

MATH 345 - PROBABILITY (KUTER)



Kristin Kuter
Saint Mary's College

Saint Mary's College
MATH 345 - Probability (Kuter)

Kristin Kuter

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 03/18/2025

TABLE OF CONTENTS

1: What is Probability?

- 1.1: Sample Spaces and Events
- 1.2: Probability Measures

2: Computing Probabilities

- 2.1: Equally Likely Outcomes and Counting Techniques (Combinatorics)
- 2.2: Conditional Probability and Bayes' Rule
- 2.3: Independent Events

3: Discrete Random Variables

- 3.1: Introduction to Random Variables
- 3.2: Probability Mass Functions (PMFs) and Cumulative Distribution Functions (CDFs) for Discrete Random Variables
- 3.3: Bernoulli and Binomial Distributions
- 3.4: Hypergeometric, Geometric, and Negative Binomial Distributions
- 3.5: Poisson Distribution
- 3.6: Expected Value of Discrete Random Variables
- 3.7: Variance of Discrete Random Variables
- 3.8: Moment-Generating Functions (MGFs) for Discrete Random Variables

4: Continuous Random Variables

- 4.1: Probability Density Functions (PDFs) and Cumulative Distribution Functions (CDFs) for Continuous Random Variables
- 4.2: Expected Value and Variance of Continuous Random Variables
- 4.3: Uniform Distributions
- 4.4: Normal Distributions
- 4.5: Exponential and Gamma Distributions
- 4.6: Weibull Distributions
- 4.7: Chi-Squared Distributions
- 4.8: Beta Distributions

5: Probability Distributions for Combinations of Random Variables

- 5.1: Joint Distributions of Discrete Random Variables
- 5.2: Joint Distributions of Continuous Random Variables
- 5.3: Conditional Probability Distributions
- 5.4: Finding Distributions of Functions of Continuous Random Variables
- 5.5: Sample Mean

Licensing

Index

Glossary

Detailed Licensing

Licensing

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

CHAPTER OVERVIEW

1: What is Probability?

Topic hierarchy

[1.1: Sample Spaces and Events](#)

[1.2: Probability Measures](#)

This page titled [1: What is Probability?](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

1.1: Sample Spaces and Events

Introduction

We begin with a definition.

Definition 1.1.1

Probability theory provides a mathematical model for chance (or random) phenomena.

While this is not a very informative definition, it does indicate the overall goal of this course, which is to develop a *formal, mathematical structure* for the fairly intuitive concept of probability. While most everyone is familiar with the notion of "chance" - we informally talk about the chance of it raining tomorrow, or the chance of getting what you want for your birthday -- when it comes to *quantifying* the chance of something happening, we need to develop a mathematical model to make things precise and calculable.

Sample Spaces and Events

Before we can formally define what the mathematical model is that we will use to make probability precise, we first establish the *structure* on which the model operates: *sample spaces* and *events*.

Definition 1.1.2

The **sample space** for a probability experiment (i.e., an experiment with random outcomes) is the set of all possible outcomes.

- The sample space is denoted S .
- An **outcome** is an *element* of S , generally denoted $s \in S$.

Example 1.1.1

Suppose we toss a coin twice and record the sequence of heads (h) and tails (t). A possible outcome of this experiment is then given by

$$s = ht$$

and the sample space is

$$S = \{hh, ht, th, tt\}. \quad (1.1.1)$$

Example 1.1.2

Suppose we record the time (t), in minutes, that a car spends waiting for a green light at a particular intersection. A possible outcome of this experiment is then given by

$$t = 1.5,$$

indicating that a particular car waited one and a half minutes for the light to turn green. The sample space consists of all non-negative numbers, since a measurement of time cannot be negative and, in theory, there is no limit on how long a car could wait for a green light. We can then write the sample space as follows:

$$S = \{t \in \mathbb{R} \mid t \geq 0\} = [0, \infty). \quad (1.1.2)$$

Definition 1.1.3

An **event** is a particular subset of the sample space.

Example 1.1.3

Continuing in the context of [Example 1.1.1](#), define A to be the event that at least one heads is recorded. We can write event A as the following subset of the sample space:


$$A = \{hh, ht, th\}.$$

Note that A is a subset of S given in Equation 1.1.1.

Example 1.1.4

Continuing in the context of Example 1.1.2, define B to be the event that a car waits at most 2 minutes for the light to turn green. We can write the event B as the following interval, i.e., a subset of the sample space S given in Equation 1.1.2:

$$B = [0, 2] = \{t \in \mathbb{R} \mid 0 \leq t \leq 2\}.$$



https://youtu.be/Y_GAW22n1Ms

Set Theory: A Brief Review

As we see from the above definitions of sample spaces and events, *sets* play the primary role in the structure of probability experiments. So, in this section, we review some of the basic definitions and notation from *set theory*. We do this in the context of sample spaces, outcomes, and events.

Definition 1.1.4

1. The **union** of two events A and B , denoted $A \cup B$, is the set of all outcomes in A or B (or both).
2. The **intersection** of two events A and B , denoted $A \cap B$, is the set of all outcomes in both A and B .

3. The **complement** of an event A , denoted A^c , is the set of all outcomes in the sample space that are not in A . This may also be written as follows:

$$A^c = \{s \in S \mid s \notin A\}.$$

4. The **empty set**, denoted \emptyset , is the set containing no outcomes.
 5. Two events A and B are **disjoint** (or **mutually exclusive**) if their intersection is the empty set, i.e., $A \cap B = \emptyset$.

