{N}$. Recall that \mathbf{X} is a discrete-time Markov chain on \mathbb{N} . Since 0 is an absorbing state, and all positive states lead to 0, we know that the positive states are transient and so are visited only finitely often with probability 1. It follows that either $X_n \rightarrow 0$ as $n \rightarrow \infty$ (extinction) or $X_n \rightarrow \infty$ as $n \rightarrow \infty$ (explosion). We have quite a bit of information about which of these events will occur from our study of Markov chains, but the martingale convergence theorems give more information.

Extinction and explosion

1. If $m \leq 1$ then $q = 1$ and extinction is certain.
2. If $m > 1$ then $q \in (0, 1)$. Either $X_n \rightarrow 0$ as $n \rightarrow \infty$ or $X_n \rightarrow \infty$ as $n \rightarrow \infty$ at an exponential rate.

Proof

The new information is the rate of divergence to ∞ in (b). The other statements are from our study of discrete-time branching Markov chains. We showed in the Introduction that $\{X_n/m^n : n \in \mathbb{N}\}$ is a martingale. Since this martingale is nonnegative, it has a limit as $n \rightarrow \infty$, and the limiting random variable takes values in \mathbb{R} . So if $m > 1$ and $X_n \rightarrow \infty$ as $n \rightarrow \infty$, then the divergence to ∞ must be at essentially the same rate as m^n .

The Beta-Bernoulli Process

Recall that the beta-Bernoulli process is constructed by randomizing the success parameter in a Bernoulli trials process with a beta distribution. Specifically, we start with a random variable P having the beta distribution with parameters $a, b \in (0, \infty)$. Next we have a sequence $\mathbf{X} = (X_1, X_2, \dots)$ of indicator variables with the property that \mathbf{X} is conditionally independent given $P = p \in (0, 1)$ with $\mathbb{P}(X_i = 1 \mid P = p) = p$ for $i \in \mathbb{N}_+$. Let $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ denote the partial sum process associated with \mathbf{X} , so that once again, $Y_n = \sum_{i=1}^n X_i$ for $n \in \mathbb{N}$. Next let $M_n = Y_n/n$ for $n \in \mathbb{N}_+$ so that M_n is the *sample mean* of (X_1, X_2, \dots, X_n) . Finally let

$$Z_n = \frac{a + Y_n}{a + b + n}, \quad n \in \mathbb{N} \quad (17.5.10)$$

We showed in the Introduction that $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} .

$M_n \rightarrow P$ and $Z_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1 and in mean.

Proof

We showed in the section on the beta-Bernoulli process that $Z_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1. Note that $0 \leq Z_n \leq 1$ for $n \in \mathbb{N}$, so the martingale \mathbf{Z} is uniformly integrable. Hence the [second martingale convergence theorem](#) applies, and the convergence is in mean also.

This is a very nice result and is reminiscent of the fact that for the ordinary Bernoulli trials sequence with success parameter $p \in (0, 1)$ we have the law of large numbers that $M_n \rightarrow p$ as $n \rightarrow \infty$ with probability 1 and in mean.

Pólya's Urn Process

Recall that in the simplest version of Pólya's urn process, we start with an urn containing a red and b green balls. At each discrete time step, we select a ball at random from the urn and then replace the ball and add c new balls of the same color to the urn. For the parameters, we need $a, b \in \mathbb{N}_+$ and $c \in \mathbb{N}$. For $i \in \mathbb{N}_+$, let X_i denote the color of the ball selected on the i th draw, where 1 means red and 0 means green. For $n \in \mathbb{N}$, let $Y_n = \sum_{i=1}^n X_i$, so that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is the partial sum process associated with $\mathbf{X} = \{X_i : i \in \mathbb{N}_+\}$. Since Y_n is the number of red balls in the urn at time $n \in \mathbb{N}_+$, the *average* number of balls at time n is $M_n = Y_n/n$. On the other hand, the total number of balls in the urn at time $n \in \mathbb{N}$ is $a + b + cn$ so the *proportion* of red balls in the urn at time n is

$$Z_n = \frac{a + cY_n}{a + b + cn} \quad (17.5.11)$$

We showed in the Introduction, that $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is a martingale. Now we are interested in the limiting behavior of M_n and Z_n as $n \rightarrow \infty$. When $c = 0$, the answer is easy. In this case, Y_n has the binomial distribution with trial parameter n and success parameter $a/(a+b)$, so by the law of large numbers, $M_n \rightarrow a/(a+b)$ as $n \rightarrow \infty$ with probability 1 and in mean. On the other hand, $Z_n = a/(a+b)$ when $c = 0$. So the interesting case is when $c > 0$.

Suppose that $c \in \mathbb{N}_+$. Then there exists a random variable P such that $M_n \rightarrow P$ and $Z_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1 and in mean. Moreover, P has the beta distribution with left parameter a/c and right parameter b/c .

Proof

In our study of Pólya's urn process we showed that when $c \in \mathbb{N}_+$ the process \mathbf{X} is a beta-Bernoulli process with parameters a/c and b/c . So the result follows from our previous theorem.

Likelihood Ratio Tests

Recall the discussion of likelihood ratio tests in the Introduction. To review, suppose that (S, \mathcal{S}, μ) is a general measure space, and that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a sequence of independent, identically distributed random variables, taking values in S , and having a common probability density function with respect to μ . The likelihood ratio test is a hypothesis test, where the null and alternative hypotheses are

- H_0 : the probability density function is g_0 .
- H_1 : the probability density function is g_1 .

We assume that g_0 and g_1 are positive on S . Also, it makes no sense for g_0 and g_1 to be the same, so we assume that $g_0 \neq g_1$ on a set of positive measure. The test is based on the *likelihood ratio test statistic*

$$L_n = \prod_{i=1}^n \frac{g_0(X_i)}{g_1(X_i)}, \quad n \in \mathbb{N} \quad (17.5.12)$$

We showed that under the alternative hypothesis H_1 , $\mathbf{L} = \{L_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathbf{X} , known as the *likelihood ratio martingale*.

Under H_1 , $L_n \rightarrow 0$ as $n \rightarrow \infty$ with probability 1.

Proof

Assume that H_1 is true. \mathbf{L} is a nonnegative martingale, so the first martingale convergence theorem applies, and hence there exists a random variable L_∞ with values in $[0, \infty)$ such that $L_n \rightarrow L_\infty$ as $n \rightarrow \infty$ with probability 1. Next note that

$$\ln(L_n) = \sum_{i=1}^n \ln \left[\frac{g_0(X_i)}{g_1(X_i)} \right] \quad (17.5.13)$$

The variables $\ln[g_0(X_i)/g_1(X_i)]$ for $i \in \mathbb{N}_+$ are also independent and identically distributed, so let m denote the common mean. The natural logarithm is concave and the martingale \mathbf{L} has mean 1, so by Jensen's inequality,

$$m = \mathbb{E} \left(\ln \left[\frac{g_0(X)}{g_1(X)} \right] \right) < \ln \left(\mathbb{E} \left[\frac{g_0(X)}{g_1(X)} \right] \right) = \ln(1) = 0 \quad (17.5.14)$$

Hence $m \in [-\infty, 0)$. By the strong law of large numbers, $\frac{1}{n} \ln(L_n) \rightarrow m$ as $n \rightarrow \infty$ with probability 1. Hence we must have $\ln(L_n) \rightarrow -\infty$ as $n \rightarrow \infty$ with probability 1. But by continuity, $\ln(L_n) \rightarrow \ln(L_\infty)$ as $n \rightarrow \infty$ with probability 1, so $L_\infty = 0$ with probability 1.

This result is good news, statistically speaking. Small values of L_n are evidence in favor of H_1 , so the decision rule is to reject H_0 in favor of H_1 if $L_n \leq l$ for a chosen critical value $l \in (0, \infty)$. If H_1 is true and the sample size n is sufficiently large, we will reject H_0 . In the proof, note that $\ln(L_n)$ must diverge to $-\infty$ at least as fast as n diverges to ∞ . Hence $L_n \rightarrow 0$ as $n \rightarrow \infty$ exponentially fast, at least. It also worth noting that \mathbf{L} is a mean 1 martingale (under H_1) so trivially $\mathbb{E}(L_n) \rightarrow 1$ as $n \rightarrow \infty$ even though $L_n \rightarrow 0$ as $n \rightarrow \infty$ with probability 1. So the likelihood ratio martingale is a good example of a sequence where the interchange of limit and expected value is not valid.

Partial Products

Suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is an independent sequence of nonnegative random variables with $\mathbb{E}(X_n) = 1$ for $n \in \mathbb{N}_+$. Let

$$Y_n = \prod_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (17.5.15)$$

so that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is the partial product process associated with \mathbf{X} . From our discussion of this process in the Introduction, we know that \mathbf{Y} is a martingale with respect to \mathbf{X} . Since \mathbf{Y} is nonnegative, the [second martingale convergence theorem](#) applies, so there exists a random variable Y_∞ such that $Y_n \rightarrow Y_\infty$ as $n \rightarrow \infty$ with probability 1. What more can we say? The following result, known as the *Kakutani product martingale theorem*, is due to Shizuo Kakutani.

Let $a_n = \mathbb{E}(\sqrt{X_n})$ for $n \in \mathbb{N}_+$ and let $A = \prod_{i=1}^{\infty} a_i$.

1. If $A > 0$ then $Y_n \rightarrow Y_\infty$ as $n \rightarrow \infty$ in mean and $\mathbb{E}(Y_\infty) = 1$.
2. If $A = 0$ then $Y_\infty = 0$ with probability 1.

Proof

Note that $a_n > 0$ for $n \in \mathbb{N}_+$ since X_n is nonnegative and $\mathbb{P}(X_n > 0) > 0$. Also, since $x \mapsto \sqrt{x}$ is concave on $(0, \infty)$ it follows from Jensen's inequality that

$$a_n = \mathbb{E}(\sqrt{X_n}) \leq \sqrt{\mathbb{E}(X_n)} = 1 \quad (17.5.16)$$

Let $A_n = \prod_{i=1}^n a_i$ for $n \in \mathbb{N}$. Since $a_n \in (0, 1]$ for $n \in \mathbb{N}_+$, it follows that $A_n \in (0, 1]$ for $n \in \mathbb{N}$ and that A_n is decreasing in $n \in \mathbb{N}$ with limit $A = \prod_{i=1}^{\infty} a_i \in [0, 1]$. Next let $Z_n = \prod_{i=1}^n \sqrt{X_i}/a_i$ for $n \in \mathbb{N}$, so that $\mathbf{Z} = \{Z_n : n \in \mathbb{N}\}$ is the partial product process associated with $\{\sqrt{X_n}/a_n : n \in \mathbb{N}\}$. Since $\mathbb{E}(\sqrt{X_n}/a_n) = 1$ for $n \in \mathbb{N}_+$, the process \mathbf{Z} is also a nonnegative martingale, so there exists a random variable Z_∞ such that $Z_n \rightarrow Z_\infty$ as $n \rightarrow \infty$ with probability 1. Note that $Z_n^2 = Y_n/A_n^2$, $Y_n = A_n^2 Z_n^2$, and $Y_n \leq Z_n^2$ for $n \in \mathbb{N}$.

1. Suppose that $A > 0$. Since the martingale \mathbf{Y} has mean 1,

$$\mathbb{E}(Z_n^2) = \mathbb{E}(Y_n/A_n^2) = 1/A_n^2 \leq 1/A^2 < \infty, \quad n \in \mathbb{N} \quad (17.5.17)$$

Let $W_n = \max\{Z_k : k \in \{0, 1, \dots, n\}\}$ for $n \in \mathbb{N}$ so that $\mathbf{W} = \{W_n : n \in \mathbb{N}\}$ is the maximal process associated with \mathbf{Z} . Also, let $W_\infty = \sup\{Z_k : k \in \mathbb{N}\}$ and note that $W_n \uparrow W_\infty$ as $n \rightarrow \infty$. By the \mathcal{L}_2 maximal inequality,

$$\mathbb{E}(W_n^2) \leq 4\mathbb{E}(Z_n^2) \leq 4/A^2, \quad n \in \mathbb{N} \quad (17.5.18)$$

By the monotone convergence theorem, $\mathbb{E}(W_\infty^2) = \lim_{n \rightarrow \infty} \mathbb{E}(W_n^2) \leq 4/A^2$. Since $x \rightarrow x^2$ is strictly increasing on $[0, \infty)$, $W_\infty^2 = \sup\{Z_n^2 : n \in \mathbb{N}\}$ and so $Y_n \leq W_\infty^2$ for $n \in \mathbb{N}$. Since $\mathbb{E}(W_\infty^2) < \infty$, it follows that the martingale \mathbf{Y} is uniformly integrable. Hence by the [third martingale convergence theorem](#) above, $Y_n \rightarrow Y_\infty$ is mean. Since convergence in mean implies that the means converge, $\mathbb{E}(Y_\infty) = \lim_{n \rightarrow \infty} \mathbb{E}(Y_n) = 1$.

2. Suppose that $A = 0$. Then $Y_n = A_n^2 Z_n^2 \rightarrow 0 \cdot Z_\infty^2 = 0$ as $n \rightarrow \infty$ with probability 1. Note that in this case, the convergence is not in mean, and trivially $\mathbb{E}(Y_\infty) = 0$.

Density Functions

This discussion continues the one on density functions in the Introduction. To review, we start with our probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a filtration $\mathfrak{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ in discrete time. Recall again that $\mathcal{F}_\infty = \sigma(\bigcup_{n=0}^{\infty} \mathcal{F}_n)$. Suppose now that μ is a finite measure on the sample space (Ω, \mathcal{F}) . For each $n \in \mathbb{N} \cup \{\infty\}$, the restriction of μ to \mathcal{F}_n is a measure on (Ω, \mathcal{F}_n) and similarly the restriction of \mathbb{P} to \mathcal{F}_n is a probability measure on (Ω, \mathcal{F}_n) . To save notation and terminology, we will refer to these as μ and \mathbb{P} on \mathcal{F}_n , respectively. Suppose now that μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F}_n for each $n \in \mathbb{N}$. By the Radon-Nikodym theorem, μ has a density function (or Radon-Nikodym derivative) $X_n : \Omega \rightarrow \mathbb{R}$ with respect to \mathbb{P} on \mathcal{F}_n for each $n \in \mathbb{N}$. The theorem and the derivative are named for Johann Radon and Otto Nikodym. In the Introduction we showed that $\mathbf{X} = \{X_n : n \in \mathbb{N}\}$ is a martingale with respect to \mathfrak{F} . Here is the convergence result:

There exists a random variable X_∞ such that $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$ with probability 1.

1. If μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F}_∞ then X_∞ is a density function of μ with respect to \mathbb{P} on \mathcal{F}_∞ .
2. If μ and \mathbb{P} are mutually singular on \mathcal{F}_∞ then $X_\infty = 0$ with probability 1.

Proof

Again, as shown in the Introduction, \mathbf{X} is a martingale with respect to \mathfrak{F} . Moreover, $\mathbb{E}(|X_n|) = \|\mu\|$ (the total variation of μ) for each $n \in \mathbb{N}$. Since μ is a finite measure, $\|\mu\| < \infty$ so the first [martingale convergence theorem](#) applies. Hence there exists a random variable X_∞ , measurable with respect to \mathcal{F}_∞ , such that $X_n \rightarrow X_\infty$ as $n \rightarrow \infty$.

1. If μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F}_∞ , then μ has a density function Y_∞ with respect to \mathbb{P} on \mathcal{F}_∞ . Our goal is to show that $X_\infty = Y_\infty$ with probability 1. By definition, Y_∞ is measurable with respect to \mathcal{F}_∞ and

$$\int_A Y_\infty d\mathbb{P} = \mathbb{E}(Y_\infty; A) = \mu(A), \quad A \in \mathcal{F}_\infty \quad (17.5.19)$$

Suppose now that $n \in \mathbb{N}$ and $A \in \mathcal{F}_n$. Then again by definition, $\mathbb{E}(X_n; A) = \mu(A)$. But $A \in \mathcal{F}_\infty$ also, so $\mathbb{E}(Y_\infty; A) = \mu(A)$. So to summarize, X_n is \mathcal{F}_n -measurable and $\mathbb{E}(X_n; A) = \mathbb{E}(Y_\infty; A)$ for each $A \in \mathcal{F}_n$. By definition, this means that $X_n = \mathbb{E}(Y_\infty | \mathcal{F}_n)$, so X is the [Doob martingale](#) associated with Y_∞ . Letting $n \rightarrow \infty$ and using the result [above](#) gives $X_\infty = \mathbb{E}(Y_\infty | \mathcal{F}_\infty) = Y_\infty$ (with probability 1, of course).

2. Suppose that μ and \mathbb{P} are mutually singular on \mathcal{F}_∞ . Assume first that μ is a positive measure, so that X_n is nonnegative for $n \in \mathbb{N} \cup \{\infty\}$. By the definition of mutual singularity, there exists $B \in \mathcal{F}_\infty$ such that $\mu_\infty(B) = 0$ and $\mathbb{P}_\infty(B^c) = 0$, so that $\mathbb{P}(B) = 1$. Our goal is to show that $\mathbb{E}(X_\infty; A) \leq \mu(A)$ for every $A \in \mathcal{F}_\infty$. Towards that end, let

$$\mathcal{M} = \{A \in \mathcal{F}_\infty : \mathbb{E}(X_\infty; A) \leq \mu(A)\} \quad (17.5.20)$$

Suppose that $A \in \bigcup_{k=0}^\infty \mathcal{F}_k$, so that $A \in \mathcal{F}_k$ for some $k \in \mathbb{N}$. Then $A \in \mathcal{F}_n$ for all $n \geq k$ and therefore $\mathbb{E}(X_n; A) = \mu(A)$ for all $n \geq k$. By Fatou's lemmas,

$$\mathbb{E}(X_\infty; A) \leq \liminf_{n \rightarrow \infty} \mathbb{E}(X_n; A) \leq \mu(A) \quad (17.5.21)$$

so $A \in \mathcal{M}$. Next, suppose that $\{A_n : n \in \mathbb{N}\}$ is an increasing or decreasing sequence in \mathcal{M} , and let $A_\infty = \lim_{n \rightarrow \infty} A_n$ (the union in the first case and the intersection in the second case). Then $\mathbb{E}(X_\infty; A_n) \leq \mu(A_n)$ for each $n \in \mathbb{N}$. By the continuity theorems, $\mathbb{E}(X_\infty; A_n) \rightarrow \mathbb{E}(X_\infty; A_\infty)$ and $\mu(A_n) \rightarrow \mu(A_\infty)$ as $n \rightarrow \infty$. Therefore $\mathbb{E}(X_\infty; A_\infty) \leq \mu(A_\infty)$ and so $A_\infty \in \mathcal{M}$. It follows that \mathcal{M} is a monotone class. Since \mathcal{M} contains the algebra $\bigcup_{n=0}^\infty \mathcal{F}_n$, it then follows from the monotone class theorem that $\mathcal{F}_\infty \subseteq \mathcal{M}$. In particular $B \in \mathcal{M}$, so $\mathbb{E}(X_\infty) = \mathbb{E}(X_\infty; B) \leq \mu(B) = 0$ and therefore $X_\infty = 0$ with probability 1. If μ is a general finite measure, then by the Jordan decomposition theorem, μ can be written uniquely in the form $\mu = \mu^+ - \mu^-$ where μ^+ and μ^- are finite positive measures. Moreover, X_n^+ is the density function of μ^+ on \mathcal{F}_n and X_n^- is the density function of μ^- on \mathcal{F}_n . By the first part of the proof, $X^+ = 0$, $X^- = 0$, and also $X = 0$, all with probability 1.