Example 1.1.5

Continuing in the context of both [Examples 1.1.1](#) & [1.1.3](#), define B to be the event that exactly one heads is recorded:

$$B = \{ht, th\}.$$

Now we can apply the set operations just defined to the events A and B :

$$A \cup B = \{hh, ht, th\} = A$$

$$A \cap B = \{ht, th\} = B$$

$$A^c = \{tt\}$$

$$B^c = \{hh, tt\}$$

Note the relationship between events A and B : every outcome in B is an outcome in A . In this case, we say that B is a **subset** of A , and write

$$B \subseteq A.$$

Note also that events A and B are not disjoint, since their intersection is not the empty set. However, if we let C be the event that no heads are recorded, then


$$C = \{tt\},$$


and

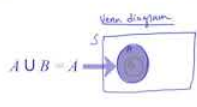
$$A \cap C = \emptyset$$

$$B \cap C = \emptyset.$$

Thus, events A and C are disjoint, and events B and C are disjoint.





ig Set Operations
 1. $\Omega = \{t_h, t_r, t_b, t_g\}$ & $A = \{t_r, t_b, t_g\}$
 text of Examples 1.1.1 & 1.1.3, define B to be the event
 is recorded:
 $B = \{t_r, t_h\}$.
 set operations just defined to the events A and B :
 union: $A \cup B = \{t_h, t_r, t_b, t_g\} = A$
 intersection: $A \cap B = \{t_r, t_h\} = B$ \Rightarrow subset: $B \subseteq A$
 " B is a subset of A "
 Venn diagram:

 $A \cup B = A$

https://youtu.be/Y_GAW22n1Ms

This page titled [1.1: Sample Spaces and Events](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

1.2: Probability Measures

Now we are ready to formally define probability.

Definition 1.2.1

A **probability measure** on the sample space S is a function, denoted P , from subsets of S to the real numbers \mathbb{R} , such that the following hold:

1. $P(S) = 1$
2. If A is any event in S , then $P(A) \geq 0$.
3. If events A_1 and A_2 are disjoint, then $P(A_1 \cup A_2) = P(A_1) + P(A_2)$.

More generally, if $A_1, A_2, \dots, A_n, \dots$ is a sequence of *pairwise disjoint* events, i.e., $A_i \cap A_j = \emptyset$, for every $i \neq j$, then

$$P(A_1 \cup A_2 \cup \dots \cup A_n \cup \dots) = P(A_1) + P(A_2) + \dots + P(A_n) + \dots$$

So essentially, we are defining probability to be an **operation** on the events of a sample space, which assigns numbers to events in such a way that the three properties stated in Definition 1.2.1 are satisfied.

Definition 1.2.1 is often referred to as the **axiomatic definition of probability**, where the three properties give the three **axioms** of probability. These three axioms are all we need to assume about the operation of probability in order for many other desirable properties of probability to hold, which we now state.

Properties of Probability Measures

Let S be a sample space with probability measure P . Also, let A and B be any events in S . Then the following hold.

1. $P(A^c) = 1 - P(A)$
2. $P(\emptyset) = 0$
3. If $A \subseteq B$, then $P(A) \leq P(B)$.
4. $P(A) \leq 1$
5. **Addition Law:** $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

Exercise 1.2.1

Can you prove the five properties of probability measures stated above using only the three axioms of probability measures stated in Definition 1.2.1?

Answer

- (1) For the first property, note that by definition of the complement of an event A we have

$$A \cup A^c = S \quad \text{and} \quad A \cap A^c = \emptyset.$$

In other words, given any event A , we can represent the sample space S as a disjoint union of A with its complement. Thus, by the first and third axioms, we derive the first property:

$$\begin{aligned} 1 = P(S) &= P(A \cup A^c) = P(A) + P(A^c) \\ &\Rightarrow P(A^c) = 1 - P(A) \end{aligned}$$

- (2) For the second property, note that we can write $S = S \cup \emptyset$, and that this is a disjoint union, since anything intersected with the empty set will necessarily be empty. So, using the first and third axioms, we derive the second property:

$$\begin{aligned} 1 = P(S) &= P(S \cup \emptyset) = P(S) + P(\emptyset) = 1 + P(\emptyset) \\ &\Rightarrow P(\emptyset) = 0 \end{aligned}$$

(3) For the third property, note that we can write $B = A \cup (B \cap A^c)$, and that this is a disjoint union, since A and A^c are disjoint. By the third axiom, we have

$$P(B) = P(A \cup (B \cap A^c)) = P(A) + P(B \cap A^c). \quad (1.2.1)$$

By the second axiom, we know that $P(B \cap A^c) \geq 0$. Thus, if we remove it from the right-hand side of Equation 1.2.1, we are left with something smaller, which proves the third property:

$$P(B) = P(A) + P(B \cap A^c) \geq P(A) \Rightarrow P(B) \geq P(A)$$

(4) For the fourth property, we will use the third property that we just proved. By definition, any event A is a subset of the sample space S , i.e., $A \subseteq S$. Thus, by the third property and the first axiom, we derive the fourth property:

$$P(A) \leq P(S) = 1 \Rightarrow P(A) \leq 1$$

(5) For the fifth property, note that we can write the union of events A and B as the union of the following two disjoint events:

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

in other words, the union of A and B is given by the union of all the outcomes in A with all the outcomes in B that are *not* in A . Furthermore, note that event B can be written as the union the following two disjoint events:

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

in other words, B is written as the disjoint union of all the outcomes in B that are also in A with the outcomes in B that are *not* in A . We can use this expression for B to find an expression for $P(A^c \cap B)$ to substitute in the expression for $A \cup B$ in order to derive the fifth property:

$$P(B) = P(A \cap B) + P(A^c \cap B) \Rightarrow P(A^c \cap B) = P(B) - P(A \cap B)$$

$$P(A \cup B) = P(A) + P(A^c \cap B) \Rightarrow P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

Note that the axiomatic definition (Definition 1.2.1) does not tell us how to *compute* probabilities. It simply defines a formal, mathematical behavior of probability. In other words, the axiomatic definition describes how probability should theoretically *behave* when applied to events. To compute probabilities, we use the properties stated above, as the next example demonstrates.

Example 1.2.1

Continuing in the context of Example 1.1.5, let's define a probability measure on S . Assuming that the coin we toss is *fair*, then the outcomes in S are **equally likely**, meaning that each outcome has the *same probability* of occurring. Since there are four outcomes, and we know that probability of the sample space must be 1 (first axiom of probability in Definition 1.2.1), it follows that the probability of each outcome is $\frac{1}{4} = 0.25$.

So, we can write

$$P(hh) = P(ht) = P(th) = P(tt) = 0.25.$$

The reader can verify that this defines a probability measure satisfying the three axioms.

With this probability measure on the outcomes we can now compute the probability of any event in S by simply *counting* the number of outcomes in the event. Thus, we find the probability of events A and B previously defined:

$$P(A) = P(\{hh, ht, th\}) = \frac{3}{4} = 0.75$$

$$P(B) = P(\{ht, th\}) = \frac{2}{4} = 0.50.$$

We consider the case of equally likely outcomes further in [Section 2.1](#).

There is another, more empirical, approach to defining probability, given by using *relative frequencies* and a version of the Law of Large Numbers.

Relative Frequency Approximation

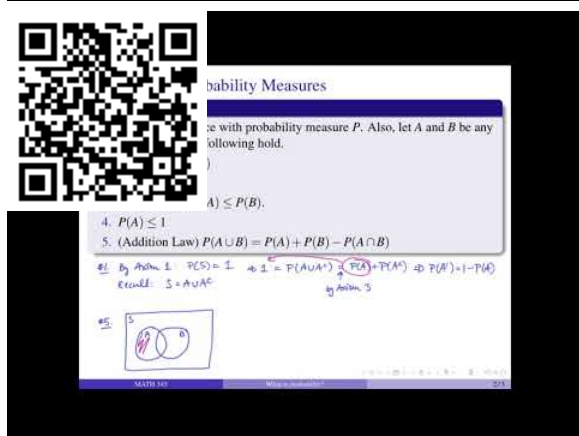

To *estimate* the probability of an event A , repeat the random experiment several times (each repetition is called a *trial*) and count the number of times A occurred, i.e., the number of times the resulting outcome is in A . Then, we approximate the probability of A using **relative frequency**:

$$P(A) \approx \frac{\text{number of times } A \text{ occurred}}{\text{number of trials}}.$$

Law of Large Numbers

As the number of trials increases, the relative frequency approximation approaches the theoretical value of $P(A)$.

This approach to defining probability is sometimes referred to as the **frequentist definition of probability**. Under this definition, probability represents a *long-run average*. The two approaches to defining probability are equivalent. It can be shown that using relative frequencies to define a probability measure satisfies the axiomatic definition.



https://youtu.be/Y_GAW22n1Ms

This page titled [1.2: Probability Measures](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

CHAPTER OVERVIEW

2: Computing Probabilities

[2.1: Equally Likely Outcomes and Counting Techniques \(Combinatorics\)](#)

[2.2: Conditional Probability and Bayes' Rule](#)

[2.3: Independent Events](#)

This page titled [2: Computing Probabilities](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

2.1: Equally Likely Outcomes and Counting Techniques (Combinatorics)

In this chapter, we consider the problem of assigning specific probabilities to outcomes in a sample space. As we saw in [Section 1.2](#), the axiomatic definition of probability ([Definition 1.2.1](#)) does not tell us how to compute probabilities. So in this section we consider the commonly encountered scenario referred to as *equally likely outcomes* and develop methods for computing probabilities in this special case.

Finite Sample Spaces

Before focusing on equally likely outcomes, we consider the more general case of *finite* sample spaces. In other words, suppose that a sample space S has a finite number of outcomes, which we can denote as N . In this case, we can represent the outcomes in S as follows:

$$S = \{s_1, s_2, \dots, s_N\}.$$

Suppose further that we denote the probability assigned to each outcome in S as $P(s_i) = p_i$, for $i = 1, \dots, N$. Then the probability of any event A in S is given by adding the probabilities corresponding to the outcomes contained in A and we can write

$$P(A) = \sum_{s_i \in A} p_i. \quad (2.1.1)$$

This follows from the third axiom of probability ([Definition 1.2.1](#)), since we can write any event as a disjoint union of the outcomes contained in the event. For example, if event A contains three outcomes, then we can write $A = \{s_1, s_2, s_3\} = \{s_1\} \cup \{s_2\} \cup \{s_3\}$. So the probability of A is given by simply summing up the probabilities assigned to s_1, s_2, s_3 . This fact will be useful in the special case of equally likely outcomes, which we consider next.

Equally Likely Outcomes

First, let's state a formal definition of what it means for the outcomes in a sample space to be *equally likely*.

Definition 2.1.1

The outcomes in a sample space S are ***equally likely*** if each outcome has the *same probability* of occurring.

In general, if outcomes in a sample space S are equally likely, then computing the probability of a single outcome or an event is very straightforward, as the following exercise demonstrates. You are encouraged to first try to answer the questions for yourself, and then click "Answer" to see the solution.

Exercise 2.1.1

Suppose that there are N outcomes in the sample space S and that the outcomes are equally likely.

- What is the probability of a single outcome in S ?
- What is the probability of an event A in S ?