The martingale approach can be used to give a probabilistic proof of the Radon-Nikodym theorem, at least in certain cases. We start with a sample set Ω . Suppose that $\mathcal{A}_n = \{A_i^n : i \in I_n\}$ is a countable partition of Ω for each $n \in \mathbb{N}$. Thus I_n is countable, $A_i^n \cap A_j^n = \emptyset$ for distinct $i, j \in I_n$, and $\bigcup_{i \in I_n} A_i^n = \Omega$. Suppose also that \mathcal{A}_{n+1} refines \mathcal{A}_n for each $n \in \mathbb{N}$ in the sense that A_i^n is a union of sets in \mathcal{A}_{n+1} for each $i \in I_n$. Let $\mathcal{F}_n = \sigma(\mathcal{A}_n)$. Thus \mathcal{F}_n is generated by a countable partition, and so the sets in \mathcal{F}_n are of the form $\bigcup_{j \in J} A_j^n$ where $J \subseteq I_n$. Moreover, by the refinement property $\mathcal{F}_n \subseteq \mathcal{F}_{n+1}$ for $n \in \mathbb{N}$, so that $\mathcal{F} = \{\mathcal{F}_n : n \in \mathbb{N}\}$ is a filtration. Let $\mathcal{F} = \mathcal{F}_\infty = \sigma(\bigcup_{n=0}^\infty \mathcal{F}_n) = \sigma(\bigcup_{n=0}^\infty \mathcal{A}_n)$, so that our sample space is (Ω, \mathcal{F}) . Finally, suppose that \mathbb{P} is a probability measure on (Ω, \mathcal{F}) with the property that $\mathbb{P}(A_i^n) > 0$ for $n \in \mathbb{N}$ and $i \in I_n$. We now have a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Interesting probability spaces that occur in applications are of this form, so the setting is not as specialized as you might think.

Suppose now that μ a finite measure on (Ω, \mathcal{F}) . From our assumptions, the only null set for \mathbb{P} on \mathcal{F}_n is \emptyset , so μ is automatically absolutely continuous with respect to \mathbb{P} on \mathcal{F}_n . Moreover, for $n \in \mathbb{N}$, we can give the density function of μ with respect to \mathbb{P} on \mathcal{F}_n explicitly:

The density function of μ with respect to \mathbb{P} on \mathcal{F}_n is the random variable X_n whose value on A_i^n is $\mu(A_i^n)/\mathbb{P}(A_i^n)$ for each $i \in I_n$. Equivalently,

$$X_n = \sum_{i \in I_n} \frac{\mu(A_i^n)}{\mathbb{P}(A_i^n)} \mathbf{1}(A_i^n) \quad (17.5.22)$$

Proof

We need to show that $\mu(A) = \mathbb{E}(X_n; A)$ for each $A \in \mathcal{F}_n$. So suppose $A = \bigcup_{j \in J} A_j^n$ where $J \subseteq I_n$. Then

$$\mathbb{E}(X_n; A) = \sum_{j \in J} \mathbb{E}(X_n; A_j^n) = \sum_{j \in J} \frac{\mu(A_j^n)}{\mathbb{P}(A_j^n)} \mathbb{P}(A_j^n) = \sum_{j \in J} \mu(A_j^n) = \mu(A) \quad (17.5.23)$$

By our theorem above, there exists a random variable X such that $X_n \rightarrow X$ as $n \rightarrow \infty$ with probability 1. If μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F} , then X is a density function of μ with respect to \mathbb{P} on \mathcal{F} . The point is that we have given a more or less explicit construction of the density.

For a concrete example, consider $\Omega = [0, 1)$. For $n \in \mathbb{N}$, let

$$\mathcal{A}_n = \left\{ \left[\frac{j}{2^n}, \frac{j+1}{2^n} \right) : j \in \{0, 1, \dots, 2^n - 1\} \right\} \quad (17.5.24)$$

This is the partition of $[0, 1)$ into 2^n subintervals of equal length $1/2^n$, based on the *dyadic rationals* (or *binary rationals*) of rank n or less. Note that every interval in \mathcal{A}_n is the union of two adjacent intervals in \mathcal{A}_{n+1} , so the refinement property holds. Let \mathbb{P} be ordinary Lebesgue measure on $[0, 1)$ so that $\mathbb{P}(A_i^n) = 1/2^n$ for $n \in \mathbb{N}$ and $i \in \{0, 1, \dots, 2^n - 1\}$. As above, let $\mathcal{F}_n = \sigma(\mathcal{A}_n)$ and $\mathcal{F} = \sigma(\bigcup_{n=0}^{\infty} \mathcal{F}_n) = \sigma(\bigcup_{n=0}^{\infty} \mathcal{A}_n)$. The dyadic rationals are dense in $[0, 1)$, so \mathcal{F} is the ordinary Borel σ -algebra on $[0, 1)$. Thus our probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is simply $[0, 1)$ with the usual Euclidean structures. If μ is a finite measure on $([0, 1), \mathcal{F})$ then the density function of μ on \mathcal{F}_n is the random variable X_n whose value on the interval $[j/2^n, (j+1)/2^n)$ is $2^n \mu[j/2^n, (j+1)/2^n)$. If μ is absolutely continuous with respect to \mathbb{P} on \mathcal{F} (so absolutely continuous in the usual sense), then a density function of μ is $X = \lim_{n \rightarrow \infty} X_n$.

This page titled [17.5: Convergence](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

17.6: Backwards Martingales

Basic Theory

A *backwards martingale* is a stochastic process that satisfies the martingale property reversed in time, in a certain sense. In some ways, backward martingales are simpler than their forward counterparts, and in particular, satisfy a convergence theorem similar to the convergence theorem for ordinary martingales. The importance of backward martingales stems from their numerous applications. In particular, some of the fundamental theorems of classical probability can be formulated in terms of backward martingales.

Definitions

As usual, we start with a stochastic process $\mathbf{Y} = \{Y_t : t \in T\}$ on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, having state space \mathbb{R} , and where the index set T (representing time) is either \mathbb{N} (discrete time) or $[0, \infty)$ (continuous time). So to review what all this means, Ω is the sample space, \mathcal{F} the σ -algebra of events, \mathbb{P} the probability measure on (Ω, \mathcal{F}) , and Y_t is a random variable with values in \mathbb{R} for each $t \in T$. But at this point our formulation diverges. Suppose that \mathcal{G}_t is a sub σ -algebra of \mathcal{F} for each $t \in T$, and that $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ is *decreasing* so that if $s, t \in T$ with $s \leq t$ then $\mathcal{G}_t \subseteq \mathcal{G}_s$. Let $\mathcal{G}_\infty = \bigcap_{t \in T} \mathcal{G}_t$. We assume that Y_t is measurable with respect to \mathcal{G}_t and that $\mathbb{E}(|Y_t|) < \infty$ for each $t \in T$.

The process $\mathbf{Y} = \{Y_t : t \in T\}$ is a *backwards martingale* (or *reversed martingale*) with respect to $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ if $\mathbb{E}(Y_s | \mathcal{G}_t) = Y_t$ for all $s, t \in T$ with $s \leq t$.

A backwards martingale can be formulated as an ordinary martingale by using negative times as the indices. Let $T^- = \{-t : t \in T\}$, so that if $T = \mathbb{N}$ (the discrete case) then T^- is the set of non-positive integers, and if $T = [0, \infty)$ (the continuous case) then $T^- = (-\infty, 0]$. Recall also that the standard martingale definitions make sense for any totally ordered index set.

Suppose again that $\mathbf{Y} = \{Y_t : t \in T\}$ is a backwards martingale with respect to $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$. Let $X_t = Y_{-t}$ and $\mathcal{F}_t = \mathcal{G}_{-t}$ for $t \in T^-$. Then $\mathbf{X} = \{X_t : t \in T^-\}$ is a martingale with respect to $\mathfrak{F} = \{\mathcal{F}_t : t \in T^-\}$.

Proof

Since \mathfrak{G} is a decreasing family of sub σ -algebras of \mathcal{F} , the collection \mathfrak{F} is an increasing family of sub σ -algebras of \mathcal{F} , and hence is a filtration. Next, $X_t = Y_{-t}$ is measurable with respect to $\mathcal{G}_{-t} = \mathcal{F}_t$ for $t \in T^-$, so \mathbf{X} is adapted to \mathfrak{F} . Finally, if $s, t \in T^-$ with $s \leq t$ then $-t \leq -s$ so

$$\mathbb{E}(X_t | \mathcal{F}_s) = \mathbb{E}(Y_{-t} | \mathcal{G}_{-s}) = Y_{-s} = X_s \quad (17.6.1)$$

Most authors define backwards martingales with negative indices, as above, in the first place. There are good reasons for doing so, since some of the fundamental theorems of martingales apply immediately to backwards martingales. However, for the *applications* of backwards martingales, this notation is artificial and clunky, so for the most part, we will use our original definition. The next result is another way to view a backwards martingale as an ordinary martingale. This one preserves nonnegative time, but introduces a finite time horizon. For $t \in T$, let $T_t = \{s \in T : s \leq t\}$, a notation we have used often before.

Suppose again that $\mathbf{Y} = \{Y_t : t \in T\}$ is a backwards martingale with respect to $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$. Fix $t \in T$ and define $X_s^t = Y_{t-s}$ and $\mathcal{F}_s^t = \mathcal{G}_{t-s}$ for $s \in T_t$. Then $\mathbf{X}^t = \{X_s^t : s \in T_t\}$ is a martingale relative to $\mathfrak{F}^t = \{\mathcal{F}_s^t : s \in T_t\}$.

Proof

The proof is essentially the same as for the previous result. Since \mathfrak{G} is a decreasing family of sub σ -algebras of \mathcal{F} , the collection \mathfrak{F}^t is an increasing family of sub σ -algebras of \mathcal{F} , and hence is a filtration. Next, $X_s^t = Y_{t-s}$ is measurable with respect to $\mathcal{G}_{t-s} = \mathcal{F}_s^t$ for $s \in T_t$, so \mathbf{X}^t is adapted to \mathfrak{F}^t . Finally, if $r, s \in T_t$ with $r \leq s$ then $t-s \leq t-r$ so

$$\mathbb{E}(X_s^t | \mathcal{F}_r^t) = \mathbb{E}(Y_{t-s} | \mathcal{G}_{t-r}) = Y_{t-r} = X_r^t \quad (17.6.2)$$

Properties

Backwards martingales satisfy a simple and important property.

Suppose that $\mathbf{Y} = \{Y_t : t \in T\}$ is a backwards martingale with respect to $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$. Then $Y_t = \mathbb{E}(Y_0 | \mathcal{G}_t)$ for $t \in T$ and hence \mathbf{Y} is uniformly integrable.

Proof

The fact that $Y_t = \mathbb{E}(Y_0 | \mathcal{G}_t)$ for $t \in T$ follows directly from the definition of a backwards martingale. Since we have assumed that $\mathbb{E}(|Y_0|) < \infty$, it follows from a basic property that \mathbf{Y} is uniformly integrable.

Here is the *Doob backwards martingale*, analogous to the ordinary Doob martingale, and of course named for Joseph Doob. In a sense, this is the converse to the previous result.

Suppose that Y is a random variable on our probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with $\mathbb{E}(|Y|) < \infty$, and that $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$ is a decreasing family of sub σ -algebras of \mathcal{F} , as above. Let $Y_t = \mathbb{E}(Y | \mathcal{G}_t)$ for $t \in T$. Then $\mathbf{Y} = \{Y_t : t \in T\}$ is a backwards martingale with respect to \mathfrak{G} .

Proof

By definition, $Y_t = \mathbb{E}(Y | \mathcal{G}_t)$ is measurable with respect to \mathcal{G}_t . Also,

$$\mathbb{E}(|Y_t|) = \mathbb{E}[\mathbb{E}(|Y| | \mathcal{G}_t)] \leq \mathbb{E}[\mathbb{E}(|Y| | \mathcal{G}_t)] = \mathbb{E}(|Y|) < \infty, \quad t \in T \quad (17.6.3)$$

Next, suppose that $s, t \in T$ with $s \leq t$. Then $\mathcal{G}_t \subseteq \mathcal{G}_s$ so by the tower property of conditional expected value,

$$\mathbb{E}(Y_s | \mathcal{G}_t) = \mathbb{E}[\mathbb{E}(Y | \mathcal{G}_s) | \mathcal{G}_t] = \mathbb{E}(Y | \mathcal{G}_t) = Y_t \quad (17.6.4)$$

The convergence theorems are the most important results for the applications of backwards martingales. Recall once again that for $k \in [1, \infty)$, the k -norm of a real-valued random variable X is

$$\|X\|_k = \left[\mathbb{E}(|X|^k) \right]^{1/k} \quad (17.6.5)$$

and the normed vector space \mathcal{L}_k consists of all X with $\|X\|_k < \infty$. Convergence in the space \mathcal{L}_1 is also referred to as *convergence in mean*, and convergence in the space \mathcal{L}_2 is referred to as *convergence in mean square*. Here is the primary *backwards martingale convergence theorem*:

Suppose again that $\mathbf{Y} = \{Y_t : t \in T\}$ is a backwards martingale with respect to $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$. Then there exists a random variable Y_∞ such that

1. $Y_t \rightarrow Y_\infty$ as $t \rightarrow \infty$ with probability 1.
2. $Y_t \rightarrow Y_\infty$ as $t \rightarrow \infty$ in mean.
3. $Y_\infty = \mathbb{E}(Y_0 | \mathcal{G}_\infty)$.

Proof

The proof is essentially the same as the ordinary martingale convergence theorem if we use the martingale constructed from \mathbf{Y} above. So, fix $t \in T$ and let $T_t = \{s \in T : s \leq t\}$. Let $X_s^t = Y_{t-s}$ and $\mathcal{F}_s^t = \mathcal{G}_{t-s}$ for $s \in T_t$, so that $\mathbf{X}^t = \{X_s^t : s \in T_t\}$ is a martingale relative to $\mathfrak{F}^t = \{\mathcal{F}_s^t : s \in T_t\}$. Now, for $a, b \in \mathbb{R}$ with $a < b$, let $U_t(a, b)$ denote the number of up-crossings of $[a, b]$ by \mathbf{X}^t on T_t . Note that $U_t(a, b)$ is also the number of down-crossings of $[a, b]$ by \mathbf{Y} on T_t . By the up-crossing inequality applied to the martingale \mathbf{X}^t ,

$$\mathbb{E}[U_t(a, b)] \leq \frac{1}{b-a} [\mathbb{E}(|X_t|) + |a|] = \frac{1}{b-a} [\mathbb{E}(|Y_0|) + |a|] \quad (17.6.6)$$

Now let $U_\infty(a, b)$ denote the number of down-crossings of $[a, b]$ by \mathbf{Y} on all of T . Since $U_t \uparrow U_\infty$ as $t \rightarrow \infty$ it follows from the monotone convergence theorem that

$$\mathbb{E}[U_\infty(a, b)] \leq \frac{1}{b-a} [\mathbb{E}(|Y_0|) + |a|] \quad (17.6.7)$$

Hence with probability 1, $U_\infty(a, b) < \infty$ for every $a, b \in \mathbb{Q}$ with $a < b$. By the characterization of convergence in terms of down-crossings (completely analogous to the one for up-crossings), there exists a random variable Y_∞ with values in $\mathbb{R}^* = \mathbb{R} \cup \{-\infty, \infty\}$ such that $Y_t \rightarrow Y_\infty$ as $t \rightarrow \infty$. By Fatou's lemma,

$$\mathbb{E}(|Y_\infty|) \leq \liminf_{t \rightarrow \infty} \mathbb{E}(|Y_t|) \leq \mathbb{E}(|Y_0|) < \infty \quad (17.6.8)$$

In particular, $\mathbb{P}(Y_\infty \in \mathbb{R}) = 1$. Since \mathbf{Y} is uniformly integrable, and $Y_\infty \in \mathcal{L}_1$, it follows that $Y_t \rightarrow Y_\infty$ as $t \rightarrow \infty$ in \mathcal{L}_1 also.

It remains to show that $Y_\infty = \mathbb{E}(Y_0 | \mathcal{G}_\infty)$. Let $A \in \mathcal{G}_\infty$. Then $A \in \mathcal{G}_t$ for every $t \in T$. Since $Y_t = \mathbb{E}(Y_0 | \mathcal{G}_t)$ it follows by definition that $\mathbb{E}(Y_t; A) = \mathbb{E}(Y_0; A)$ for every $t \in T$. Letting $t \rightarrow \infty$ and using the dominated convergence theorem, gives $\mathbb{E}(Y_\infty; A) = \mathbb{E}(Y_0; A)$. Hence $Y_\infty = \mathbb{E}(Y_0 | \mathcal{G}_\infty)$.

As a simple extension of the last result, if $Y_0 \in \mathcal{L}_k$ for some $k \in [1, \infty)$ then the convergence is in \mathcal{L}_k also.

Suppose again that $\mathbf{Y} = \{Y_t : t \in T\}$ is a backwards martingale relative to $\mathfrak{G} = \{\mathcal{G}_t : t \in T\}$. If $Y_0 \in \mathcal{L}_k$ for some $k \in [1, \infty)$ then $Y_t \rightarrow Y_\infty$ as $t \rightarrow \infty$ in \mathcal{L}_k .

Proof

The previous result applies, of course, so we know that there exists a random variable $Y_\infty \in \mathcal{L}_1$ such that $Y_t \rightarrow Y_\infty$ as $t \rightarrow \infty$ with probability 1 and in \mathcal{L}_1 . The function $x \mapsto |x|^k$ is convex on \mathbb{R} so by Jensen's inequality for conditional expected value,

$$\mathbb{E}(|Y_t|^k) = \mathbb{E}[\mathbb{E}(|Y_t|^k | \mathcal{G}_t)] \leq \mathbb{E}[\mathbb{E}(|Y_0|^k | \mathcal{G}_t)] = \mathbb{E}(|Y_0|^k) < \infty \quad (17.6.9)$$

so $Y_t \in \mathcal{L}_k$ for every $t \in T$. By Fatou's lemma,

$$\mathbb{E}(|Y_\infty|^k) \leq \liminf_{t \rightarrow \infty} \mathbb{E}(|Y_t|^k) \leq \mathbb{E}(|Y_0|^k) < \infty \quad (17.6.10)$$

so $Y_\infty \in \mathcal{L}_k$ also. Next, since $Y_t = \mathbb{E}(Y_0 | \mathcal{G}_t)$ and Y_∞ is measurable with respect to \mathcal{G}_t , we can use Jensen's inequality again to get

$$|Y_t - Y_\infty|^k = |\mathbb{E}(Y_0 - Y_\infty | \mathcal{G}_t)|^k \leq \mathbb{E}(|Y_0 - Y_\infty|^k | \mathcal{G}_t) \quad (17.6.11)$$

It follows that the family of random variables $\{|Y_t - Y_\infty|^k : t \in T\}$ is uniformly integrable, and hence $\mathbb{E}(|Y_t - Y_\infty|^k) \rightarrow 0$ as $t \rightarrow \infty$.

Applications

The Strong Law of Large Numbers

The strong law of large numbers is one of the fundamental theorems of classical probability. Our previous proof required that the underlying distribution have finite variance. Here we present an elegant proof using backwards martingales that does not require this extra assumption. So, suppose that $\mathbf{X} = \{X_n : n \in \mathbb{N}_+\}$ is a sequence of independent, identically distributed random variables with common mean $\mu \in \mathbb{R}$. In statistical terms, \mathbf{X} corresponds to *sampling* from the underlying distribution. Next let

$$Y_n = \sum_{i=1}^n X_i, \quad n \in \mathbb{N} \quad (17.6.12)$$

so that $\mathbf{Y} = \{Y_n : n \in \mathbb{N}\}$ is the partial sum process associated with \mathbf{X} . Recall that the sequence \mathbf{Y} is also a discrete-time random walk. Finally, let $M_n = Y_n/n$ for $n \in \mathbb{N}_+$ so that $\mathbf{M} = \{M_n : n \in \mathbb{N}_+\}$ is the sequence of *sample means*.

The law of large numbers

1. $M_n \rightarrow \mu$ as $n \rightarrow \infty$ with probability 1.
2. $M_n \rightarrow \mu$ as $n \rightarrow \infty$ in mean.

Proof

As usual, let $(\Omega, \mathcal{F}, \mathbb{P})$ denote the underlying probability space. Also, equalities involving random variables (and particularly conditional expected values) are assumed to hold with probability 1. Now, for $n \in \mathbb{N}$, let

$$\mathcal{G}_n = \sigma\{Y_n, Y_{n+1}, Y_{n+2}, \dots\} = \sigma\{Y_n, X_{n+1}, X_{n+2}, \dots\} \quad (17.6.13)$$

so that $\mathcal{G} = \{\mathcal{G}_n : n \in \mathbb{N}\}$ is a decreasing family of sub σ -algebras of \mathcal{F} . The core of the proof is to show that \mathbf{M} is a backwards martingale relative to \mathcal{G} . Let $n \in \mathbb{N}_+$. Clearly M_n is measurable with respect to \mathcal{G}_n . By independence, $\mathbb{E}(X_i | \mathcal{G}_n) = \mathbb{E}(X_i | Y_n)$ for $i \in \{1, 2, \dots, n\}$. By symmetry (the sequence \mathbf{X} is exchangeable), $\mathbb{E}(X_i | Y_n) = \mathbb{E}(X_j | Y_n)$ for $i, j \in \{1, 2, \dots, n\}$. Hence for $i \in \{1, 2, \dots, n\}$

$$Y_n = \mathbb{E}(Y_n | \mathcal{G}_n) = \sum_{j=1}^n \mathbb{E}(X_j | \mathcal{G}_n) = \sum_{j=1}^n \mathbb{E}(X_i | \mathcal{G}_n) = n\mathbb{E}(X_i | \mathcal{G}_n) \quad (17.6.14)$$

so that $\mathbb{E}(X_i | \mathcal{G}_n) = Y_n/n = M_n$ for each $i \in \{1, 2, \dots, n\}$. Next,

$$\mathbb{E}(Y_n | \mathcal{G}_{n+1}) = \mathbb{E}(Y_{n+1} - X_{n+1} | \mathcal{G}_{n+1}) = Y_{n+1} - \mathbb{E}(X_{n+1} | \mathcal{G}_{n+1}) = Y_{n+1} - \frac{1}{n+1} Y_{n+1} = \frac{n}{n+1} Y_{n+1} \quad (17.6.15)$$

Dividing by n gives $\mathbb{E}(M_n | \mathcal{G}_{n+1}) = M_{n+1}$ and hence \mathbf{M} is a backwards martingale with respect to \mathcal{G} . From the backwards martingale convergence theorem, there exists M_∞ such that $M_n \rightarrow M_\infty$ as $n \rightarrow \infty$ with probability 1 and in mean. Next, for $n, k \in \mathbb{N}_+$ simple algebra gives

$$M_{n+k} = \frac{1}{n+k} \sum_{i=1}^k X_i + \frac{n}{n+k} \frac{1}{n} \sum_{i=k+1}^{k+n} X_i \quad (17.6.16)$$

Letting $n \rightarrow \infty$ then shows that

$$M_\infty = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=k+1}^{k+n} X_i \quad (17.6.17)$$

for every $k \in \mathbb{N}_+$. Hence M_∞ is a tail random variable for the IID sequence \mathbf{X} . From the Kolmogorov 0-1 law, M_∞ must be a constant. Finally, convergence in mean implies that the means converge, and since $\mathbb{E}(M_n) = \mu$ for each n , it follows that $M_\infty = \mu$.

Exchangeable Variables

We start with a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and another measurable space (S, \mathcal{S}) . Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is a sequence of random variables each taking values in S . Recall that \mathbf{X} is *exchangeable* if for every $n \in \mathbb{N}$, every permutation of (X_1, X_2, \dots, X_n) has the same distribution on (S^n, \mathcal{S}^n) (where \mathcal{S}^n is the n -fold product σ -algebra). Clearly if \mathbf{X} is a sequence of independent, identically distributed variables, then \mathbf{X} is exchangeable. Conversely, if \mathbf{X} is exchangeable then the variables are identically distributed (by definition), but are not necessarily independent. The most famous example of a sequence that is exchangeable but not independent is Pólya's urn process, named for George Pólya. On the other hand, *conditionally* independent and identically distributed sequences *are* exchangeable. Thus suppose that (T, \mathcal{T}) is another measurable space and that Θ is a random variable taking values in T .

If \mathbf{X} is conditionally independent and identically distributed given Θ , then \mathbf{X} is exchangeable.

Proof

Implicit in the statement is that the variables in the sequence have a regular conditional distribution μ_Θ given Θ . Then for every $n \in \mathbb{N}_+$, the conditional distribution of every permutation of (X_1, X_2, \dots, X_n) , given Θ , is μ_Θ^n on (S^n, \mathcal{S}^n) , where μ_Θ^n is the n -fold product measure. Unconditionally, the distribution of any permutation is $B \mapsto \mathbb{E}[\mu_\Theta^n(B)]$ for $B \in \mathcal{S}^n$.

Often the setting of this theorem arises when we start with a sequence of independent, identically distributed random variables that are governed by a parametric distribution, and then randomize one of the parameters. In a sense, we can always think of the setting in this way: Imagine that $\theta \in T$ is a parameter for a distribution on S . A special case is the beta-Bernoulli process, in which the success parameter p in sequence of Bernoulli trials is randomized with the beta distribution. On the other hand, Pólya's urn process is an example of an exchangeable sequence that does not at first seem to have anything to do with randomizing parameters. But in fact, we know that Pólya's urn process is a special case of the beta-Bernoulli process. This connection gives a hint of *de Finetti's theorem*, named for Bruno de Finetti, which we consider next. This theorem states any exchangeable sequence of indicator random variables corresponds to randomizing the success parameter in a sequence of Bernoulli trials.

de Finetti's Theorem. Suppose that $\mathbf{X} = (X_1, X_2, \dots)$ is an exchangeable sequence of random variables, each taking values in $\{0, 1\}$. Then there exists a random variable P with values in $[0, 1]$, such that given $P = p \in [0, 1]$, \mathbf{X} is a sequence of Bernoulli trials with success parameter p .

Proof

As usual, we need some notation. First recall the falling power notation $r^{(j)} = r(r-1) \cdots (r-j+1)$ for $r \in \mathbb{R}$ and $j \in \mathbb{N}$. Next for $n \in \mathbb{N}_+$ and $k \in \{0, 1, \dots, n\}$, let

$$B_k^n = \left\{ (x_1, x_2, \dots, x_n) \in \{0, 1\}^n : \sum_{i=1}^n x_i = k \right\} \quad (17.6.18)$$

That is, B_k^n is the set of bit strings of length n with 1 occurring exactly k times. Of course, $\#(B_k^n) = \binom{n}{k} = n^{(k)}/k!$.

Suppose now that $\mathbf{X} = (X_1, X_2, \dots)$ is an exchangeable sequence of variables with values in $\{0, 1\}$. For $n \in \mathbb{N}_+$ let $Y_n = \sum_{i=1}^n X_i$ and $M_n = Y_n/n$. So $\mathbf{Y} = \{Y_n : n \in \mathbb{N}_+\}$ is the partial sum process associated with \mathbf{X} and $\mathbf{M} = \{M_n : n \in \mathbb{N}_+\}$ the sequence of sample means. Let $\mathcal{G}_n = \sigma\{Y_n, Y_{n+1}, \dots\}$ and $\mathcal{G}_\infty = \bigcap_{n=0}^\infty \mathcal{G}_n$. The family of σ -algebras $\mathcal{G} = \{\mathcal{G}_n : n \in \mathbb{N}_+\}$ is decreasing. The key to the proof is to find two backwards martingales and use the [backwards martingale convergence theorem](#).

Let $m \in \mathbb{N}_+$ and $k \in \{0, 1, \dots, m\}$. The crucial insight is that by exchangeability, given $Y_m = k$, the random vector (X_1, X_2, \dots, X_m) is uniformly distributed on B_k^m . So if $n \in \mathbb{N}_+$ and $n \leq m$, the random vector (X_1, X_2, \dots, X_n) , again given $Y_m = k$, fits the hypergeometric model: a sample of size n chosen at random and without replacement from a population of m objects of which k are type 1 and $m-k$ are type 0. Thus, if $j \in \{0, 1, \dots, n\}$ and $(x_1, x_2, \dots, x_n) \in B_j^n$ then

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid Y_m = k) = \frac{k^{(j)}(m-k)^{(n-j)}}{m^{(n)}} \quad (17.6.19)$$

Equivalently,

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid Y_m) = \frac{Y_m^{(j)}(m - Y_m)^{(n-j)}}{m^{(n)}} \quad (17.6.20)$$

Given Y_m , the variables $(Y_{m+1}, Y_{m+2}, \dots)$ give no additional information about the distribution of (X_1, X_2, \dots, X_n) and hence

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid \mathcal{G}_m) = \mathbb{E}[\mathbf{1}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) \mid \mathcal{G}_m] = \frac{Y_m^{(j)}(m - Y_m)^{(n-j)}}{m^{(n)}} \quad (17.6.21)$$

For fixed n, j , and $(x_1, x_2, \dots, x_n) \in B_j^n$, the conditional expected value in the middle of the displayed equation, as a function of m , is a [Doob backward martingale](#) with respect to \mathcal{G} and hence converges to $\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid \mathcal{G}_\infty)$ as $m \rightarrow \infty$.

Next we show that \mathbf{M} is a backwards martingale with respect to \mathcal{G} . Trivially M_n is measurable with respect to \mathcal{G}_n and $\mathbb{E}(M_n) \leq 1$ for each $n \in \mathbb{N}$. Thus we need to show that $\mathbb{E}(M_n \mid \mathcal{G}_m) = M_m$ for $m, n \in \mathbb{N}_+$ with $n \leq m$. From our previous work with (X_1, X_2, \dots, X_n) we know that the conditional distribution of Y_n given $Y_m = k$ is hypergeometric with parameters m, k , and n :

$$\mathbb{P}(Y_n = j \mid Y_m = k) = \binom{n}{j} \frac{k^{(j)}(m - k)^{(n-j)}}{m^{(n)}}, \quad j \in \{0, 1, \dots, n\} \quad (17.6.22)$$

Recall that the mean of the hypergeometric distribution is the sample size times the proportion of type 1 objects in the population. Thus,

$$\mathbb{E}(M_n \mid Y_m = k) = \frac{1}{n} \mathbb{E}(Y_n \mid Y_m = k) = \frac{1}{n} n \frac{k}{m} = \frac{k}{m} \quad (17.6.23)$$

Or equivalently, $\mathbb{E}(M_n \mid Y_m) = Y_m/m = M_m$. Once again, given Y_m , the variables Y_{m+1}, Y_{m+2} give no additional information and so $\mathbb{E}(Y_n \mid \mathcal{G}_m) = Y_m$. Hence \mathbf{M} is a backwards martingale with respect to \mathcal{G} . From the [backwards martingale convergence theorem](#), there exists a random variable P such that $M_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1.

It just remains to connect the dots. Suppose now that $n \in \mathbb{N}_+$ and $j \in \{0, 1, \dots, n\}$ and that $m \in \mathbb{N}_+$ and $k_m \in \{0, 1, \dots, m\}$. From simple calculus, if n and j are fixed and $k_m/m \rightarrow p \in [0, 1]$ as $m \rightarrow \infty$ then

$$\frac{k_m^{(j)}(m - k_m)^{(n-j)}}{m^{(n)}} \rightarrow p^j(1 - p)^{n-j} \text{ as } m \rightarrow \infty \quad (17.6.24)$$

(You may recall that this computation is used in the proof of the convergence of the hypergeometric distribution to the binomial.) Returning to the joint distribution, recall that if $(x_1, x_2, \dots, x_n) \in B_j^n$ then

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid \mathcal{G}_m) = \frac{Y_m^{(j)}(m - Y_m)^{(n-j)}}{m^{(n)}} \quad (17.6.25)$$

Let $m \rightarrow \infty$. Since $Y_m/m \rightarrow P$ as $m \rightarrow \infty$ we get

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid \mathcal{G}_\infty) = P^j(1 - P)^{n-j} \quad (17.6.26)$$

Random variable P is measurable with respect to \mathcal{G}_∞ so

$$\mathbb{P}(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n \mid P) = P^j(1 - P)^{n-j} \text{ as } m \rightarrow \infty \quad (17.6.27)$$

Given $P = p \in [0, 1]$, \mathbf{X} is a sequence of Bernoulli trials with success parameter p .

De Finetti's theorem has been extended to much more general sequences of exchangeable variables. Basically, if $\mathbf{X} = (X_1, X_2, \dots)$ is an exchangeable sequence of random variables, each taking values in a significantly nice measurable space (S, \mathcal{S}) then there exists a random variable Θ such that \mathbf{X} is independent and identically distributed given Θ . In the proof, the result that $M_n \rightarrow P$ as $n \rightarrow \infty$ with probability 1, where $M_n = \frac{1}{n} \sum_{i=1}^n X_i$, is known as *de Finetti's strong law of large numbers*. De Finetti's theorem, and its generalizations are important in Bayesian statistical inference. For an exchangeable sequence of random variables (our observations in a statistical experiment), there is a hidden, random parameter Θ . Given $\Theta = \theta$, the variables are independent and identically distributed. We gain information about Θ by imposing a *prior distribution* on Θ and then updating this, based on our observations and using Baye's theorem, to a *posterior distribution*.

Stated more in terms of distributions, de Finetti's theorem states that the distribution of n distinct variables in the exchangeable sequence is a *mixture* of product measures. That is, if μ_θ is the distribution of a generic X on (S, \mathcal{S}) given $\Theta = \theta$, and ν is the distribution of Θ on (T, \mathcal{T}) , then the distribution of n of the variables on (S^n, \mathcal{S}^n) is

$$B \mapsto \int_T \mu_\theta^n(B) d\nu(\theta) \quad (17.6.28)$$

This page titled [17.6: Backwards Martingales](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

CHAPTER OVERVIEW

18: Brownian Motion

Brownian motion is a stochastic process of great theoretical importance, and as the basic building block of a variety of other processes, of great practical importance as well. In this chapter we study Brownian motion and a number of random processes that can be constructed from Brownian motion. We also study the Ito stochastic integral and the resulting calculus, as well as two remarkable representation theorems involving stochastic integrals.

[18.1: Standard Brownian Motion](#)

[18.2: Brownian Motion with Drift and Scaling](#)

[18.3: The Brownian Bridge](#)

[18.4: Geometric Brownian Motion](#)

This page titled [18: Brownian Motion](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

18.1: Standard Brownian Motion

Basic Theory

History

In 1827, the botanist Robert Brown noticed that tiny particles from pollen, when suspended in water, exhibited continuous but very jittery and erratic motion. In his “miracle year” in 1905, Albert Einstein explained the behavior physically, showing that the particles were constantly being bombarded by the molecules of the water, and thus helping to firmly establish the atomic theory of matter. Brownian motion as a mathematical random process was first constructed in rigorous way by Norbert Wiener in a series of papers starting in 1918. For this reason, the Brownian motion process is also known as the *Wiener process*.

Run the two-dimensional Brownian motion simulation several times in single-step mode to get an idea of what Mr. Brown may have observed under his microscope.

Along with the Bernoulli trials process and the Poisson process, the Brownian motion process is of central importance in probability. Each of these processes is based on a set of idealized assumptions that lead to a rich mathematical theory. In each case also, the process is used as a building block for a number of related random processes that are of great importance in a variety of applications. In particular, Brownian motion and related processes are used in applications ranging from physics to statistics to economics.

Definition

A *standard Brownian motion* is a random process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with state space \mathbb{R} that satisfies the following properties:

1. $X_0 = 0$ (with probability 1).
2. \mathbf{X} has stationary increments. That is, for $s, t \in [0, \infty)$ with $s < t$, the distribution of $X_t - X_s$ is the same as the distribution of X_{t-s} .
3. \mathbf{X} has independent increments. That is, for $t_1, t_2, \dots, t_n \in [0, \infty)$ with $t_1 < t_2 < \dots < t_n$, the random variables $X_{t_1}, X_{t_2} - X_{t_1}, \dots, X_{t_n} - X_{t_{n-1}}$ are independent.
4. X_t is normally distributed with mean 0 and variance t for each $t \in (0, \infty)$.
5. With probability 1, $t \mapsto X_t$ is continuous on $[0, \infty)$.

To understand the assumptions physically, let's take them one at a time.

1. Suppose that we measure the position of a Brownian particle in one dimension, starting at an arbitrary time which we designate as $t = 0$, with the initial position designated as $x = 0$. Then this assumption is satisfied by convention. Indeed, occasionally, it's convenient to relax this assumption and allow X_0 to have other values.
2. This is a statement of *time homogeneity*: the underlying dynamics (namely the jostling of the particle by the molecules of water) do not change over time, so the distribution of the displacement of the particle in a time interval $[s, t]$ depends only on the length of the time interval.
3. This is an idealized assumption that would hold approximately if the time intervals are large compared to the tiny times between collisions of the particle with the molecules.
4. This is another idealized assumption based on the central limit theorem: the position of the particle at time t is the result of a very large number of collisions, each making a very small contribution. The fact that the mean is 0 is a statement of *spatial homogeneity*: the particle is no more or less likely to be jostled to the right than to the left. Next, recall that the assumptions of stationary, independent increments means that $\text{var}(X_t) = \sigma^2 t$ for some positive constant σ^2 . By a change in time scale, we can assume $\sigma^2 = 1$, although we will consider more general Brownian motions in the next section.
5. Finally, the continuity of the sample paths is an essential assumption, since we are modeling the position of a physical particle as a function of time.

Of course, the first question we should ask is whether there *exists* a stochastic process satisfying the definition. Fortunately, the answer is yes, although the proof is complicated.

There exists a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a stochastic process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ on this probability space satisfying the assumptions in the [definition](#).

Proof sketch

The assumptions in the definition lead to a consistent set of finite dimensional distributions (which are given [below](#)). Thus by Kolmogorov existence theorem, there exists a stochastic process $\mathbf{U} = \{U_t : t \in [0, \infty)\}$ that has these finite dimensional distributions. However, \mathbf{U} does not have continuous sample paths, but we can construct from \mathbf{U} an equivalent process that does have continuous sample paths.