Answer

First, note that we can represent the outcomes in S as follows:

$$S = \{s_1, s_2, \dots, s_N\}.$$

For each outcome in S , note that we can denote its probability as

$$P(s_i) = c, \text{ for } i = 1, 2, \dots, N,$$

where c is some constant. This follows from the fact that the outcomes of S are equally likely and so have the same probability of occurring. With this set-up and using the axioms of probability ([Definition 1.2.1](#)), we have the following:

$$\begin{aligned}
 1 = P(S) &= P(\{s_1\} \cup \cdots \cup \{s_N\}) \\
 &= P(s_1) + \cdots + P(s_N) \\
 &= c + \cdots + c \\
 &= N \times c \\
 \Rightarrow c &= \frac{1}{N}.
 \end{aligned}$$

Thus, the probability of a single outcome is given by 1 divided by the number of outcomes in S .

Now, for an event A in S , suppose it has n outcomes, where n is an integer such that $0 \leq n \leq N$. We can represent the outcomes in A as follows:

$$A = \{a_1, \dots, a_n\}.$$

Using Equation 2.1.1, we compute the probability of A as follows:

$$\begin{aligned}
 P(A) &= \sum_{i=1}^n P(a_i) = \sum_{i=1}^n \frac{1}{N} \\
 &= \frac{1}{N} + \cdots + \frac{1}{N} \\
 &= n \left(\frac{1}{N} \right) \\
 &= \frac{n}{N}.
 \end{aligned}$$



Thus, the probability of an event in S is equal to the number of outcomes in the event divided by the total number of outcomes in S .

We have already seen an example of a sample space with equally likely outcomes in [Example 1.2.1](#). You are encouraged to revisit that example and connect it to the results of Exercise 2.1.1.

In general, Exercise 2.1.1 shows that if a sample space S has equally likely outcomes, then the probability of an event A in the sample space is given by

$$P(A) = \frac{\text{number of outcomes in } A}{\text{number of outcomes in } S}. \quad (2.1.2)$$

From this result, we see that in the context of equally likely outcomes calculating probabilities of events reduces to simply *counting* the number of outcomes in the event and the sample space. So, we take a break from our discussion of probability, and briefly introduce some counting techniques.



<https://youtu.be/YdWaiH3oTq8>

Counting Techniques

First, let's consider the general context of performing multi-step experiments. The following tells us how to count the number of outcomes in such scenarios.

Multiplication Principle

If one probability experiment has m outcomes and another has n outcomes, then there are $m \times n$ total outcomes for the two experiments.

More generally, if there are k many probability experiments with the first experiment having n_1 outcomes, the second with n_2 , etc., then there are $n_1 \times n_2 \times \cdots \times n_k$ total outcomes for the k experiments.

Example 2.1.1

To demonstrate the Multiplication Principle, consider again the example of tossing a coin twice (see [Example 1.2.1](#)). Each toss is a probability experiment and on each toss, there are two possible outcomes: h or t . Thus, for two tosses, there are $2 \times 2 = 4$ total outcomes.

If we toss the coin a third time, there are $2 \times 2 \times 2 = 8$ total outcomes.

Next we define two commonly encountered situations, *permutations* and *combinations*, and consider how to count the number of ways in which they can occur.

Definition 2.1.2

A **permutation** is an *ordered* arrangement of objects. For example, "MATH" is a permutation of four letters from the alphabet.

A **combination** is an *unordered* collection of r objects from n total objects. For example, a group of three students chosen from a class of 10 students.

In order to count the number of possible permutations in a given setting, the Multiplication Principle is applied. For example, if we want to know the number of possible permutations of the four letters in "MATH", we compute

$$4 \times 3 \times 2 \times 1 = 4! = 24,$$

since there are four letters to select for the first position, three letters for the second, two for the third, leaving only one letter for the last. In other words, we treat each letter selection as an experiment in a multi-step process.

Counting Permutations

The number of permutations of n distinct objects is given by the following:

$$n \times (n-1) \times \cdots \times 2 \times 1 = n!$$

Counting combinations is a little more complicated, since we are not interested in the order in which objects are selected and so the Multiplication Principle does not directly apply. Consider the example that a group of three students are chosen from a class of 10. The group is the same regardless of the order in which the three students are selected. This implies that if we want to count the number of possible combinations, we need to be careful not to include permutations, i.e., rearrangements, of a certain selection. This leads to the following result that the number of possible combinations of size r selected from a total of n objects is given by binomial coefficients.

Counting Combinations

The number of combinations of r objects selected without replacement from n distinct objects is given by

$$\binom{n}{r} = \frac{n!}{r! \times (n-r)!}.$$

Note that $\binom{n}{r}$, read as " n choose r ", is also referred to as a **binomial coefficient**, since it appears in the [Binomial Theorem](#).

Using the above, we can compute the number of possible ways to select three students from a class of 10:

$$\binom{10}{3} = \frac{10!}{3! \times 7!} = 120$$

Example 2.1.2

Consider the example of tossing a coin three times. Note that an outcome is a sequence of heads and tails. Suppose that we are interested in the number of outcomes with exactly two heads, not in the actual sequence. To find the number of outcomes with exactly two heads, we need to determine the number of ways to select positions in the sequence for the heads, then the remaining position will be a tails. If we toss the coin three times, there are three positions to select from, and we want to select two. Since the order that we make the selection of placements does not matter, we are counting the number of combinations of 2 positions from a total of 3 positions, i.e.,

$$\binom{3}{2} = \frac{3!}{2! \times 1!} = 3.$$

Of course, this example is small enough that we could have arrived at the answer of 3 using brute force by just listing the possibilities. However, if we toss the coin a higher number of times, say 50, then the brute force approach becomes infeasible and we need to make use of binomial coefficients.