First recall that a *binary rational* (or *dyadic rational*) in $[0, \infty)$ is a number of the form $k/2^n$ where $k, n \in \mathbb{N}$. Let \mathbb{D}_+ denote the set of all binary rationals in $[0, \infty)$, and recall that \mathbb{D}_+ is countable but also *dense* in $[0, \infty)$ (that is, if $t \in [0, \infty) \setminus \mathbb{D}_+$ then there exists $t_n \in \mathbb{D}_+$ for $n \in \mathbb{N}_+$ such that $t_n \rightarrow t$ as $n \rightarrow \infty$).

Now, for $n \in \mathbb{N}_+$, let $X_n(t) = U_t$ if t is a binary rational of the form $k/2^n$ for some $k \in \mathbb{N}$. If t is not such a binary rational, define $X_n(t)$ by linear interpolation between the the closest binary rationals of this form on either side of t . Then $X_n(t) \rightarrow U(t)$ as $n \rightarrow \infty$ for every $t \in \mathbb{D}_+$, and with probability 1, the convergence is uniform on $\mathbb{D}_+ \cap [0, T]$ for each $T > 0$. It then follows that \mathbf{U} is continuous on \mathbb{D}_+ with probability 1.

For the last step, let $X_t = \lim_{s \rightarrow t, s \in \mathbb{D}_+} U_s$ for $t \in [0, \infty)$. The limit exists since \mathbf{U} is continuous on \mathbb{D}_+ with probability 1. The process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is continuous on $[0, \infty)$ with probability 1, and has the same finite dimensional distributions as \mathbf{U} .

Run the simulation of the standard Brownian motion process a few times in single-step mode. Note the qualitative behavior of the sample paths. Run the simulation 1000 times and compare the empirical density function and moments of X_t to the true probability density function and moments.

Brownian Motion as a Limit of Random Walks

Clearly the underlying dynamics of the Brownian particle being knocked about by molecules suggests a random walk as a possible model, but with tiny time steps and tiny spatial jumps. Let $\mathbf{X} = (X_0, X_1, X_2, \dots)$ be the symmetric simple random walk. Thus, $X_n = \sum_{i=1}^n U_i$ where $\mathbf{U} = (U_1, U_2, \dots)$ is a sequence of independent variables with $\mathbb{P}(U_i = 1) = \mathbb{P}(U_i = -1) = \frac{1}{2}$ for each $i \in \mathbb{N}_+$. Recall that $\mathbb{E}(X_n) = 0$ and $\text{var}(X_n) = n$ for $n \in \mathbb{N}$. Also, since \mathbf{X} is the partial sum process associated with an IID sequence, \mathbf{X} has stationary, independent increments (but of course in discrete time). Finally, recall that by the central limit theorem, X_n/\sqrt{n} converges to the standard normal distribution as $n \rightarrow \infty$. Now, for $h, d \in (0, \infty)$ the continuous time process

$$\mathbf{X}_{h,d} = \{dX_{\lfloor t/h \rfloor} : t \in [0, \infty)\} \quad (18.1.1)$$

is a jump process with jumps at $\{0, h, 2h, \dots\}$ and with jumps of size $\pm d$. Basically we would like to let $h \downarrow 0$ and $d \downarrow 0$, but this cannot be done arbitrarily. Note that $\mathbb{E}[X_{h,d}(t)] = 0$ but $\text{var}[X_{h,d}(t)] = d^2 \lfloor t/h \rfloor$. Thus, by the central limit theorem, if we take $d = \sqrt{h}$ then the distribution of $X_{h,d}(t)$ will converge to the normal distribution with mean 0 and variance t as $h \downarrow 0$. More generally, we might hope that all of requirements in the [definition](#) are satisfied by the limiting process, and if so, we have a standard Brownian motion.

Run the simulation of the random walk process for increasing values of n . In particular, run the simulation several times with $n = 100$. Compare the qualitative behavior with the standard Brownian motion process. Note that the scaling of the random walk in time and space is effectivly accomplished by scaling the horizontal and vertical axes in the graph window.

Finite Dimensional Distributions

Let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be a standard Brownian motion. It follows from part (d) of the [definition](#) that X_t has probability density function f_t given by

$$f_t(x) = \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{x^2}{2t}\right), \quad x \in \mathbb{R} \quad (18.1.2)$$

This family of density functions determines the finite dimensional distributions of \mathbf{X} .

If $t_1, t_2, \dots, t_n \in (0, \infty)$ with $0 < t_1 < t_2 < \dots < t_n$ then $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ has probability density function f_{t_1, t_2, \dots, t_n} given by

$$f_{t_1, t_2, \dots, t_n}(x_1, x_2, \dots, x_n) = f_{t_1}(x_1) f_{t_2 - t_1}(x_2 - x_1) \cdots f_{t_n - t_{n-1}}(x_n - x_{n-1}), \quad (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (18.1.3)$$

Proof

This follows because \mathbf{X} has stationary, independent increments.

\mathbf{X} is a Gaussian process with mean function mean function $m(t) = 0$ for $t \in [0, \infty)$ and covariance function $c(s, t) = \min\{s, t\}$ for $s, t \in [0, \infty)$.

Proof

The fact that \mathbf{X} is a Gaussian process follows because X_t is normally distributed for each $t \in T$ and \mathbf{X} has stationary, independent increments. The mean function is 0 by assumption. For the covariance function, suppose $s, t \in [0, \infty)$ with $s \leq t$. Since X_s and $X_t - X_s$ are independent, we have

$$\text{cov}(X_s, X_t) = \text{cov}[X_s, X_s + (X_t - X_s)] = \text{var}(X_s) + 0 = s \quad (18.1.4)$$

Recall that for a Gaussian process, the finite dimensional (multivariate normal) distributions are completely determined by the mean function m and the covariance function c . Thus, it follows that a standard Brownian motion is characterized as a continuous Gaussian process with the mean and covariance functions in the last theorem. Note also that

$$\text{cor}(X_s, X_t) = \frac{\min\{s, t\}}{\sqrt{st}} = \sqrt{\frac{\min\{s, t\}}{\max\{s, t\}}}, \quad (s, t) \in [0, \infty)^2 \quad (18.1.5)$$

We can also give the higher moments and the moment generating function for X_t .

For $n \in \mathbb{N}$ and $t \in [0, \infty)$,

1. $\mathbb{E}(X_t^{2n}) = 1 \cdot 3 \cdots (2n-1)t^n = (2n)!t^n / (n!2^n)$
2. $\mathbb{E}(X_t^{2n+1}) = 0$

Proof

These moments follow from standard results, since X_t is normally distributed with mean 0 and variance t .

For $t \in [0, \infty)$, X_t has moment generating function given by

$$\mathbb{E}(e^{uX_t}) = e^{tu/2}, \quad u \in \mathbb{R} \quad (18.1.6)$$

Proof

Again, this is a standard result for the normal distribution.

Simple Transformations

There are several simple transformations that preserve standard Brownian motion and will give us insight into some of its properties. As usual, our starting place is a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Our first result is that reflecting the paths of \mathbf{X} in the line $x = 0$ gives another standard Brownian motion

Let $Y_t = -X_t$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \geq 0\}$ is also a standard Brownian motion.

Proof

Clearly the new process is still a Gaussian process, with mean function $\mathbb{E}(-X_t) = -\mathbb{E}(X_t) = 0$ for $t \in [0, \infty)$ and covariance function $\text{cov}(-X_s, -X_t) = \text{cov}(X_s, X_t) = \min\{s, t\}$ for $(s, t) \in [0, \infty)^2$. Finally, since \mathbf{X} is continuous, so is \mathbf{Y} .

Our next result is related to the Markov property, which we explore in more detail below. If we “restart” Brownian motion at a fixed time s , and shift the origin to X_s , then we have another standard Brownian motion. This means that Brownian motion is both *temporally* and *spatially* homogeneous.

Fix $s \in [0, \infty)$ and define $Y_t = X_{s+t} - X_s$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Proof

Since \mathbf{X} has stationary, independent increments, the process \mathbf{Y} is equivalent in distribution to \mathbf{X} . Clearly also \mathbf{Y} is continuous since \mathbf{X} is.

Our next result is a simple time reversal, but to state this result, we need to restrict the time parameter to a bounded interval of the form $[0, T]$ where $T > 0$. The upper endpoint T is sometimes referred to as a *finite time horizon*. Note that $\{X_t : t \in [0, T]\}$ still satisfies the [definition](#), but with the time parameters restricted to $[0, T]$.

Define $Y_t = X_{T-t} - X_T$ for $0 \leq t \leq T$. Then $\mathbf{Y} = \{Y_t : t \in [0, T]\}$ is also a standard Brownian motion on $[0, T]$.

Proof

\mathbf{Y} is a Gaussian process, since a finite, linear combination of variables from this process reduces to a finite, linear combination of variables from \mathbf{X} . Next, $\mathbb{E}(Y_t) = \mathbb{E}(X_{T-t}) - \mathbb{E}(X_T) = 0$. Next, if $s, t \in [0, T]$ with $s \leq t$ then

$$\text{cov}(Y_s, Y_t) = \text{cov}(X_{T-s} - X_T, X_{T-t} - X_T) = \text{cov}(X_{T-s}, X_{T-t}) - \text{cov}(X_{T-s}, X_T) - \text{cov}(X_T, X_{T-t}) + \text{cov}(X_T, X_T) \quad (18.1.7)$$

$$= (T-t) - (T-s) - (T-t) + T = s \quad (18.1.8)$$

Finally, $t \mapsto Y_t$ is continuous on $[0, T]$ with probability 1, since $t \mapsto X_t$ is continuous on $[0, T]$ with probability 1.

Our next transformation involves scaling \mathbf{X} both temporally and spatially, and is known as *self-similarity*.

Let $a > 0$ and define $Y_t = \frac{1}{a} X_{a^2 t}$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Proof

Once again, \mathbf{Y} is a Gaussian process, since finite, linear combinations of variables in \mathbf{Y} reduce to finite, linear combinations of variables in \mathbf{X} . Next, $\mathbb{E}(Y_t) = a\mathbb{E}(X_{a^2 t}) = 0$ for $t > 0$, and for $s, t > 0$ with $s < t$,

$$\text{cov}(Y_s, Y_t) = \text{cov}\left(\frac{1}{a}X_{a^2s}, \frac{1}{a}X_{a^2t}\right) = \frac{1}{a^2}\text{cov}(X_{a^2s}, X_{a^2t}) = \frac{1}{a^2}a^2s = s \quad (18.1.9)$$

Finally \mathbf{Y} is a continuous process since \mathbf{X} is continuous.

Note that the graph of \mathbf{Y} can be obtained from the graph of \mathbf{X} by scaling the time axis t by a factor of a^2 and scaling the spatial axis x by a factor of a . The fact that the temporal scale factor must be the square of the spatial scale factor is clearly related to Brownian motion as the limit of random walks. Note also that this transformation amounts to “zooming in or out” of the graph of \mathbf{X} and hence Brownian motion has a self-similar, fractal quality, since the graph is unchanged by this transformation. This also suggests that, although continuous, $t \mapsto X_t$ is highly irregular. We consider this in the next subsection.

Our final transformation is referred to as *time inversion*.

Let $Y_0 = 0$ and $Y_t = tX_{1/t}$ for $t > 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Proof

Clearly \mathbf{Y} is a Gaussian process, since finite, linear combinations of variables in \mathbf{Y} reduce to finite, linear combinations of variables in \mathbf{X} .

Next, $\mathbb{E}(Y_t) = t\mathbb{E}(X_{1/t}) = 0$ for $t > 0$, and for $s, t > 0$ with $s < t$,

$$\text{cov}(Y_s, Y_t) = \text{cov}(sX_{1/s}, tX_{1/t}) = st \text{cov}(X_{1/s}, X_{1/t}) = st \frac{1}{t} = s \quad (18.1.10)$$

Since $t \mapsto X_t$ is continuous on $[0, \infty)$ with probability 1, $t \mapsto Y_t$ is continuous on $(0, \infty)$ with probability 1. Thus, all that remains is to show continuity at $t = 0$. Thus we need to show that with probability 1, $tX_{1/t} \rightarrow 0$ as $t \downarrow 0$, or equivalently, $X_s/s \rightarrow 0$ as $s \uparrow \infty$. But this last statement holds by the [law of the iterated logarithm](#), given below.

Irregularity

The [defining properties](#) suggest that standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ cannot be a smooth, differentiable function. Consider the usual difference quotient at t ,

$$\frac{X_{t+h} - X_t}{h} \quad (18.1.11)$$

By the stationary increments property, if $h > 0$, the numerator has the same distribution as X_h , while if $h < 0$, the numerator has the same distribution as $-X_{-h}$, which in turn has the same distribution as X_{-h} . So, in both cases, the difference quotient has the same distribution as $X_{|h|}/h$, and this variable has the normal distribution with mean 0 and variance $|h|/h^2 = 1/|h|$. So the variance of the difference quotient diverges to ∞ as $h \rightarrow 0$, and hence the difference quotient does not even converge in distribution, the weakest form of convergence.

The [temporal-spatial transformation](#) above also suggests that Brownian motion cannot be differentiable. The intuitive meaning of *differentiable* at t is that the function is *locally linear* at t —as we zoom in, the graph near t begins to look like a line (whose slope, of course, is the derivative). But as we zoom in on Brownian motion, (in the sense of the transformation), it always looks the same, and in particular, just as jagged. More formally, if \mathbf{X} is differentiable at t , then so is the transformed process \mathbf{Y} , and the chain rule gives $Y'(t) = aX'(a^2t)$. But \mathbf{Y} is also a standard Brownian motion for every $a > 0$, so something is clearly wrong. While not rigorous, these examples are motivation for the following theorem:

With probability 1, \mathbf{X} is nowhere differentiable on $[0, \infty)$.

Run the simulation of the standard Brownian motion process. Note the continuity but very jagged quality of the sample paths. Of course, the simulation cannot really capture Brownian motion with complete fidelity.

The following theorems give a more precise measure of the irregularity of standard Brownian motion.

Standard Brownian motion \mathbf{X} has Hölder exponent $\frac{1}{2}$. That is, \mathbf{X} is Hölder continuous with exponent α for every $\alpha < \frac{1}{2}$, but is not Hölder continuous with exponent α for any $\alpha > \frac{1}{2}$.

In particular, \mathbf{X} is not Lipschitz continuous, and this shows again that it is not differentiable. The following result states that in terms of Hausdorff dimension, the graph of standard Brownian motion lies midway between a simple curve (dimension 1) and the plane (dimension 2).

The graph of standard Brownian motion has Hausdorff dimension $\frac{3}{2}$.

Yet another indication of the irregularity of Brownian motion is that it has infinite total variation on any interval of positive length.

Suppose that $a, b \in \mathbb{R}$ with $a < b$. Then the total variation of \mathbf{X} on $[a, b]$ is ∞ .

The Markov Property and Stopping Times

As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Recall that a Markov process has the property that the future is independent of the past, given the present state. Because of the stationary, independent increments property, Brownian motion has the property. As a minor note, to view \mathbf{X} as a Markov process, we sometimes need to relax Assumption 1 and let X_0 have an arbitrary value in \mathbb{R} . Let $\mathcal{F}_t = \sigma\{X_s : 0 \leq s \leq t\}$, the sigma-algebra generated by the process up to time $t \in [0, \infty)$. The family of σ -algebras $\mathcal{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is known as a filtration.

Standard Brownian motion is a time-homogeneous Markov process with transition probability density p given by

$$p_t(x, y) = f_t(y - x) = \frac{1}{\sqrt{2\pi t}} \exp\left[-\frac{(y - x)^2}{2t}\right], \quad t \in (0, \infty); x, y \in \mathbb{R} \quad (18.1.12)$$

Proof

Fix $s \in [0, \infty)$. The theorem follows from the fact that the process $\{X_{s+t} - X_s : t \in [0, \infty)\}$ is another standard Brownian motion, as shown [above](#), and is independent of \mathcal{F}_s .

The transition density p satisfies the following *diffusion equations*. The first is known as the *forward equation* and the second as the *backward equation*.

$$\frac{\partial}{\partial t} p_t(x, y) = \frac{1}{2} \frac{\partial^2}{\partial y^2} p_t(x, y) \quad (18.1.13)$$

$$\frac{\partial}{\partial t} p_t(x, y) = \frac{1}{2} \frac{\partial^2}{\partial x^2} p_t(x, y) \quad (18.1.14)$$

Proof

These results follow from standard calculus.

The diffusion equations are so named, because the spatial derivative in the first equation is with respect to y , the state *forward* at time t , while the spatial derivative in the second equation is with respect to x , the state *backward* at time 0.

Recall that a random time τ taking values in $[0, \infty]$ is a *stopping time* with respect to the process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ if $\{\tau \leq t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$. Informally, we can determine whether or not $\tau \leq t$ by observing the process up to time t . An important special case is the first time that our Brownian motion hits a specified state. Thus, for $x \in \mathbb{R}$ let $\tau_x = \inf\{t \in [0, \infty) : X_t = x\}$. The random time τ_x is a stopping time.

For a stopping time τ , we need the σ -algebra of events that can be defined in terms of the process up to the random time τ , analogous to \mathcal{F}_t , the σ -algebra of events that can be defined in terms of the process up to a fixed time t . The appropriate definition is

$$\mathcal{F}_\tau = \{B \in \mathcal{F} : B \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \geq 0\} \quad (18.1.15)$$

See the section on Filtrations and Stopping Times for more information on filtrations, stopping times, and the σ -algebra associated with a stopping time.

The *strong Markov property* is the Markov property generalized to stopping times. Standard Brownian motion \mathbf{X} is also a strong Markov process. The best way to say this is by a generalization of the [temporal and spatial homogeneity result](#) above.

Suppose that τ is a stopping time and define $Y_t = X_{\tau+t} - X_\tau$ for $t \in [0, \infty)$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is a standard Brownian motion and is independent of \mathcal{F}_τ .

The Reflection Principle

Many interesting properties of Brownian motion can be obtained from a clever idea known as the *reflection principle*. As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Let τ be a stopping time for \mathbf{X} . Define

$$W_t = \begin{cases} X_t, & 0 \leq t < \tau \\ 2X_\tau - X_t, & \tau \leq t < \infty \end{cases} \quad (18.1.16)$$

Thus, the graph of $\mathbf{W} = \{W_t : t \in [0, \infty)\}$ can be obtained from the graph of \mathbf{X} by reflecting in the line $x = X_\tau$ after time τ . In particular, if the stopping time τ is τ_a , the first time that the process hits a specified state $a > 0$, then the graph of \mathbf{W} is obtained from the graph of \mathbf{X} by reflecting in the line $x = a$ after time τ_a .

Open the simulation of reflecting Brownian motion. This app shows the process \mathbf{W} corresponding to the stopping time τ_a , the time of first visit to a positive state a . Run the simulation in single step mode until you see the reflected process several times. Make sure that you understand how the process \mathbf{W} works.

The reflected process $\mathbf{W} = \{W_t : t \in [0, \infty)\}$ is also a standard Brownian motion.

Run the simulation of the reflected Brownian motion process 1000 times. Compare the empirical density function and moments of W_t to the true probability density function and moments.