Strategies for Analyzing a Counting Problem

Before we return to our discussion of probability, the following outlines an approach for tackling problems in which we need to count the number of possible outcomes.

- ***Are there cases? To describe all possible outcomes, must one consider specific ways the desired event can occur?***

If so, make a note of each case and add the number of possibilities for each case.

Example: How many ways can a hand of 5 cards have at least 3 hearts?

Solution: To examine the ways this can happen, we observe that there are 3 specific ways a hand could have at least 3 hearts:

- (i) a hand could have exactly 3 hearts and 2 other cards (that are not hearts) OR
- (ii) a hand could have exactly 4 hearts and 1 other card (that is not hearts) OR
- (iii) a hand could have exactly 5 hearts.

So there are 3 cases. The total number of ways a hand of 5 cards has at least 3 hearts = the number of ways to have exactly 3 hearts and 2 other cards + the number of ways to have exactly 4 hearts and 1 other card + the number of ways to have 5 hearts.

Non-Example: How many groups of 10 students have 4 members from Indiana and 6 from Michigan? This event is clearly described and already very specific: 4 from one state and 6 from the other. There are no other options, possibilities, or cases.

- ***Are repetitions allowed?***

If the answer is yes, then use the Multiplication Rule.

- ***Are there steps?***

If so, put a slot for each step and a dot (multiplication symbol) between the slots (Multiplication Principle).

Example: How many ways can a hand of 5 cards have exactly 3 hearts?

Solution: There are two steps:

1. get 3 hearts AND
2. get 2 other cards (that are not hearts).

Then the number of hands = # of ways to get 3 hearts · # of ways to get 2 other cards (not hearts).

Non-Example: How many different groups of 4 books can be selected from a shelf that has 12 books? There is only one step: grab the 4 books.

- ***Does order matter?***

If so, use permutations or the Multiplication Rule.

Example: There are 4 distinct cleaning jobs: wash the windows, vacuum, dust, wash the kitchen floor. In how many ways could the jobs be assigned to 4 people from a set of 6?

Solution: Order matters. A person is assigned a particular job.

Non-Example: See the next point.

- ***If order doesn't matter and the end result is a clump or group, then you are counting combinations.***

Example: How many ways can one choose 5 people from a set of 9 to help with some chores?

Solution: Order doesn't matter. We are merely choosing the lucky folks who will help us get some work done. We have NOT assigned particular tasks.

This page titled [2.1: Equally Likely Outcomes and Counting Techniques \(Combinatorics\)](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

2.2: Conditional Probability and Bayes' Rule

In many situations, additional information about the result of a probability experiment is known (or at least assumed to be known) and given that information the probability of some other event is desired. For this scenario, we compute what is referred to as *conditional probability*.

Definition 2.2.1

For events A and B , with $P(B) > 0$, the **conditional probability** of A given B , denoted $P(A | B)$, is given by

$$P(A | B) = \frac{P(A \cap B)}{P(B)}.$$

In computing a conditional probability we *assume* that we know the outcome of the experiment is in event B and then, given that additional information, we calculate the probability that the outcome is also in event A . This is useful in practice given that partial information about the outcome of an experiment is often known, as the next example demonstrates.

Example 2.2.1

Continuing in the context of [Example 1.2.1](#), where we considered tossing a fair coin twice, define D to be the event that at least one tails is recorded:

$$D = \{ht, th, tt\}$$

Let's calculate the conditional probability of A given D , i.e., the probability that at least one heads is recorded (event A) assuming that at least one tails is recorded (event D). Recalling that outcomes in this sample space are equally likely, we apply the definition of conditional probability ([Definition 2.2.1](#)) and find

$$P(A | D) = \frac{P(A \cap D)}{P(D)} = \frac{P(\{ht, th\})}{P(\{ht, th, tt\})} = \frac{(2/4)}{(3/4)} = \frac{2}{3} \approx 0.67.$$

Note that in [Example 1.2.1](#) we found the **un**-conditional probability of A to be $P(A) = 0.75$. So, knowing that at least one tails was recorded, i.e., assuming event D occurred, the conditional probability of A given D decreased. This is because, if event D occurs, then the outcome hh in A cannot occur, thereby decreasing the chances that event A occurs.

Exercise 2.2.1

Suppose we randomly draw a card from a [standard deck of 52 playing cards](#).

- If we know that the card is a King, what is the probability that the card is a club?
- If we instead know that the card is black, what is the probability that the card is a club?

Answer

In order to compute the necessary probabilities, first note that the sample space is given by the set of cards in a standard deck of playing cards. So the number of outcomes in the sample space is 52. Next, note that the outcomes are equally likely, since we are *randomly* drawing the card from the deck.

For part (a), we are looking for the conditional probability that the randomly selected card is club, given that it is a King. If we let C denote the event that the card is a club and K the event that it is a King, then we are looking to compute

$$P(C | K) = \frac{P(C \cap K)}{P(K)}. \quad (2.2.1)$$

To compute these probabilities, we count the number of outcomes in the following events:

of outcomes in C = # of clubs in standard deck = 13

of outcomes in K = # of Kings in standard deck = 4

of outcomes in $C \cap K$ = # of King of clubs in standard deck = 1

The probabilities in Equation 2.2.1 are then given by dividing the counts of outcomes in each event by the total number of outcomes in the sample space (by the boxed Equation 2.1.8 in Section 2.1):

$$P(C | K) = \frac{P(C \cap K)}{P(K)} = \frac{(1/52)}{(4/52)} = \frac{1}{4} = 0.25.$$