Martingales

As usual, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be a standard Brownian motion, and let $\mathcal{F}_t = \sigma\{X_s : 0 \leq s \leq t\}$ for $t \in [0, \infty)$, so that $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is the natural filtration for \mathbf{X} . There are several important martingales associated with \mathbf{X} . We will study a couple of them in this section, and others in subsequent sections. Our first result is that \mathbf{X} itself is a martingale, simply by virtue of having stationary, independent increments and 0 mean.

\mathbf{X} is a martingale with respect to \mathfrak{F} .

Proof

Again, this is true of any process with stationary, independent increments and 0 mean, but we give the proof anyway, for completeness. Let $s, t \in [0, \infty)$ with $s < t$. Since X_s is measurable with respect to \mathcal{F}_s and $X_t - X_s$ is independent of \mathcal{F}_s we have

$$\mathbb{E}(X_t | \mathcal{F}_s) = \mathbb{E}[X_s + (X_t - X_s) | \mathcal{F}_s] = X_s + \mathbb{E}(X_t - X_s) = X_s \quad (18.1.17)$$

The next martingale is a little more interesting.

Let $Y_t = X_t^2 - t$ for $t \in [0, \infty)$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is a martingale with respect to \mathfrak{F} .

Proof

Let $s, t \in [0, \infty)$ with $s < t$. Then

$$Y_t = X_t^2 - t = [X_s + (X_t - X_s)]^2 - t = X_s^2 + 2X_s(X_t - X_s) + (X_t - X_s)^2 - t \quad (18.1.18)$$

Since X_s is measurable with respect to \mathcal{F}_s and $X_t - X_s$ is independent of \mathcal{F}_s we have

$$\mathbb{E}(Y_t | \mathcal{F}_s) = X_s^2 + 2X_s \mathbb{E}(X_t - X_s) + \mathbb{E}[(X_t - X_s)^2] - t \quad (18.1.19)$$

But $\mathbb{E}(X_t - X_s) = 0$ and $\mathbb{E}[(X_t - X_s)^2] = \text{var}(X_t - X_s) = t - s$ so $\mathbb{E}(Y_t | \mathcal{F}_s) = X_s^2 - s = Y_s$.

Maximums and Hitting Times

As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. For $y \in [0, \infty)$ recall that $\tau_y = \min\{t \geq 0 : X_t = y\}$ is the first time that the process hits state y . Of course, $\tau_0 = 0$. For $t \in [0, \infty)$, let $Y_t = \max\{X_s : 0 \leq s \leq t\}$, the maximum value of \mathbf{X} on the interval $[0, t]$. Note that Y_t is well defined by the continuity of \mathbf{X} , and of course $Y_0 = 0$. Thus we have two new stochastic processes: $\{\tau_y : y \in [0, \infty)\}$ and $\{Y_t : t \in [0, \infty)\}$. Both have index set $[0, \infty)$ and (as we will see) state space $[0, \infty)$. Moreover, the processes are inverses of each other in a sense:

For $t, y \in (0, \infty)$, $\tau_y \leq t$ if and only if $Y_t \geq y$.

Proof

Since standard Brownian motion starts at 0 and is continuous, both events mean that the process hits state y in the interval $[0, t]$.

Thus, if we can compute the distribution of Y_t for each $t \in (0, \infty)$ then we can compute the distribution of τ_y for each $y \in (0, \infty)$, and conversely.

For $y > 0$, τ_y has the same distribution as y^2/Z^2 , where Z is a standard normal variable. The probability density function g_y is given by

$$g_y(t) = \frac{y}{\sqrt{2\pi t^3}} \exp\left(-\frac{y^2}{2t}\right), \quad t \in (0, \infty) \quad (18.1.20)$$

Proof

Let $t > 0$. From the [previous result](#), note that $X_t \geq y \implies Y_t \geq y \implies \tau_y \leq t$. Hence

$$\mathbb{P}(X_t \geq y) = \mathbb{P}(X_t \geq y, \tau_y \leq t) = \mathbb{P}(X_t \geq y | \tau_y \leq t) \mathbb{P}(\tau_y \leq t) \quad (18.1.21)$$

But from the [strong Markov property](#) above, $s \mapsto X(\tau_y + s) - y$ is another standard Brownian motion. Hence $\mathbb{P}(X_t \geq y | \tau_y \leq t) = \frac{1}{2}$. Therefore

$$\mathbb{P}(\tau_y \leq t) = 2\mathbb{P}(X_t \geq y) = \frac{2}{\sqrt{2\pi t}} \int_y^\infty e^{-x^2/2t} dx = \frac{2}{\sqrt{2\pi}} \int_{y/\sqrt{t}}^\infty e^{-z^2/2} dz \quad (18.1.22)$$

The second integral follows from the first by the change of variables $z = x/\sqrt{t}$. We can recognize this integral as $\mathbb{P}(y^2/Z^2 \leq t)$ where Z has a standard normal distribution. Taking the derivative of the integral with respect to t gives the PDF.

The distribution of τ_y is the *Lévy distribution* with scale parameter y^2 , and is named for the French mathematician Paul Lévy. The Lévy distribution is studied in more detail in the chapter on special distributions.

Open the hitting time experiment. Vary y and note the shape and location of the probability density function of τ_y . For selected values of the parameter, run the simulation in single step mode a few times. Then run the experiment 1000 times and compare the empirical density function to the probability density function.

Open the special distribution simulator and select the Lévy distribution. Vary the parameters and note the shape and location of the probability density function. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function to the probability density function.

Standard Brownian motion is *recurrent*. That is, $\mathbb{P}(\tau_y < \infty) = 1$ for every $y \in \mathbb{R}$.

Proof

Suppose first that $y > 0$. From the proof of the [last theorem](#),

$$\mathbb{P}(\tau_y < \infty) = \lim_{t \rightarrow \infty} \mathbb{P}(\tau_y \leq t) = \frac{2}{\sqrt{2\pi}} \int_0^\infty e^{-z^2/2} dz = 1 \quad (18.1.23)$$

Note that the integral above is equivalent to the integral of the standard normal PDF over \mathbb{R} . In particular, the function g_y given above really is a valid PDF. If $y < 0$ then by symmetry, τ_y has the same distribution as τ_{-y} , so $\mathbb{P}(\tau_y < \infty) = 1$. Trivially, $\tau_0 = 0$.

Thus, for each $y \in \mathbb{R}$, \mathbf{X} eventually hits y with probability 1. Actually we can say more:

With probability 1, \mathbf{X} visits every point in \mathbb{R} .

Proof

By continuity, if \mathbf{X} reaches $y > 0$ then \mathbf{X} visits every point in $[0, y]$. By symmetry, a similar statement holds for $y < 0$. Thus the event that \mathbf{X} visits every point in \mathbb{R} is $\bigcap_{n=1}^\infty (\{\tau_n < \infty\} \cap \{\tau_{-n} < \infty\})$. The probability of a countable intersection of events with probability 1 still has probability 1.

On the other hand,

Standard Brownian motion is *null recurrent*. That is, $\mathbb{E}(\tau_y) = \infty$ for every $y \in \mathbb{R} \setminus \{0\}$.

Proof

By symmetry, it suffices to consider $y > 0$. From the result above on the [distribution of \$\tau_y\$](#) ,

$$\mathbb{E}(\tau_y) = \int_0^\infty \mathbb{P}(\tau_y > t) dt = \frac{2}{\sqrt{2\pi}} \int_0^\infty \int_0^{y/\sqrt{t}} e^{-z^2/2} dz dt \quad (18.1.24)$$

Changing the order of integration gives

$$\mathbb{E}(\tau_y) = \frac{2}{\sqrt{2\pi}} \int_0^\infty \int_0^{y^2/z^2} e^{-z^2/2} dt dz = \frac{2y^2}{\sqrt{2\pi}} \int_0^\infty \frac{1}{z^2} e^{-z^2/2} dz \quad (18.1.25)$$

Next we get a lower bound on the last integral by integrating over the interval $[0, 1]$ and noting that $e^{-z^2/2} \geq e^{-1/2}$ on this interval. Thus,

$$\mathbb{E}(\tau_y) \geq \frac{2y^2 e^{-1/2}}{\sqrt{2\pi}} \int_0^1 \frac{1}{z^2} dz = \infty \quad (18.1.26)$$

The process $\{\tau_x : x \in [0, \infty)\}$ has stationary, independent increments.

Proof

The proof relies on the temporal and spatial homogeneity of Brownian motion and the strong Markov property. Suppose that $x, y \in [0, \infty)$ with $x < y$. By continuity, \mathbf{X} must reach x before reaching y . Thus, $\tau_y = \tau_x + (\tau_y - \tau_x)$. But $\tau_y - \tau_x$ is the hitting time to $y - x$ for the process $t \mapsto X(\tau_x + t) - x$, and as [shown above](#), this process is also a standard Brownian motion, independent of $\mathcal{F}(\tau_x)$. Hence $\tau_y - \tau_x$ is independent of $\mathcal{F}(\tau_x)$ and has the same distribution as τ_{y-x} .

The family of probability density functions $\{g_x : x \in (0, \infty)\}$ is closed under convolution. That is, $g_x * g_y = g_{x+y}$ for $x, y \in (0, \infty)$.

Proof

This follows immediately from the previous theorem. A direct proof is an interesting exercise.

Now we turn our attention to the maximum process $\{Y_t : t \in [0, \infty)\}$, the “inverse” of the hitting process $\{\tau_y : y \in [0, \infty)\}$.

For $t > 0$, Y_t has the same distribution as $|X_t|$, known as the *half-normal distribution* with scale parameter t . The probability density function is

$$h_t(y) = \sqrt{\frac{2}{\pi t}} \exp\left(-\frac{y^2}{2t}\right), \quad y \in [0, \infty) \quad (18.1.27)$$

Proof

From the [inverse relation](#) and the [distribution](#) of τ_y , $\mathbb{P}(Y_t \geq y) = \mathbb{P}(\tau_y \leq t) = 2\mathbb{P}(X_t \geq y) = \mathbb{P}(|X_t| \geq y)$ for $y \geq 0$. By definition, $|X_t|$ has the half-normal distribution with parameter t . In particular,

$$\mathbb{P}(Y_t \geq y) = \frac{2}{\sqrt{2\pi t}} \int_y^\infty e^{-x^2/2t} dx \quad (18.1.28)$$

Taking the negative derivative of the integral above, with respect to y , gives the PDF.

The half-normal distribution is a special case of the folded normal distribution, which is studied in more detail in the chapter on special distributions.

For $t \geq 0$, the mean and variance of Y_t are

1. $\mathbb{E}(Y_t) = \sqrt{\frac{2t}{\pi}}$
2. $\text{var}(Y_t) = t \left(1 - \frac{2}{\pi}\right)$

Proof

These follow from standard results for the half-normal distribution.

In the standard Brownian motion simulation, select the maximum value. Vary the parameter t and note the shape of the probability density function and the location and size of the mean-standard deviation bar. Run the simulation 1000 times and compare the empirical density and moments to the true probability density function and moments.

Open the special distribution simulator and select the folded-normal distribution. Vary the parameters and note the shape and location of the probability density function and the size and location of the mean-standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function and moments to the true density function and moments.

Zeros and Arcsine Laws

As usual, we start with a standard Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$. Study of the zeros of \mathbf{X} lead to a number of probability laws referred to as *arcsine laws*, because as we might guess, the probabilities and distributions involve the arcsine function.

For $s, t \in [0, \infty)$ with $s < t$, let $E(s, t)$ be the event that \mathbf{X} has a zero in the time interval (s, t) . That is, $E(s, t) = \{X_u = 0 \text{ for some } u \in (s, t)\}$. Then

$$\mathbb{P}[E(s, t)] = 1 - \frac{2}{\pi} \arcsin\left(\sqrt{\frac{s}{t}}\right) \quad (18.1.29)$$

Proof

Conditioning on X_s and using symmetry gives

$$\mathbb{P}[E(s, t)] = \int_{-\infty}^{\infty} \mathbb{P}[E(s, t) \mid X_s = x] f_s(x) dx = 2 \int_{-\infty}^0 \mathbb{P}[E(s, t) \mid X_s = x] f_s(x) dx \quad (18.1.30)$$

But by the homogeneity of \mathbf{X} in time and space, note that for $x > 0$, $\mathbb{P}[E(s, t) \mid X_s = -x] = \mathbb{P}(\tau_x < t - s)$. That is, a process in state $-x$ at time s that hits 0 before time t is the same as a process in state 0 at time 0 reaching state x before time $t - s$. Hence

$$\mathbb{P}[E(s, t)] = \int_0^\infty \int_0^{t-s} g_x(u) f_s(-x) du dx \quad (18.1.31)$$

where g_x is the PDF of τ_x given [above](#). Substituting gives

$$\mathbb{P}[E(s, t)] = \frac{1}{\pi\sqrt{s}} \int_0^{t-s} u^{-3/2} \int_0^\infty x \exp\left[-\frac{1}{2}x^2 \left(\frac{u+s}{us}\right)\right] dx du = \frac{\sqrt{s}}{\pi} \int_0^{t-s} \frac{1}{(u+s)\sqrt{u}} du \quad (18.1.32)$$

Finally substituting $v = \sqrt{u/s}$ in the last integral give

$$\mathbb{P}[E(s, t)] = \frac{2}{\pi} \int_0^{\sqrt{t/s-1}} \frac{1}{v^2+1} dv = \frac{2}{\pi} \arctan\left(\sqrt{\frac{t}{s}-1}\right) = 1 - \frac{2}{\pi} \arcsin\left(\sqrt{\frac{s}{t}}\right) \quad (18.1.33)$$

In particular, $\mathbb{P}[E(0, t)] = 1$ for every $t > 0$, so with probability 1, \mathbf{X} has a zero in $(0, t)$. Actually, we can say a bit more:

For $t > 0$, \mathbf{X} has infinitely many zeros in $(0, t)$ with probability 1.

Proof

The event that \mathbf{X} has infinitely many zeros in $(0, t)$ is $\bigcap_{n=1}^{\infty} E(0, t/n)$. The intersection of a countable collection of events with probability 1 still has probability 1.

The last result is further evidence of the very strange and irregular behavior of Brownian motion. Note also that $\mathbb{P}[E(s, t)]$ depends only on the ratio s/t . Thus, $\mathbb{P}[E(s, t)] = \mathbb{P}[E(1/t, 1/s)]$ and $\mathbb{P}[E(s, t)] = \mathbb{P}[E(cs, ct)]$ for every $c > 0$. So, for example the probability of at least one zero in the interval $(2, 5)$ is the same as the probability of at least one zero in $(1/5, 1/2)$ the same as the probability of at least one zero in $(6, 15)$, and the same as the probability of at least one zero in $(200, 500)$.

For $t > 0$, let Z_t denote the time of the last zero of \mathbf{X} before time t . That is, $Z_t = \max\{s \in [0, t] : X_s = 0\}$. Then Z_t has the *arcsine distribution* with parameter t . The distribution function H_t and the probability density function h_t are given by

$$H_t(s) = \frac{2}{\pi} \arcsin\left(\sqrt{\frac{s}{t}}\right), \quad 0 \leq s \leq t \quad (18.1.34)$$

$$h_t(s) = \frac{1}{\pi \sqrt{s(t-s)}}, \quad 0 < s < t \quad (18.1.35)$$

Proof

For $0 \leq s < t$, the event $Z_t \leq s$ is the same as $\neg [E(s, t)]^c$, that there are no zeros in the interval (s, t) . Hence the formula for H_t follows from the [result above](#). Taking the derivative of H_t and simplifying gives the formula for h_t .

The density function of Z_t is *u-shaped* and symmetric about the midpoint $t/2$, so the points with the largest density are those near the endpoints 0 and t , a surprising result at first. The arcsine distribution is studied in more detail in the chapter on special distributions.

The mean and variance of Z_t are

1. $\mathbb{E}(Z_t) = t/2$
2. $\mathbb{E}(Z_t) = t^2/8$

Proof

These are standard results for the arcsine distribution. That the mean is the midpoint $t/2$ also follows from symmetry, of course.

In the simulation of standard Brownian motion, select the last zero variable. Vary the parameter t and note the shape of the probability density function and the size and location of the mean-standard deviation bar. For selected values of t run the simulation in single step mode a few times and note the position of the last zero. Finally, run the simulation 1000 times and compare the empirical density function and moments to the true probability density function and moments.

Open the special distribution simulator and select the arcsine distribution. Vary the parameters and note the shape and location of the probability density function and the size and location of the mean-standard deviation bar. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function and moments to the true density function and moments.

Now let $Z = \{t \in [0, \infty) : X_t = 0\}$ denote the set of zeros of \mathbf{X} , so that Z is a random subset of $[0, \infty)$. The theorem below gives some of the strange properties of the random set Z , but to understand these, we need to review some definitions. A *nowhere dense* set is a set whose closure has empty interior. A *perfect set* is a set with no isolated points. As usual, we let λ denote Lebesgue measure on \mathbb{R} .

With probability 1,

1. Z is closed.
2. $\lambda(Z) = 0$
3. Z is nowhere dense.
4. Z is perfect.

Proof

1. Note that Z is the inverse image of the closed set $\{0\}$ under the function $t \mapsto X_t$. Since this function is continuous with probability 1, Z is closed with probability 1.

2. For each $t \in (0, \infty)$ note that $\mathbb{P}(t \in Z) = \mathbb{P}(X_t = 0) = 0$ since X_t has a continuous distribution. Using Fubini's theorem

$$\mathbb{E}[\lambda(Z)] = \mathbb{E}\left[\int_0^\infty \mathbf{1}_Z(t) d\lambda(t)\right] = \int_0^\infty \mathbb{E}[\mathbf{1}_Z(t)] d\lambda(t) = 0 \quad (18.1.36)$$

and hence $\mathbb{P}[\lambda(Z) = 0] = 1$,

3. Since Z is closed and has Lebesgue measure 0, its interior is empty (all of these statements with probability 1).

4. Suppose that $s \in Z$. Then by the [temporal and spatial homogeneity](#) properties, $t \mapsto X_{s+t}$ is also a standard Brownian motion. But then by the result above on [zeros](#), with probability 1, \mathbf{X} has a zero in the interval $(s, s + 1/n)$ for every $n \in \mathbb{N}_+$. Hence s is not an isolated point of Z .

The following theorem gives a deeper property of Z . The Hausdorff dimension of Z is midway between that of a point (dimension 0) and a line (dimension 1).

Z has Hausdorff dimension $\frac{1}{2}$.

The Law of the Iterated Logarithm

As usual, let $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ be standard Brownian motion. By definition, we know that X_t has the normal distribution with mean 0 and standard deviation \sqrt{t} , so the function $x = \sqrt{t}$ gives some idea of how the process grows in time. The precise growth rate is given by the famous *law of the iterated logarithm*

With probability 1,

$$\limsup_{t \rightarrow \infty} \frac{X_t}{\sqrt{2t \ln \ln t}} = 1 \quad (18.1.37)$$

Computational Exercises

In the following exercises, $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a standard Brownian motion process.

Explicitly find the probability density function, covariance matrix, and correlation matrix of $(X_{0.5}, X_1, X_{2.3})$.