For part (b), we are looking for the conditional probability that the randomly selected card is club, given that it is instead black. If we let B denote the event that the card is black, then we are looking to compute

$$P(C | B) = \frac{P(C \cap B)}{P(B)}. \quad (2.2.2)$$

To compute these probabilities, we count the number of outcomes in the following events:

of outcomes in B = # of black cards in standard deck = 26

of outcomes in $C \cap B$ = # of black clubs in standard deck = 13

The probabilities in Equation 2.2.2 are then given by dividing the counts of outcomes in each event by the total number of outcomes in the sample space:

$$P(C | B) = \frac{P(C \cap B)}{P(B)} = \frac{(13/52)}{(26/52)} = \frac{13}{26} = 0.5.$$

Remark: Exercise 2.2.1 demonstrates the following fact. For sample spaces with equally likely outcomes, conditional probabilities are calculated using

$$P(A | B) = \frac{\text{number of outcomes in } A \cap B}{\text{number of outcomes in } B}. \quad (2.2.3)$$

In other words, if we know that the outcome of the probability experiment is in the event B , then we restrict our focus to the outcomes in that event that are also in A . We can think of this as event B taking the place of the sample space, since we know the outcome must lie in that event.

Properties of Conditional Probability

As with unconditional probability, we also have some useful properties for conditional probabilities. The first property below, referred to as the *Multiplication Law*, is simply a rearrangement of the probabilities used to define conditional probability. The Multiplication Law provides a way for computing the probability of an intersection of events when the conditional probabilities are known.

Multiplication Law

$$P(A \cap B) = P(A | B)P(B) = P(B | A)P(A)$$

The next two properties are useful when a *partition* of the sample space exists, where a partition is a way of dividing up the outcomes in the sample space into non-overlapping sets. A partition is formally defined in the *Law of Total Probability* below. In many cases, when a partition exists, it is easy to compute the conditional probability of an event in the sample space given an event in the partition. The Law of Total Probability then provides a way of using those conditional probabilities of an event, given the partition to compute the unconditional probability of the event. Following the Law of Total Probability, we state *Bayes' Rule*, which is really just an application of the Multiplication Law. Bayes' Rule is used to calculate what are informally referred to as "reverse conditional probabilities", which are the conditional probabilities of an event in a partition of the sample space, given any other event.

Law of Total Probability

Suppose events B_1, B_2, \dots, B_n , satisfy the following:

1. $S = B_1 \cup B_2 \cup \dots \cup B_n$
2. $B_i \cap B_j = \emptyset$, for every $i \neq j$
3. $P(B_i) > 0$, for $i = 1, \dots, n$

We say that the events B_1, B_2, \dots, B_n , **partition** the sample space S . Then for any event A , we can write

$$P(A) = P(A | B_1)P(B_1) + \dots + P(A | B_n)P(B_n).$$

Bayes' Rule

Let B_1, B_2, \dots, B_n , partition the sample space S and let A be an event with $P(A) > 0$. Then, for $j = 1, \dots, n$, we have

$$P(B_j | A) = \frac{P(A | B_j)P(B_j)}{P(A)}.$$

A common application of the Law of Total Probability and Bayes' Rule is in the context of medical diagnostic testing.

Example 2.2.2

Consider a test that can diagnose kidney cancer. The **test correctly detects when a patient has cancer 90% of the time**. Also, **if a person does not have cancer, the test correctly indicates so 99.9% of the time**. Finally, suppose it is known that **1 in every 10,000 individuals has kidney cancer**. We find the probability that a patient has kidney cancer, given that the test indicates she does.

First, note that we are finding a conditional probability. If we let A denote the event that the patient tests positive for cancer, and we let B_1 denote the event that the patient actually has cancer, then we want

$$P(B_1 | A).$$

If we let $B_2 = B_1^c$, then we have a partition of all patients (which is the sample space) given by B_1 and B_2 .

In the first paragraph of this example, we are given the following probabilities:

$$\text{test correctly detects cancer 90\% of time: } P(A | B_1) = 0.9$$

$$\text{test correctly detects no cancer 99.9\% of time: } P(A^c | B_2) = 0.999 \Rightarrow P(A | B_2) = 1 - P(A^c | B_2) = 0.001$$

$$\text{1 in every 10,000 individuals has cancer: } P(B_1) = 0.0001 \Rightarrow P(B_2) = 1 - P(B_1) = 0.9999$$



Since we have a partition of the sample space, we apply the Law of Total Probability to find $P(A)$:

$$P(A) = P(A | B_1)P(B_1) + P(A | B_2)P(B_2) = (0.9)(0.0001) + (0.001)(0.9999) = 0.0010899$$

Next, we apply Bayes' Rule to find the desired conditional probability:

$$P(B_1 | A) = \frac{P(A | B_1)P(B_1)}{P(A)} = \frac{(0.9)(0.0001)}{0.0010899} \approx 0.08$$

This implies that only about 8% of patients that test positive under this particular test actually have kidney cancer, which is not very good.



This page titled [2.2: Conditional Probability and Bayes' Rule](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

2.3: Independent Events

In this section we consider a property of events that relates to conditional probability, namely *independence*. First, we define what it means for a pair of events to be independent, and then we consider collections of more than two events.

Independence for Pairs of Events

The following definition provides an intuitive definition of the concept of independence for two events, and then we look at an example that provides a computational way for determining when events are independent.

Definition 2.3.1

Events A and B are **independent** if knowing that one occurs does not affect the probability that the other occurs, i.e.,

$$P(A | B) = P(A) \quad \text{and} \quad P(B | A) = P(B). \quad (2.3.1)$$

Using the definition of conditional probability ([Definition 2.2.1](#)), we can derive an alternate way to the Equations [2.3.1](#) for determining when two events are independent, as the following example demonstrates.