This page titled [18.1: Standard Brownian Motion](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist](#) ([Random Services](#)) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

18.2: Brownian Motion with Drift and Scaling

Basic Theory

Definition

We start with the assumptions that govern standard Brownian motion, except that we relax the restrictions on the parameters of the normal distribution.

Suppose that $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$. *Brownian motion* with drift parameter μ and scale parameter σ is a random process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with state space \mathbb{R} that satisfies the following properties:

1. $X_0 = 0$ (with probability 1).
2. \mathbf{X} has stationary increments. That is, for $s, t \in [0, \infty)$ with $s < t$, the distribution of $X_t - X_s$ is the same as the distribution of X_{t-s} .
3. \mathbf{X} has independent increments. That is, for $t_1, t_2, \dots, t_n \in [0, \infty)$ with $t_1 < t_2 < \dots < t_n$, the random variables $X_{t_1}, X_{t_2} - X_{t_1}, \dots, X_{t_n} - X_{t_{n-1}}$ are independent.
4. X_t has the normal distribution with mean μt and variance $\sigma^2 t$ for $t \in [0, \infty)$.
5. With probability 1, $t \mapsto X_t$ is continuous on $[0, \infty)$.

Note that we cannot assign the parameters of the normal distribution of X_t arbitrarily. We know that since \mathbf{X} has stationary, independent increments, $\mathbb{E}(X_t)$ and $\text{var}(X_t)$ must be linear functions of $t \in [0, \infty)$.

Open the simulation of Brownian motion with drift and scaling. Run the simulation in single step mode several times for various values of the parameters. Note the behavior of the sample paths. For selected values of the parameters, run the simulation 1000 times and compare the empirical density function and moments to the true density function and moments.

It's easy to construct Brownian motion with drift and scaling from a standard Brownian motion, so we don't have to worry about the existence question.

Relation to standard Brownian motion.

1. Suppose that $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is a standard Brownian motion, and that $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$. Let $X_t = \mu t + \sigma Z_t$ for $t \in [0, \infty)$. Then $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Brownian motion with drift parameter μ and scale parameter σ .
2. Conversely, suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Brownian motion with drift parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$. Let $Z_t = (X_t - \mu t)/\sigma$ for $t \in [0, \infty)$. Then $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is a standard Brownian motion.

Proof

It's straightforward to show that the processes \mathbf{X} and \mathbf{Z} satisfy the appropriate set of assumptions.

In differential form, part (a) can be written as

$$dX_t = \mu dt + \sigma dZ_t, \quad X_0 = 0 \quad (18.2.1)$$

Finite Dimensional Distributions

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is Brownian motion with drift parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$. It follows from part (d) of the [definition](#) that X_t has probability density function f_t given by

$$f_t(x) = \frac{1}{\sigma\sqrt{2\pi t}} \exp\left[-\frac{1}{2\sigma^2 t}(x - \mu t)^2\right], \quad x \in \mathbb{R} \quad (18.2.2)$$

This family of density functions determines the finite dimensional distributions of \mathbf{X} .

If $t_1, t_2, \dots, t_n \in (0, \infty)$ with $0 < t_1 < t_2 < \dots < t_n$ then $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ has probability density function f_{t_1, t_2, \dots, t_n} given by

$$f_{t_1, t_2, \dots, t_n}(x_1, x_2, \dots, x_n) = f_{t_1}(x_1) f_{t_2 - t_1}(x_2 - x_1) \cdots f_{t_n - t_{n-1}}(x_n - x_{n-1}), \quad (x_1, x_2, \dots, x_n) \in \mathbb{R}^n \quad (18.2.3)$$

Proof

This follows because \mathbf{X} has stationary, independent increments.

\mathbf{X} is a Gaussian process with mean function mean function m and covariance function c given by

1. $m(t) = \mu t$ for $t \in [0, \infty)$
2. $c(s, t) = \sigma^2 \min\{s, t\}$ for $s, t \in [0, \infty)$.

Proof

The fact that \mathbf{X} is a Gaussian process follows from the construction $X_t = \mu t + \sigma Z_t$ for $t \in [0, \infty)$, where \mathbf{Z} is a standard Brownian motion. We know that \mathbf{Z} is a Gaussian process. The form of the mean and covariance functions follow because \mathbf{X} has stationary, independent increments. Note that μ and σ^2 are the mean and variance of X_1 .

The correlation function is independent of the parameters, and thus is the same as for standard Brownian motion. This is hardly surprising since correlation is a standardized measure of association.

$$\text{cor}(X_s, X_t) \frac{\sigma^2 \min\{s, t\}}{\sigma s \sigma t} = \frac{\min\{s, t\}}{st} = \sqrt{\frac{\min\{s, t\}}{\max\{s, t\}}}, \quad (s, t) \in [0, \infty)^2 \quad (18.2.4)$$

Transformations

There are a couple simple transformations that preserve Brownian motion, but perhaps change the drift and scale parameters. Our starting place is a Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with drift parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$. Our first result involves scaling \mathbf{X} in time and space (and possibly reflecting in the spatial origin).

Let $a \in \mathbb{R} \setminus \{0\}$ and $b \in (0, \infty)$. Define $Y_t = aX_{bt}$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \geq 0\}$ is also a Brownian motion with drift parameter $ab\mu$ and scale parameter $|a|\sqrt{b}\sigma$.

Proof

Clearly the new process is still a Gaussian process. The mean function is $\mathbb{E}(Y_t) = a\mathbb{E}(X_{bt}) = ab\mu t$ for $t \in [0, \infty)$. The covariance function is $\text{cov}(Y_s, Y_t) = a^2 \text{cov}(X_{bs}, X_{bt}) = a^2 \sigma^2 \min\{bs, bt\} = a^2 b \sigma^2 \min\{s, t\}$ for $(s, t) \in [0, \infty)^2$. Finally, since \mathbf{X} is continuous, so is \mathbf{Y} .

Suppose that $a > 0$ in the previous theorem, so that we are scaling temporally and spatially. In order to preserve the original drift parameter μ we must have $ab = 1$ (if $\mu \neq 0$). In order to preserve the original scale parameter σ , we must have $a\sqrt{b} = 1$. We can't have both unless $\mu = 0$, which leads to a slight generalization of one of our results for standard Brownian motion:

Suppose that \mathbf{X} is a Brownian motion with drift parameter $\mu = 0$ and scale parameter $\sigma > 0$. Suppose also that $c > 0$ and let $Y_t = \frac{1}{c} X_{c^2 t}$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a Brownian motion with drift parameter 0 and scale parameter σ .

Our next result is related to the Markov property, which we explore in more detail below. We return to the general case where $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a Brownian motion with drift parameter $\mu \in \mathbb{R}$ and scale parameter $\sigma \in (0, \infty)$. If we “restart” Brownian motion at a fixed time s , and shift the origin to X_s , then we have another Brownian motion with the same parameters.

Fix $s \in [0, \infty)$ and define $Y_t = X_{s+t} - X_s$ for $t \geq 0$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is also a Brownian motion with the same drift and scale parameters.

Proof

Clearly \mathbf{Y} is also a Gaussian process. Moreover, $\mathbb{E}(Y_t) = \mathbb{E}(X_{s+t}) - \mathbb{E}(X_s) = \mu(s+t) - \mu s = \mu t$ for $t \in [0, \infty)$. Also, if $r, t \in [0, \infty)$ with $r \leq t$ then

$$\text{cov}(Y_r, Y_t) = \text{cov}(X_{s+r} - X_s, X_{s+t} - X_s) \quad (18.2.5)$$

$$= \text{cov}(X_{s+r}, X_{s+t}) - \text{cov}(X_{s+r}, X_s) - \text{cov}(X_s, X_{s+t}) + \text{cov}(X_s, X_s) \quad (18.2.6)$$

$$= \sigma^2(s+r) - \sigma^2 s - \sigma^2 s + \sigma^2 s = \sigma^2 r \quad (18.2.7)$$

Finally, \mathbf{Y} is continuous by the continuity of \mathbf{X} .

The Markov Property and Stopping Times

As usual, we start with a Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ with drift parameter μ and scale parameter σ . Recall again that a Markov process has the property that the future is independent of the past, given the present state. Because of the stationary, independent increments property, Brownian motion has the property. As a minor note, to view \mathbf{X} as a Markov process, we sometimes need to relax Assumption 1 and let X_0 have an arbitrary value in \mathbb{R} . Let $\mathcal{F}_t = \sigma\{X_s : 0 \leq s \leq t\}$, the sigma-algebra generated by the process up to time $t \in [0, \infty)$. The family of σ -algebras $\mathcal{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is known as a filtration.

Brownian motion is a time-homogeneous Markov process with transition probability density p given by

$$p_t(x, y) = f_t(y - x) = \frac{1}{\sigma\sqrt{2\pi t}} \exp\left[-\frac{1}{2\sigma^2 t}(y - x - \mu t)^2\right], \quad t \in (0, \infty); \quad x, y \in \mathbb{R} \quad (18.2.8)$$

Proof

Fix $s \in [0, \infty)$. The theorem follows from the fact that the process $\{X_{s+t} - X_s : t \in [0, \infty)\}$ is another standard Brownian motion, as noted [above](#), and is independent of \mathcal{F}_s .

The transition density p satisfies the following *diffusion equations*. The first is known as the *forward equation* and the second as the *backward equation*.

$$\frac{\partial}{\partial t} p_t(x, y) = -\mu \frac{\partial}{\partial y} p_t(x, y) + \frac{1}{2} \sigma^2 \frac{\partial^2}{\partial y^2} p_t(x, y) \quad (18.2.9)$$

$$\frac{\partial}{\partial t} p_t(x, y) = \mu \frac{\partial}{\partial x} p_t(x, y) + \frac{1}{2} \sigma^2 \frac{\partial^2}{\partial x^2} p_t(x, y) \quad (18.2.10)$$

Proof

These results follow from standard calculus.

The diffusion equations are so named, because the spatial derivative in the first equation is with respect to y , the state *forward* at time t , while the spatial derivative in the second equation is with respect to x , the state *backward* at time 0.

Recall again that a random time τ taking values in $[0, \infty]$ is a *stopping time* with respect to the process \mathbf{X} if $\{\tau \leq t\} \in \mathcal{F}_t$ for every $t \in [0, \infty)$. The σ -algebra associated with τ is

$$\mathcal{F}_\tau = \{B \in \mathcal{F} : B \cap \{\tau \leq t\} \in \mathcal{F}_t \text{ for all } t \geq 0\} \quad (18.2.11)$$

See the section on Filtrations and Stopping Times for more information on filtrations, stopping times, and the σ -algebra associated with a stopping time. Brownian motion \mathbf{X} is also a strong Markov process.

Suppose that τ is a stopping time and define $Y_t = X_{\tau+t} - X_\tau$ for $t \in [0, \infty)$. Then $\mathbf{Y} = \{Y_t : t \in [0, \infty)\}$ is a Brownian motion with the same drift and scale parameters, and is independent of \mathcal{F}_τ .

This page titled [18.2: Brownian Motion with Drift and Scaling](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

18.3: The Brownian Bridge

Basic Theory

Definition and Constructions

In the most common formulation, the *Brownian bridge process* is obtained by taking a standard Brownian motion process \mathbf{X} , restricted to the interval $[0, 1]$, and conditioning on the event that $X_1 = 0$. Since $X_0 = 0$ also, the process is “tied down” at both ends, and so the process in between forms a “bridge” (albeit a very jagged one). The Brownian bridge turns out to be an interesting stochastic process with surprising applications, including a very important application to statistics. In terms of a definition, however, we will give a list of characterizing properties as we did for standard Brownian motion and for Brownian motion with drift and scaling.

A *Brownian bridge* is a stochastic process $\mathbf{X} = \{X_t : t \in [0, 1]\}$ with state space \mathbb{R} that satisfies the following properties:

1. $X_0 = 0$ and $X_1 = 0$ (each with probability 1).
2. \mathbf{X} is a Gaussian process.
3. $\mathbb{E}(X_t) = 0$ for $t \in [0, 1]$.
4. $\text{cov}(X_s, X_t) = \min\{s, t\} - st$ for $s, t \in [0, 1]$.
5. With probability 1, $t \mapsto X_t$ is continuous on $[0, 1]$.

So, in short, a Brownian bridge \mathbf{X} is a continuous Gaussian process with $X_0 = X_1 = 0$, and with mean and covariance functions given in (c) and (d), respectively. Naturally, the first question is whether there exists such a process. The answer is yes, of course, otherwise why would we be here? But in fact, we will see several ways of constructing a Brownian bridge from a standard Brownian motion. To help with the proofs, recall that a standard Brownian motion process $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is a continuous Gaussian process with $Z_0 = 0$, $\mathbb{E}(Z_t) = 0$ for $t \in [0, \infty)$ and $\text{cov}(Z_s, Z_t) = \min\{s, t\}$ for $s, t \in [0, \infty)$. Here is our first construction:

Suppose that $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is a standard Brownian motion, and let $X_t = Z_t - tZ_1$ for $t \in [0, 1]$. Then $\mathbf{X} = \{X_t : t \in [0, 1]\}$ is a Brownian bridge.

Proof

1. Note that $X_0 = Z_0 = 0$ and $X_1 = Z_1 - Z_1 = 0$.
2. Linear combinations of the variables in \mathbf{X} reduce to linear combinations of the variables in \mathbf{Z} and hence have normal distributions. Thus \mathbf{X} is a Gaussian process.
3. $\mathbb{E}(X_t) = \mathbb{E}(Z_t) - t\mathbb{E}(Z_1) = 0$ for $t \in [0, 1]$
4. $\text{cov}(X_s, X_t) = \text{cov}(Z_s - sZ_1, Z_t - tZ_1) = \text{cov}(Z_s, Z_t) - t\text{cov}(Z_s, Z_1) - s\text{cov}(Z_1, Z_t) + st\text{cov}(Z_1, Z_1) = \min\{s, t\} - st - st + st$ for $s, t \in [0, 1]$.
5. $t \mapsto X_t$ is continuous on $[0, 1]$ since $t \mapsto Z_t$ is continuous on $[0, 1]$.

Let's see the Brownian bridge in action.

Run the simulation of the Brownian bridge process in single step mode a few times.

For the Brownian bridge \mathbf{X} , note in particular that X_t is normally distributed with mean 0 and variance $t(1-t)$ for $t \in [0, 1]$. Thus, the variance increases and then decreases on $[0, 1]$ reaching a maximum of $1/4$ at $t = 1/2$. Of course, the variance is 0 at $t = 0$ and $t = 1$, since $X_0 = X_1 = 0$ deterministically.

Open the simulation of the Brownian bridge process. Vary t and note the change in the probability density function and moments. For various values of t , run the simulation 1000 times and compare the empirical density function and moments to the true density function and moments.

Conversely to the [construction above](#), we can build a standard Brownian motion on the time interval $[0, 1]$ from a Brownian bridge.

Suppose that $\mathbf{X} = \{X_t : t \in [0, 1]\}$ is a Brownian bridge, and suppose that Z is a random variable with a standard normal distribution, independent of \mathbf{X} . Let $Z_t = X_t + tZ$ for $t \in [0, 1]$. Then $\mathbf{Z} = \{Z_t : t \in [0, 1]\}$ is a standard Brownian motion on $[0, 1]$.

Proof

1. Note that $Z_0 = X_0 = 0$.
2. Linear combinations of the variables in \mathbf{Z} reduce to linear combinations of the variables in \mathbf{X} and hence have normal distributions. Thus \mathbf{Z} is a Gaussian process.
3. $\mathbb{E}(Z_t) = \mathbb{E}(X_t) + t\mathbb{E}(Z) = 0$ for $t \in [0, 1]$.
4. $\text{cov}(Z_s, Z_t) = \text{cov}(X_s + sZ, X_t + tZ) = \text{cov}(X_s, X_t) + t\text{cov}(X_s, Z) + s\text{cov}(X_t, Z) + st\text{var}(Z) = \min\{s, t\} - st + 0 + 0 + st = \min\{s, t\}$ for $s, t \in [0, 1]$.
5. $t \mapsto Z_t$ is continuous on $[0, 1]$ since $t \mapsto X_t$ is continuous on $[0, 1]$.

Here's another way to construct a Brownian bridge from a standard Brownian motion.

Suppose that $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is a standard Brownian motion. Define $X_1 = 0$ and

$$X_t = (1-t)Z\left(\frac{t}{1-t}\right), \quad t \in [0, 1) \quad (18.3.1)$$

Then $\mathbf{X} = \{X_t : t \in [0, 1]\}$ is a Brownian bridge.

Proof

1. Note that $X_0 = Z_0 = 0$ and by definition, $X_1 = 0$.
2. Linear combinations of variables in \mathbf{X} reduce to linear combinations of variables in \mathbf{Z} and hence have normal distributions. Thus \mathbf{X} is a Gaussian process.
3. For $t \in [0, 1]$,

$$\mathbb{E}(X_t) = (1-t)\mathbb{E}\left[Z\left(\frac{t}{1-t}\right)\right] = 0 \quad (18.3.2)$$

4. If $s, t \in [0, 1]$ with $s < t$ then $s/(1-s) < t/(1-t)$ so

$$\text{cov}(X_s, X_t) = \text{cov}\left[(1-s)Z\left(\frac{s}{1-s}\right), (1-t)Z\left(\frac{t}{1-t}\right)\right] = (1-s)(1-t)\frac{s}{1-s} = s(1-t) \quad (18.3.3)$$

5. Finally, $t \mapsto X_t$ is continuous with probability 1 on $[0, 1)$, and with probability 1, $X_t = (1-t)Z[t/(1-t)] \rightarrow 0$ as $t \uparrow 1$.

Conversely, we can construct a standard Brownian motion from a Brownian bridge.

Suppose that $\mathbf{X} = \{X_t : t \in [0, 1]\}$ is a Brownian bridge. Define

$$Z_t = (1+t)X\left(\frac{t}{1+t}\right), \quad t \in [0, \infty) \quad (18.3.4)$$

Then $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is a standard Brownian motion process.

Proof

1. Note that $Z_0 = X_0 = 0$
2. Linear combinations of the variables in \mathbf{Z} reduce to linear combinations of the variables in \mathbf{X} , and hence have normal distributions. Thus \mathbf{Z} is a Gaussian process.
3. For $t \in [0, \infty)$,

$$\mathbb{E}(Z_t) = (1+t)\mathbb{E}\left[X\left(\frac{t}{1+t}\right)\right] = 0 \quad (18.3.5)$$

4. If $s, t \in [0, 1]$ with $s < t$ Then $s/(1+s) < t/(1+t)$ so

$$\text{cov}(Z_s, Z_t) = \text{cov}\left[(1+s)X\left(\frac{s}{1+s}\right), (1+t)X\left(\frac{t}{1+t}\right)\right] = (1+s)(1+t)\left[\frac{s}{1+s} - \frac{s}{1+s}\frac{t}{1+t}\right] = s \quad (18.3.6)$$

5. Since $t \mapsto X_t$ is continuous, $t \mapsto Z_t$ is continuous

We return to the comments at the beginning of this section, on conditioning a standard Brownian motion to be 0 at time 1. Unlike the previous two constructions, note that we are not transforming the random variables, rather we are changing the underlying *probability measure*.

Suppose that $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is a standard Brownian motion. Then conditioned on $X_1 = 0$, the process $\{X_t : t \in [0, 1]\}$ is a Brownian bridge process.

Proof

Part of the argument is based on properties of the multivariate normal distribution. The conditioned process is still continuous and is still a Gaussian process. In particular, suppose that $s, t \in [0, 1]$ with $s < t$. Then (X_t, X_1) has a joint normal distribution with parameters specified by the mean and covariance functions of \mathbf{X} . By standard computations, the conditional distribution of X_t given $X_1 = 0$ is normal with mean 0 and variance $t(1-t)$. Similarly, the joint distribution of (X_s, X_t, X_1) is normal with parameters specified by the mean and covariance functions of \mathbf{X} . Again, by standard computations, the conditional distribution of (X_s, X_t) given $X_1 = 0$ is bivariate normal with 0 means and with $\text{cov}(X_s, X_t | X_1 = 0) = s(1-t)$.

Finally, the Brownian bridge can be defined in terms a stochastic integral

Suppose that $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is standard Brownian motions. Define $X_1 = 1$ and

$$X_t = (1-t) \int_0^t \frac{1}{1-s} dZ_s, \quad t \in [0, 1) \quad (18.3.7)$$

Then $\mathbf{X} = \{X_t : t \in [0, 1]\}$ is a Brownian bridge process.

Proof

1. Note that $X_0 = 0$ and by definition, $X_1 = 0$.
2. Since the integrand in the stochastic integral is deterministic, \mathbf{X} is a Gaussian process.
3. \mathbf{X} is continuous on $[0, 1]$ with probability 1, as a basic property of stochastic integrals. Moreover, $X_t \rightarrow 0$ as $t \uparrow 1$ as a consequence of the martingale inequality.
4. $\mathbb{E}(X_t) = 0$ since the stochastic integral has mean 0.
5. Suppose that $s, t \in [0, 1]$ with $s \leq t$. Then

$$\text{cov}(X_s, X_t) = \text{cov} \left[(1-s) \int_0^s \frac{1}{1-u} dZ_u, (1-t) \left(\int_0^s \frac{1}{1-u} dZ_u + \int_s^t \frac{1}{1-u} dZ_u \right) \right] \quad (18.3.8)$$

But $\int_0^s \frac{1}{1-u} dZ_u$ and $\int_s^t \frac{1}{1-u} dZ_u$ are independent,

$$\text{cov}(X_s, X_t) = (1-s)(1-t) \text{var} \left(\int_0^s \frac{1}{1-u} dZ_u \right) \quad (18.3.9)$$

But then by the Ito isometry,

$$\text{cov}(X_s, X_t) = (1-s)(1-t) \int_0^s \frac{1}{(1-u)^2} du = (1-s)(1-t) \left(\frac{1}{1-s} - 1 \right) = (1-t)s \quad (18.3.10)$$

In differential form, the process above can be written as

$$dX_t = \frac{X_t}{1-t} dt + dZ_t, \quad X_0 = 0 \quad (18.3.11)$$

The General Brownian Bridge

The processes constructed above (in several ways!) is the *standard* Brownian bridge. it's a simple matter to generalize the process so that it starts at a and ends at b , for arbitrary $a, b \in \mathbb{R}$.

Suppose that $\mathbf{Z} = \{Z_t : t \in [0, 1]\}$ is a standard Brownian bridge process. Let $a, b \in \mathbb{R}$ and define $X_t = (1-t)a + tb + Z_t$ for $t \in [0, 1]$. Then $\mathbf{X} = \{X_t : t \in [0, 1]\}$ is a *Brownian bridge process from a to b* .

Of course, any of the constructions above for standard Brownian bridge can be modified to produce a general Brownian bridge. Here are the characterizing properties.

The Brownian bridge process $\mathbf{X} = \{X_t : t \in [0, 1]\}$ from a to b is characterized by the following properties:

1. $X_0 = a$ and $X_1 = b$ (each with probability 1).
2. \mathbf{X} is a Gaussian process.
3. $\mathbb{E}(X_t) = (1-t)a + tb$ for $t \in [0, 1]$.
4. $\text{cov}(X_s, X_t) = \min\{s, t\} - st$ for $s, t \in [0, 1]$.
5. With probability 1, $t \mapsto X_t$ is continuous on $[0, 1]$.

Applications

The Empirical Distribution Function

We start with a problem that is one of the most basic in statistics. Suppose that T is a real-valued random variable with an unknown distribution. Let F denote the distribution function of T , so that $F(t) = \mathbb{P}(T \leq t)$ for $t \in \mathbb{R}$. Our goal is to construct an estimator of F , so naturally our first step is to *sample* from the distribution of T . This generates a sequence $\mathbf{T} = (T_1, T_2, \dots)$ of independent variables, each with the distribution of T (and so with distribution function F). Think of \mathbf{T} as a sequence of independent *copies* of T . For $n \in \mathbb{N}_+$ and $t \in \mathbb{R}$, the natural estimator of $F(t)$ based on the first n sample values is

$$F_n(t) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(T_i \leq t) \quad (18.3.12)$$

which is simply the proportion of the first n sample values that fall in the interval $(-\infty, t]$. Appropriately enough, F_n is known as the *empirical distribution function* corresponding to the sample of size n . Note that $(\mathbf{1}(T_1 \leq t), \mathbf{1}(T_2 \leq t), \dots)$ is a sequence of independent, identically distributed indicator variables (and hence is a sequence of Bernoulli trials), and corresponds to sampling from the distribution of $\mathbf{1}(T \leq t)$. The estimator $F_n(t)$ is simply the sample mean of the first n of these variables. The numerator, the number of the original sample variables with values in $(-\infty, t]$, has the binomial distribution with parameters n and $F(t)$. Like all sample means from independent, identically distributed samples, $F_n(t)$ satisfies some basic and important properties. A summary is given below, but to make sense of some of these facts, you need to recall the mean and variance of the indicator variable that we are sampling from: $\mathbb{E}[\mathbf{1}(T \leq t)] = F(t)$, $\text{var}[\mathbf{1}(T \leq t)] = F(t)[1 - F(t)]$

For fixed $t \in \mathbb{R}$,

1. $\mathbb{E}[F_n(t)] = F(t)$ so $F_n(t)$ is an unbiased estimator of $F(t)$
2. $\text{var}[F_n(t)] = F(t)[1 - F(t)]/n$ so $F_n(t)$ is a consistent estimator of $F(t)$

3. $F_n(t) \rightarrow F(t)$ as $n \rightarrow \infty$ with probability 1, the strong law of large numbers.
4. $\sqrt{n} [F_n(t) - F(t)]$ has mean 0 and variance $F(t) [1 - F(t)]$ and converges to the normal distribution with these parameters as $n \rightarrow \infty$, the central limit theorem.

The theorem above gives us a great deal of information about $F_n(t)$ for fixed t , but now we want to let t vary and consider the expression in (d), namely $t \mapsto \sqrt{n} [F_n(t) - F(t)]$, as a random process for each $n \in \mathbb{N}_+$. The key is to consider a very special distribution first.

Suppose that T has the standard uniform distribution, that is, the continuous uniform distribution on the interval $[0, 1]$. In this case the distribution function is simply $F(t) = t$ for $t \in [0, 1]$, so we have the sequence of stochastic processes $\mathbf{X}_n = \{X_n(t) : t \in [0, 1]\}$ for $n \in \mathbb{N}_+$, where

$$X_n(t) = \sqrt{n} [F_n(t) - t] \quad (18.3.13)$$

Of course, the previous results apply, so the process \mathbf{X}_n has mean function 0, variance function $t \mapsto t(1-t)$, and for fixed $t \in [0, 1]$, the distribution $X_n(t)$ converges to the corresponding normal distribution as $n \rightarrow \infty$. Here is the new bit of information, the covariance function of \mathbf{X}_n is the same as that of the Brownian bridge!

$$\text{cov} [X_n(s), X_n(t)] = \min\{s, t\} - st \text{ for } s, t \in [0, 1].$$

Proof

Suppose that $s \leq t$. From basic properties of covariance,

$$\text{cov} [X_n(s), X_n(t)] = n \text{cov} [F_n(s), F_n(t)] = \frac{1}{n} \text{cov} \left(\sum_{i=1}^n \mathbf{1}(T_i \leq s), \sum_{j=1}^n \mathbf{1}(T_j \leq t) \right) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \text{cov} [\mathbf{1}(T_i \leq s), \mathbf{1}(T_j \leq t)] \quad (18.3.14)$$

But if $i \neq j$, the variables $\mathbf{1}(T_i \leq s)$ and $\mathbf{1}(T_j \leq t)$ are independent, and hence have covariance 0. On the other hand,

$$\text{cov} [\mathbf{1}(T_i \leq s), \mathbf{1}(T_i \leq t)] = \mathbb{P}(T_i \leq s, T_i \leq t) - \mathbb{P}(T_i \leq s)\mathbb{P}(T_i \leq t) = \mathbb{P}(T_i \leq s) - \mathbb{P}(T_i \leq s)\mathbb{P}(T_i \leq t) = s - st \quad (18.3.15)$$

hence

$$\text{cov} [X_n(s), X_n(t)] = \frac{1}{n} \sum_{i=1}^n \text{cov} [\mathbf{1}(T_i \leq s), \mathbf{1}(T_i \leq t)] = s - st \quad (18.3.16)$$

This page titled [18.3: The Brownian Bridge](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

18.4: Geometric Brownian Motion

Basic Theory

Geometric Brownian motion, and other stochastic processes constructed from it, are often used to model population growth, financial processes (such as the price of a stock over time), subject to random “noise”.

Definition

Suppose that $\mathbf{Z} = \{Z_t : t \in [0, \infty)\}$ is standard Brownian motion and that $\mu \in \mathbb{R}$ and $\sigma \in (0, \infty)$. Let

$$X_t = \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) t + \sigma Z_t \right], \quad t \in [0, \infty) \quad (18.4.1)$$

The stochastic process $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ is *geometric Brownian motion* with *drift parameter* μ and *volatility parameter* σ .

Note that the stochastic process

$$\left\{ \left(\mu - \frac{\sigma^2}{2} \right) t + \sigma Z_t : t \in [0, \infty) \right\} \quad (18.4.2)$$

is Brownian motion with drift parameter $\mu - \sigma^2/2$ and scale parameter σ , so geometric Brownian motion is simply the exponential of this process. In particular, the process is always positive, one of the reasons that geometric Brownian motion is used to model financial and other processes that cannot be negative. Note also that $X_0 = 1$, so the process starts at 1, but we can easily change this. For $x_0 \in (0, \infty)$, the process $\{x_0 X_t : t \in [0, \infty)\}$ is geometric Brownian motion *starting at* x_0 . You may well wonder about the particular combination of parameters $\mu - \sigma^2/2$ in the definition. The short answer to the question is given in the following theorem:

Geometric Brownian motion $\mathbf{X} = \{X_t : t \in [0, \infty)\}$ satisfies the stochastic differential equation

$$dX_t = \mu X_t dt + \sigma X_t dZ_t \quad (18.4.3)$$

Note that the deterministic part of this equation is the standard differential equation for exponential growth or decay, with rate parameter μ .

Run the simulation of geometric Brownian motion several times in single step mode for various values of the parameters. Note the behavior of the process.

Distributions

For $t \in (0, \infty)$, X_t has the lognormal distribution with parameters $\left(\mu - \frac{\sigma^2}{2} \right) t$ and $\sigma\sqrt{t}$. The probability density function f_t is given by

$$f_t(x) = \frac{1}{\sqrt{2\pi t} \sigma x} \exp \left(- \frac{[\ln(x) - (\mu - \frac{\sigma^2}{2})t]^2}{2\sigma^2 t} \right), \quad x \in (0, \infty) \quad (18.4.4)$$

1. f increases and then decreases with mode at $x = \exp \left[\left(\mu - \frac{3}{2} \sigma^2 \right) t \right]$
2. f is concave upward, then downward, then upward again with inflection points at $x = \exp \left[\left(\mu - \sigma^2 \right) t \pm \frac{1}{2} \sigma \sqrt{\sigma^2 t^2 + 4t} \right]$

Proof

Since the variable $U_t = \left(\mu - \frac{\sigma^2}{2} \right) t + \sigma Z_t$ has the normal distribution with mean $\left(\mu - \frac{\sigma^2}{2} \right) t$ and standard deviation $\sigma\sqrt{t}$, it follows that $X_t = \exp(U_t)$ has the lognormal distribution with these parameters. These result for the PDF then follow directly from the corresponding results for the lognormal PDF.

In particular, geometric Brownian motion is not a Gaussian process.

Open the simulation of geometric Brownian motion. Vary the parameters and note the shape of the probability density function of X_t . For various values of the parameters, run the simulation 1000 times and compare the empirical density function to the true probability density function.

For $t \in (0, \infty)$, the distribution function F_t of X_t is given by

$$F_t(x) = \Phi \left[\frac{\ln(x) - (\mu - \sigma^2/2)t}{\sigma\sqrt{t}} \right], \quad x \in (0, \infty) \quad (18.4.5)$$

where Φ is the standard normal distribution function.

Proof

Again, this follows directly from the CDF of the lognormal distribution.

For $t \in (0, \infty)$, the quantile function F_t^{-1} of X_t is given by

$$F_t^{-1}(p) = \exp[(\mu - \sigma^2/2)t + \sigma\sqrt{t}\Phi^{-1}(p)], \quad p \in (0, 1) \quad (18.4.6)$$

where Φ^{-1} is the standard normal quantile function.

Proof

This follows directly from the lognormal quantile function.

Moments

For $n \in \mathbb{N}$ and $t \in [0, \infty)$,

$$\mathbb{E}(X_t^n) = \exp \left\{ \left[n\mu + \frac{\sigma^2}{2}(n^2 - n) \right] t \right\} \quad (18.4.7)$$

Proof

This follows from the formula for the moments of the lognormal distribution.

In terms of the order of the moment n , the dominant term inside the exponential is $\sigma^2 n^2/2$. If $n > 1 - 2\mu/\sigma^2$ then $n\mu + \frac{\sigma^2}{2}(n^2 - n) > 0$ so $\mathbb{E}(X_t^n) \rightarrow \infty$ as $t \rightarrow \infty$. The mean and variance follow easily from the general moment result.

For $t \in [0, \infty)$,

1. $\mathbb{E}(X_t) = e^{\mu t}$
2. $\text{var}(X_t) = e^{2\mu t} (e^{\sigma^2 t} - 1)$

In particular, note that the mean function $m(t) = \mathbb{E}(X_t) = e^{\mu t}$ for $t \in [0, \infty)$ satisfies the deterministic part of the [stochastic differential equation](#) above. If $\mu > 0$ then $m(t) \rightarrow \infty$ as $t \rightarrow \infty$. If $\mu = 0$ then $m(t) = 1$ for all $t \in [0, \infty)$. If $\mu < 0$ then $m(t) \rightarrow 0$ as $t \rightarrow \infty$.

Open the simulation of geometric Brownian motion. The graph of the mean function m is shown as a blue curve in the main graph box. For various values of the parameters, run the simulation 1000 times and note the behavior of the random process in relation to the mean function.

Open the simulation of geometric Brownian motion. Vary the parameters and note the size and location of the mean \pm standard deviation bar for X_t . For various values of the parameter, run the simulation 1000 times and compare the empirical mean and standard deviation to the true mean and standard deviation.

Properties

The parameter $\mu - \sigma^2/2$ determines the asymptotic behavior of geometric Brownian motion.

Asymptotic behavior:

1. If $\mu > \sigma^2/2$ then $X_t \rightarrow \infty$ as $t \rightarrow \infty$ with probability 1.
2. If $\mu < \sigma^2/2$ then $X_t \rightarrow 0$ as $t \rightarrow \infty$ with probability 1.
3. If $\mu = \sigma^2/2$ then X_t has no limit as $t \rightarrow \infty$ with probability 1.

Proof

These results follow from the law of the iterative logarithm. Asymptotically, the term $(\mu - \sigma^2/2)t$ dominates the term σZ_t as $t \rightarrow \infty$.

It's interesting to compare this result with the asymptotic behavior of the mean function, given above, which depends only on the parameter μ . When the drift parameter is 0, geometric Brownian motion is a martingale.

If $\mu = 0$, geometric Brownian motion \mathbf{X} is a martingale with respect to the underlying Brownian motion \mathbf{Z} .

Proof from stochastic integrals

This is the simplest proof. When $\mu = 0$, \mathbf{X} satisfies the stochastic differential equation $dX_t = \sigma X_t dZ_t$ and therefore

$$X_t = 1 + \sigma \int_0^t X_s dZ_s, \quad t \geq 0 \quad (18.4.8)$$

The process associated with a stochastic integral is always a martingale, assuming the usual assumptions on the integrand process (which are satisfied here).

Direct proof

Let $\mathcal{F}_t = \sigma\{Z_s : 0 \leq s \leq t\}$ for $t \in [0, \infty)$, so that $\mathfrak{F} = \{\mathcal{F}_t : t \in [0, \infty)\}$ is the natural filtration associated with \mathbf{Z} . Let $s, t \in [0, \infty)$ with $s \leq t$. We use our usual trick of writing $Z_t = Z_s + (Z_t - Z_s)$, to take advantage of the stationary and independent increments properties of Brownian motion. Thus,

$$X_t = \exp\left[-\frac{\sigma^2}{2}t + \sigma Z_s + \sigma(Z_t - Z_s)\right] \quad (18.4.9)$$

Since Z_s is measurable with respect to \mathcal{F}_s and $Z_t - Z_s$ is independent of \mathcal{F}_s we have

$$\mathbb{E}(X_t | \mathcal{F}_s) = \exp\left(-\frac{\sigma^2}{2}t + \sigma Z_s\right) \mathbb{E}\{\exp[\sigma(Z_t - Z_s)]\} \quad (18.4.10)$$

But $Z_t - Z_s$ has the normal distribution with mean 0 and variance $t - s$, so from the formula for the moment generating function of the normal distribution, we have

$$\mathbb{E}\{\exp[\sigma(Z_t - Z_s)]\} = \exp\left[\frac{\sigma^2}{2}(t - s)\right] \quad (18.4.11)$$

Substituting gives

$$\mathbb{E}(X_t | \mathcal{F}_s) = \exp\left(-\frac{\sigma^2}{2}s + \sigma Z_s\right) = X_s \quad (18.4.12)$$

This page titled [18.4: Geometric Brownian Motion](#) is shared under a [CC BY 2.0](#) license and was authored, remixed, and/or curated by [Kyle Siegrist \(Random Services\)](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform.