Example 2.3.1

Suppose that events A and B are independent. We rewrite Equations [2.3.1](#) using the definition of conditional probability:

$$P(A | B) = P(A) \quad \Rightarrow \quad \frac{P(A \cap B)}{P(B)} = P(A)$$

and

$$P(B | A) = P(B) \quad \Rightarrow \quad \frac{P(A \cap B)}{P(A)} = P(B)$$

In each of the expressions on the right-hand side above we isolate $P(A \cap B)$:

$$\frac{P(A \cap B)}{P(B)} = P(A) \quad \Rightarrow \quad P(A \cap B) = P(A)P(B)$$

and

$$\frac{P(A \cap B)}{P(A)} = P(B) \quad \Rightarrow \quad P(A \cap B) = P(A)P(B)$$

Both expressions result in $P(A \cap B) = P(A)P(B)$. Thus, we have shown that if events A and B are independent, then the probability of their intersection is equal to the product of their individual probabilities. We state this fact in the next definition.

Definition 2.3.2

Events A and B are **independent** if

$$P(A \cap B) = P(A)P(B).$$

Generally speaking, Definition 2.3.2 tends to be an easier condition than Definition 2.3.1 to verify when checking whether two events are independent.

Example 2.3.2

Consider the context of [Exercise 2.2.1](#), where we randomly draw a card from a standard deck of 52 and C denotes the event of drawing a club, K the event of drawing a King, and B the event of drawing a black card.

Are C and K independent events? Recall that $P(C \cap K) = 1/52$, and note that $P(C) = 13/52$ and $P(K) = 4/52$. Thus, we have

$$P(C \cap K) = \frac{1}{52} = P(C)P(K) = \frac{13}{52} \times \frac{4}{52},$$

indicating that C and K **are independent**.

Are C and B independent events? Recall that $P(C \cap B) = 13/52$, and note that $P(B) = 26/52$. Thus, we have

$$P(C \cap B) = \frac{13}{52} \neq P(C)P(B) = \frac{13}{52} \times \frac{26}{52},$$

indicating that C and B **are not independent**.

Let's think about the results of this example intuitively. To say that C and K are independent means that knowing that one of the events occurs does not affect the probability of the other event occurring. In other words, knowing that the card drawn is a King does not influence the probability of the card being a club. The proportion of clubs in the entire deck of 52 is the same as the proportion of clubs in just the collection of Kings: $1/4$. On the other hand, C and B are not independent (AKA *dependent*) because knowing that the card drawn is club indicates that the card *must be black*, i.e., the probability that the card is black is 1. Alternately, knowing that the card drawn is black increases the probability that the card is a club, since the proportion of clubs in the entire deck is $1/4$, but the proportion of clubs in the collection of black cards is $1/2$.

Independence for 3 or More Events

For collections of 3 or more events, there are two different types of independence.

Definition 2.3.3

Let A_1, A_2, \dots, A_k , where $k \geq 3$, be a collection of events.

1. The events are **pairwise independent** if every pair of events in the collection is independent.
2. The events are **mutually independent** if every sub-collection of events, say $A_{i_1}, A_{i_2}, \dots, A_{i_n}$, satisfy the following:

$$P(A_{i_1} \cap A_{i_2} \cap \dots \cap A_{i_n}) = P(A_{i_1}) \times P(A_{i_2}) \times \dots \times P(A_{i_n})$$

Mutually independent is a stronger type of independence, since it *implies* pairwise independent. But pairwise independence does NOT imply mutual independence, as the following example will demonstrate.

Example 2.3.3

Consider again the context of [Example 1.1.1](#), i.e., tossing a fair coin twice, and define the following events:

A = first toss is heads

B = second toss is heads

C = exactly one head is recorded

We show that this collection of events - A, B, C - is pairwise independent, but NOT mutually independent. First, we note that the individual probabilities of each event are 0.5:

$$P(A) = P(\{hh, ht\}) = 0.5$$

$$P(B) = P(\{hh, th\}) = 0.5$$

$$P(C) = P(\{ht, th\}) = 0.5$$

Next, we look at the probabilities of all pairwise intersections to establish pairwise independence:

$$P(A \cap B) = P(hh) = 0.25 = P(A)P(B)$$

$$P(A \cap C) = P(ht) = 0.25 = P(A)P(C)$$

$$P(B \cap C) = P(th) = 0.25 = P(B)P(C)$$

However, note that the three events do not have any outcomes in common, i.e., $A \cap B \cap C = \emptyset$. Thus, we have

$$P(A \cap B \cap C) = 0 \neq P(A)P(B)P(C),$$

and so the events are not mutually independent.

This page titled [2.3: Independent Events](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

CHAPTER OVERVIEW

3: Discrete Random Variables

3.1: Introduction to Random Variables

3.2: Probability Mass Functions (PMFs) and Cumulative Distribution Functions (CDFs) for Discrete Random Variables

3.3: Bernoulli and Binomial Distributions

3.4: Hypergeometric, Geometric, and Negative Binomial Distributions

3.5: Poisson Distribution

3.6: Expected Value of Discrete Random Variables

3.7: Variance of Discrete Random Variables

3.8: Moment-Generating Functions (MGFs) for Discrete Random Variables

This page titled [3: Discrete Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.1: Introduction to Random Variables

Now that we have formally defined probability and the underlying structure, we add another layer: *random variables*. Random variables allow characterization of outcomes, so that we do not need to focus on each outcome specifically. We begin with the formal definition.