Index

A

arcsine distribution
[5.19: The Arcsine Distribution](#)

B

Bayesian estimation
[7.4: Bayesian Estimation](#)
Benford's law
[5.39: Benford's Law](#)
Bernoulli Trials
[11.1: Introduction to Bernoulli Trials](#)
Bertrand's paradox
[10.2: Bertrand's Paradox](#)
Bertrand's problem
[10.2: Bertrand's Paradox](#)
beta distribution
[5.17: The Beta Distribution](#)
beta function
[5.17: The Beta Distribution](#)
beta prime distribution
[5.18: The Beta Prime Distribution](#)
binomial distribution
[2.5: Independence](#)
bold play
[13.10: Bold Play](#)
bonus number
[13.7: Lotteries](#)
Brownian motion
[18: Brownian Motion](#)
[18.1: Standard Brownian Motion](#)

C

Chebyshev's Inequality
[4.3: Variance](#)
chi distribution
[5.9: Chi-Square and Related Distribution](#)
conditional expected value
[4.7: Conditional Expected Value](#)
conditional probability
[2.4: Conditional Probability](#)
continuous distributions
[3.2: Continuous Distributions](#)
correlation
[2.4: Conditional Probability](#)
craps
[13.4: Craps](#)

D

disjointness
[2.5: Independence](#)
distribution function
[3.6: Distribution and Quantile Functions](#)
[3.9: General Distribution Functions](#)

E

Ehrenfest chains
[16.8: The Ehrenfest Chains](#)
equivalence class
[1.5: Equivalence Relations](#)
extremal elements
[1.4: Partial Orders](#)

F

Feller Markov processes
[16.2: Potentials and Generators for General Markov Processes](#)

G

gamma distribution
[5.8: The Gamma Distribution](#)
[14.3: The Gamma Distribution](#)
gamma function
[5.8: The Gamma Distribution](#)
geometric distribution
[11.3: The Geometric Distribution](#)

H

Hasse graph
[1.4: Partial Orders](#)

I

incomplete beta function
[5.17: The Beta Distribution](#)
independence
[2.5: Independence](#)
infinitely divisible distributions
[5.4: Infinitely Divisible Distributions](#)
isomorphism
[1.4: Partial Orders](#)

J

joint distributions
[3.4: Joint Distributions](#)

K

keno
[13.7: Lotteries](#)
kernels
[4.13: Kernels and Operators](#)
Kolmogorov axioms
[2.3: Probability Measures](#)
kurtosis
[4.4: Skewness and Kurtosis](#)

L

Lévy distribution
[5.16: The Lévy Distribution](#)
law of the iterated logarithm
[18.1: Standard Brownian Motion](#)
likelihood ratio
[9.5: Likelihood Ratio Tests](#)
lottery
[13.7: Lotteries](#)

M

Martingales
[17: Martingales](#)
matching problem
[12.5: The Matching Problem](#)
metric spaces
[1.10: Metric Spaces](#)
moments
[7.2: The Method of Moments](#)

Monty Hall problem

[13.6: The Monty Hall Problem](#)
multivariate normal distribution
[5.7: The Multivariate Normal Distribution](#)

N

normal distribution
[5.6: The Normal Distribution](#)
normal model
[9.2: Tests in the Normal Model](#)

O

operators
[4.13: Kernels and Operators](#)

P

Pólya's urn process
[12.8: Pólya's Urn Process](#)
[17.1: Introduction to Martingales](#)
Pareto distribution
[3.2: Continuous Distributions](#)
[5.36: The Pareto Distribution](#)
partial orders
[1.4: Partial Orders](#)
poker dice
[13.3: Simple Dice Games](#)
power series distributions
[5.5: Power Series Distributions](#)

Q

queuing theory
[16.22: Continuous-Time Queuing Chains](#)

R

random experiment
[2.1: Random Experiments](#)
Rayleigh distribution
[5.14: The Rayleigh Distribution](#)
Red and Black game
[13.8: The Red and Black Game](#)
roulette
[13.5: Roulette](#)

S

second central moment
[4.3: Variance](#)
second moment
[4.3: Variance](#)
skewness
[4.4: Skewness and Kurtosis](#)
stable distributions
[5.3: Stable Distributions](#)
standard triangle distribution
[5.24: The Triangle Distribution](#)
Stirling's approximation
[5.8: The Gamma Distribution](#)
strong law of large numbers
[2.6: Convergence](#)
Student t distribution
[5.10: The Student t Distribution](#)

T

The F distribution

[5.11: The F Distribution](#)

timid play

[13.9: Timid Play](#)

V

variance

[4.3: Variance](#)

W

Wald distribution

[5.37: The Wald Distribution](#)

weak law of large numbers

[2.6: Convergence](#)

Weibull distribution

[5.38: The Weibull Distribution](#)

Wiener process

[18.1: Standard Brownian Motion](#)

Z

zeta distribution

[5.40: The Zeta Distribution](#)

Zipf distribution

[5.40: The Zeta Distribution](#)

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: Probability, Mathematical Statistics, and Stochastic Processes (Siegrist)

Webpages: 226

All licenses found:

- [CC BY 2.0](#): 98.7% (223 pages)
- [Undeclared](#): 1.3% (3 pages)

By Page

- Probability, Mathematical Statistics, and Stochastic Processes (Siegrist) - [CC BY 2.0](#)
 - Front Matter - [CC BY 2.0](#)
 - TitlePage - [CC BY 2.0](#)
 - InfoPage - [CC BY 2.0](#)
 - Introduction - [CC BY 2.0](#)
 - Table of Contents - [Undeclared](#)
 - Licensing - [Undeclared](#)
 - Table of Contents - [CC BY 2.0](#)
 - Object Library - [CC BY 2.0](#)
 - Credits - [CC BY 2.0](#)
 - Sources and Resources - [CC BY 2.0](#)
 - 1: Foundations - [CC BY 2.0](#)
 - 1.1: Sets - [CC BY 2.0](#)
 - 1.2: Functions - [CC BY 2.0](#)
 - 1.3: Relations - [CC BY 2.0](#)
 - 1.4: Partial Orders - [CC BY 2.0](#)
 - 1.5: Equivalence Relations - [CC BY 2.0](#)
 - 1.6: Cardinality - [CC BY 2.0](#)
 - 1.7: Counting Measure - [CC BY 2.0](#)
 - 1.8: Combinatorial Structures - [CC BY 2.0](#)
 - 1.9: Topological Spaces - [CC BY 2.0](#)
 - 1.10: Metric Spaces - [CC BY 2.0](#)
 - 1.11: Measurable Spaces - [CC BY 2.0](#)
 - 1.12: Special Set Structures - [CC BY 2.0](#)
 - 2: Probability Spaces - [CC BY 2.0](#)
 - 2.1: Random Experiments - [CC BY 2.0](#)
 - 2.2: Events and Random Variables - [CC BY 2.0](#)
 - 2.3: Probability Measures - [CC BY 2.0](#)
 - 2.4: Conditional Probability - [CC BY 2.0](#)
 - 2.5: Independence - [CC BY 2.0](#)
 - 2.6: Convergence - [CC BY 2.0](#)
 - 2.7: Measure Spaces - [CC BY 2.0](#)
 - 2.8: Existence and Uniqueness - [CC BY 2.0](#)
 - 2.9: Probability Spaces Revisited - [CC BY 2.0](#)
 - 2.10: Stochastic Processes - [CC BY 2.0](#)
 - 2.11: Filtrations and Stopping Times - [CC BY 2.0](#)
 - 3: Distributions - [CC BY 2.0](#)
 - 3.1: Discrete Distributions - [CC BY 2.0](#)
 - 3.2: Continuous Distributions - [CC BY 2.0](#)
 - 3.3: Mixed Distributions - [CC BY 2.0](#)
 - 3.4: Joint Distributions - [CC BY 2.0](#)
 - 3.5: Conditional Distributions - [CC BY 2.0](#)
 - 3.6: Distribution and Quantile Functions - [CC BY 2.0](#)
 - 3.7: Transformations of Random Variables - [CC BY 2.0](#)
 - 3.8: Convergence in Distribution - [CC BY 2.0](#)
 - 3.9: General Distribution Functions - [CC BY 2.0](#)
 - 3.10: The Integral With Respect to a Measure - [CC BY 2.0](#)
 - 3.11: Properties of the Integral - [CC BY 2.0](#)
 - 3.12: General Measures - [CC BY 2.0](#)
 - 3.13: Absolute Continuity and Density Functions - [CC BY 2.0](#)
 - 3.14: Function Spaces - [CC BY 2.0](#)
 - 4: Expected Value - [CC BY 2.0](#)
 - 4.1: Definitions and Basic Properties - [CC BY 2.0](#)
 - 4.2: Additional Properties - [CC BY 2.0](#)
 - 4.3: Variance - [CC BY 2.0](#)
 - 4.4: Skewness and Kurtosis - [CC BY 2.0](#)
 - 4.5: Covariance and Correlation - [CC BY 2.0](#)
 - 4.6: Generating Functions - [CC BY 2.0](#)
 - 4.7: Conditional Expected Value - [CC BY 2.0](#)
 - 4.8: Expected Value and Covariance Matrices - [CC BY 2.0](#)
 - 4.9: Expected Value as an Integral - [CC BY 2.0](#)
 - 4.10: Conditional Expected Value Revisited - [CC BY 2.0](#)
 - 4.11: Vector Spaces of Random Variables - [CC BY 2.0](#)
 - 4.12: Uniformly Integrable Variables - [CC BY 2.0](#)
 - 4.13: Kernels and Operators - [CC BY 2.0](#)
 - 5: Special Distributions - [CC BY 2.0](#)
 - 5.1: Location-Scale Families - [CC BY 2.0](#)
 - 5.2: General Exponential Families - [CC BY 2.0](#)
 - 5.3: Stable Distributions - [CC BY 2.0](#)
 - 5.4: Infinitely Divisible Distributions - [CC BY 2.0](#)
 - 5.5: Power Series Distributions - [CC BY 2.0](#)

- 5.6: The Normal Distribution - CC BY 2.0
- 5.7: The Multivariate Normal Distribution - CC BY 2.0
- 5.8: The Gamma Distribution - CC BY 2.0
- 5.9: Chi-Square and Related Distribution - CC BY 2.0
- 5.10: The Student t Distribution - CC BY 2.0
- 5.11: The F Distribution - CC BY 2.0
- 5.12: The Lognormal Distribution - CC BY 2.0
- 5.13: The Folded Normal Distribution - CC BY 2.0
- 5.14: The Rayleigh Distribution - CC BY 2.0
- 5.15: The Maxwell Distribution - CC BY 2.0
- 5.16: The Lévy Distribution - CC BY 2.0
- 5.17: The Beta Distribution - CC BY 2.0
- 5.18: The Beta Prime Distribution - CC BY 2.0
- 5.19: The Arcsine Distribution - CC BY 2.0
- 5.20: General Uniform Distributions - CC BY 2.0
- 5.21: The Uniform Distribution on an Interval - CC BY 2.0
- 5.22: Discrete Uniform Distributions - CC BY 2.0
- 5.23: The Semicircle Distribution - CC BY 2.0
- 5.24: The Triangle Distribution - CC BY 2.0
- 5.25: The Irwin-Hall Distribution - CC BY 2.0
- 5.26: The U-Power Distribution - CC BY 2.0
- 5.27: The Sine Distribution - CC BY 2.0
- 5.28: The Laplace Distribution - CC BY 2.0
- 5.29: The Logistic Distribution - CC BY 2.0
- 5.30: The Extreme Value Distribution - CC BY 2.0
- 5.31: The Hyperbolic Secant Distribution - CC BY 2.0
- 5.32: The Cauchy Distribution - CC BY 2.0
- 5.33: The Exponential-Logarithmic Distribution - CC BY 2.0
- 5.34: The Gompertz Distribution - CC BY 2.0
- 5.35: The Log-Logistic Distribution - CC BY 2.0
- 5.36: The Pareto Distribution - CC BY 2.0
- 5.37: The Wald Distribution - CC BY 2.0
- 5.38: The Weibull Distribution - CC BY 2.0
- 5.39: Benford's Law - CC BY 2.0
- 5.40: The Zeta Distribution - CC BY 2.0
- 5.41: The Logarithmic Series Distribution - CC BY 2.0
- 6: Random Samples - CC BY 2.0
 - 6.1: Introduction - CC BY 2.0
 - 6.2: The Sample Mean - CC BY 2.0
 - 6.3: The Law of Large Numbers - CC BY 2.0
 - 6.4: The Central Limit Theorem - CC BY 2.0
 - 6.5: The Sample Variance - CC BY 2.0
 - 6.6: Order Statistics - CC BY 2.0
 - 6.7: Sample Correlation and Regression - CC BY 2.0
 - 6.8: Special Properties of Normal Samples - CC BY 2.0
- 7: Point Estimation - CC BY 2.0
 - 7.1: Estimators - CC BY 2.0
 - 7.2: The Method of Moments - CC BY 2.0
 - 7.3: Maximum Likelihood - CC BY 2.0
 - 7.4: Bayesian Estimation - CC BY 2.0
 - 7.5: Best Unbiased Estimators - CC BY 2.0
 - 7.6: Sufficient, Complete and Ancillary Statistics - CC BY 2.0
- 8: Set Estimation - CC BY 2.0
 - 8.1: Introduction to Set Estimation - CC BY 2.0
 - 8.2: Estimation the Normal Model - CC BY 2.0
 - 8.3: Estimation in the Bernoulli Model - CC BY 2.0
 - 8.4: Estimation in the Two-Sample Normal Model - CC BY 2.0
 - 8.5: Bayesian Set Estimation - CC BY 2.0
- 9: Hypothesis Testing - CC BY 2.0
 - 9.1: Introduction to Hypothesis Testing - CC BY 2.0
 - 9.2: Tests in the Normal Model - CC BY 2.0
 - 9.3: Tests in the Bernoulli Model - CC BY 2.0
 - 9.4: Tests in the Two-Sample Normal Model - CC BY 2.0
 - 9.5: Likelihood Ratio Tests - CC BY 2.0
 - 9.6: Chi-Square Tests - CC BY 2.0
- 10: Geometric Models - CC BY 2.0
 - 10.1: Buffon's Problems - CC BY 2.0
 - 10.2: Bertrand's Paradox - CC BY 2.0
 - 10.3: Random Triangles - CC BY 2.0
- 11: Bernoulli Trials - CC BY 2.0
 - 11.1: Introduction to Bernoulli Trials - CC BY 2.0
 - 11.2: The Binomial Distribution - CC BY 2.0
 - 11.3: The Geometric Distribution - CC BY 2.0
 - 11.4: The Negative Binomial Distribution - CC BY 2.0
 - 11.5: The Multinomial Distribution - CC BY 2.0
 - 11.6: The Simple Random Walk - CC BY 2.0
 - 11.7: The Beta-Bernoulli Process - CC BY 2.0
- 12: Finite Sampling Models - CC BY 2.0
 - 12.1: Introduction to Finite Sampling Models - CC BY 2.0
 - 12.2: The Hypergeometric Distribution - CC BY 2.0
 - 12.3: The Multivariate Hypergeometric Distribution - CC BY 2.0
 - 12.4: Order Statistics - CC BY 2.0
 - 12.5: The Matching Problem - CC BY 2.0
 - 12.6: The Birthday Problem - CC BY 2.0
 - 12.7: The Coupon Collector Problem - CC BY 2.0
 - 12.8: Pólya's Urn Process - CC BY 2.0
 - 12.9: The Secretary Problem - CC BY 2.0
- 13: Games of Chance - CC BY 2.0
 - 13.1: Introduction to Games of Chance - CC BY 2.0
 - 13.2: Poker - CC BY 2.0
 - 13.3: Simple Dice Games - CC BY 2.0

- 13.4: Craps - *CC BY 2.0*
- 13.5: Roulette - *CC BY 2.0*
- 13.6: The Monty Hall Problem - *CC BY 2.0*
- 13.7: Lotteries - *CC BY 2.0*
- 13.8: The Red and Black Game - *CC BY 2.0*
- 13.9: Timid Play - *CC BY 2.0*
- 13.10: Bold Play - *CC BY 2.0*
- 13.11: Optimal Strategies - *CC BY 2.0*
- 14: The Poisson Process - *CC BY 2.0*
 - 14.1: Introduction to the Poisson Process - *CC BY 2.0*
 - 14.2: The Exponential Distribution - *CC BY 2.0*
 - 14.3: The Gamma Distribution - *CC BY 2.0*
 - 14.4: The Poisson Distribution - *CC BY 2.0*
 - 14.5: Thinning and Superposition - *CC BY 2.0*
 - 14.6: Non-homogeneous Poisson Processes - *CC BY 2.0*
 - 14.7: Compound Poisson Processes - *CC BY 2.0*
 - 14.8: Poisson Processes on General Spaces - *CC BY 2.0*
- 15: Renewal Processes - *CC BY 2.0*
 - 15.1: Introduction - *CC BY 2.0*
 - 15.2: Renewal Equations - *CC BY 2.0*
 - 15.3: Renewal Limit Theorems - *CC BY 2.0*
 - 15.4: Delayed Renewal Processes - *CC BY 2.0*
 - 15.5: Alternating Renewal Processes - *CC BY 2.0*
 - 15.6: Renewal Reward Processes - *CC BY 2.0*
- 16: Markov Processes - *CC BY 2.0*
 - 16.1: Introduction to Markov Processes - *CC BY 2.0*
 - 16.2: Potentials and Generators for General Markov Processes - *CC BY 2.0*
 - 16.3: Introduction to Discrete-Time Chains - *CC BY 2.0*
 - 16.4: Transience and Recurrence for Discrete-Time Chains - *CC BY 2.0*
 - 16.5: Periodicity of Discrete-Time Chains - *CC BY 2.0*
 - 16.6: Stationary and Limiting Distributions of Discrete-Time Chains - *CC BY 2.0*
 - 16.7: Time Reversal in Discrete-Time Chains - *CC BY 2.0*
 - 16.8: The Ehrenfest Chains - *CC BY 2.0*
 - 16.9: The Bernoulli-Laplace Chain - *CC BY 2.0*
 - 16.10: Discrete-Time Reliability Chains - *CC BY 2.0*
 - 16.11: Discrete-Time Branching Chain - *CC BY 2.0*
 - 16.12: Discrete-Time Queuing Chains - *CC BY 2.0*
 - 16.13: Discrete-Time Birth-Death Chains - *CC BY 2.0*
 - 16.14: Random Walks on Graphs - *CC BY 2.0*
 - 16.15: Introduction to Continuous-Time Markov Chains - *CC BY 2.0*
 - 16.16: Transition Matrices and Generators of Continuous-Time Chains - *CC BY 2.0*
 - 16.17: Potential Matrices - *CC BY 2.0*
 - 16.18: Stationary and Limiting Distributions of Continuous-Time Chains - *CC BY 2.0*
 - 16.19: Time Reversal in Continuous-Time Chains - *CC BY 2.0*
 - 16.20: Chains Subordinate to the Poisson Process - *CC BY 2.0*
 - 16.21: Continuous-Time Birth-Death Chains - *CC BY 2.0*
 - 16.22: Continuous-Time Queuing Chains - *CC BY 2.0*
 - 16.23: Continuous-Time Branching Chains - *CC BY 2.0*
- 17: Martingales - *CC BY 2.0*
 - 17.1: Introduction to Martingales - *CC BY 2.0*
 - 17.2: Properties and Constructions - *CC BY 2.0*
 - 17.3: Stopping Times - *CC BY 2.0*
 - 17.4: Inequalities - *CC BY 2.0*
 - 17.5: Convergence - *CC BY 2.0*
 - 17.6: Backwards Martingales - *CC BY 2.0*
- 18: Brownian Motion - *CC BY 2.0*
 - 18.1: Standard Brownian Motion - *CC BY 2.0*
 - 18.2: Brownian Motion with Drift and Scaling - *CC BY 2.0*
 - 18.3: The Brownian Bridge - *CC BY 2.0*
 - 18.4: Geometric Brownian Motion - *CC BY 2.0*
- Back Matter - *CC BY 2.0*
 - Index - *CC BY 2.0*
 - Glossary - *CC BY 2.0*
 - Detailed Licensing - *Undeclared*