Definition 3.1.1

A **random variable** is a function from a sample space S to the real numbers \mathbb{R} . We denote random variables with capital letters, e.g.,

$$X : S \rightarrow \mathbb{R}.$$

Informally, a random variable assigns numbers to outcomes in the sample space. So, instead of focusing on the outcomes themselves, we highlight a specific characteristic of the outcomes.

Example 3.1.1

Consider again the context of [Example 1.1.1](#), where we recorded the sequence of heads and tails in two tosses of a fair coin. The sample space for this random experiment is given by

$$S = \{hh, ht, th, tt\}.$$

Suppose we are only interested in tosses that result in heads. We can define a random variable X that tracks the number of heads obtained in an outcome. So, if outcome hh is obtained, then X will equal 2. Formally, we denote this as follows:

$$\begin{aligned} X : S &\rightarrow \mathbb{R} \\ s &\mapsto \text{number of } h\text{'s in } s \end{aligned}$$

Since there are only four outcomes in S , we can list the value of X for each outcome individually:

inputs: S	function: X	outputs: \mathbb{R}
hh	\mapsto	2
th	\mapsto	1
ht	\mapsto	1
tt	\mapsto	0

We can also write the above as follows:

$$X(hh) = 2, \quad X(ht) = X(th) = 1, \quad X(tt) = 0.$$

The advantage to defining the random variable X in this context is that the two outcomes ht and th are both assigned a value of 1, meaning we are not focused on the actual sequence of heads and tails that resulted in obtaining one heads.

In [Example 3.1.1](#), note that the random variable we defined only equals one of three possible values: 0, 1, 2. This is an example of what we call a *discrete* random variable. We will also encounter another type of random variable: *continuous*. The next definitions make precise what we mean by these two types.

Definition 3.1.2

A **discrete random variable** is a random variable that has only a finite or countably infinite (think integers or whole numbers) number of possible values.

Definition 3.1.3

A **continuous random variable** is a random variable with infinitely many possible values (think an interval of real numbers, e.g., $[0, 1]$).

In this chapter, we take a closer look at discrete random variables, then in [Chapter 4](#) we consider continuous random variables.

This page titled [3.1: Introduction to Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.2: Probability Mass Functions (PMFs) and Cumulative Distribution Functions (CDFs) for Discrete Random Variables

Since random variables simply assign values to outcomes in a sample space and we have defined probability measures on sample spaces, we can also talk about probabilities for random variables. Specifically, we can compute the probability that a discrete random variable equals a specific value (*probability mass function*) and the probability that a random variable is less than or equal to a specific value (*cumulative distribution function*).

Probability Mass Functions (PMFs)

In the following example, we compute the probability that a discrete random variable equals a specific value.

Example 3.2.1

Continuing in the context of [Example 3.1.1](#), we compute the *probability* that the random variable X equals 1. There are two outcomes that lead to X taking the value 1, namely ht and th . So, the probability that $X = 1$ is given by the probability of the event ht, th , which is 0.5:

$$P(X = 1) = P(\{ht, th\}) = \frac{\# \text{ outcomes in } \{ht, th\}}{\# \text{ outcomes in } S} = \frac{2}{4} = 0.5$$

In [Example 3.2.1](#), the probability that the random variable X equals 1, $P(X = 1)$, is referred to as the *probability mass function* of X evaluated at 1. In other words, the specific value 1 of the random variable X is associated with the probability that X equals that value, which we found to be 0.5. The process of assigning probabilities to specific values of a discrete random variable is what the probability mass function is and the following definition formalizes this.

Definition 3.2.1

The **probability mass function (pmf)** (or **frequency function**) of a discrete random variable X assigns probabilities to the possible values of the random variable. More specifically, if x_1, x_2, \dots denote the possible values of a random variable X , then the probability mass function is denoted as p and we write

$$p(x_i) = P(X = x_i) = P(\underbrace{\{s \in S \mid X(s) = x_i\}}_{\text{set of outcomes resulting in } X=x_i}). \quad (3.2.1)$$

Note that, in [Equation 3.2.1](#), $p(x_i)$ is *shorthand* for $P(X = x_i)$, which represents the probability of the *event* that the random variable X equals x_i .

As we can see in [Definition 3.2.1](#), the probability mass function of a random variable X depends on the probability measure of the underlying sample space S . Thus, pmf's inherit some properties from the axioms of probability ([Definition 1.2.1](#)). In fact, in order for a function to be a *valid* pmf it must satisfy the following properties.

Properties of Probability Mass Functions

Let X be a discrete random variable with possible values denoted $x_1, x_2, \dots, x_i, \dots$. The probability mass function of X , denoted p , must satisfy the following:

1. $\sum_{x_i} p(x_i) = p(x_1) + p(x_2) + \dots = 1$
2. $p(x_i) \geq 0$, for all x_i

Furthermore, if A is a subset of the possible values of X , then the probability that X takes a value in A is given by

$$P(X \in A) = \sum_{x_i \in A} p(x_i). \quad (3.2.2)$$

Note that the first property of pmf's stated above follows from the first axiom of probability, namely that the probability of the sample space equals 1: $P(S) = 1$. The second property of pmf's follows from the second axiom of probability, which states that all probabilities are non-negative.

We now apply the formal definition of a pmf and verify the properties in a specific context.

Example 3.2.2

Returning to [Example 3.2.1](#), now using the notation of [Definition 3.2.1](#), we found that the pmf for X at 1 is given by

$$p(1) = P(X = 1) = P(\{ht, th\}) = 0.5.$$

Similarly, we find the pmf for X at the other possible values of the random variable:

$$\begin{aligned} p(0) &= P(X = 0) = P(\{tt\}) = 0.25 \\ p(2) &= P(X = 2) = P(\{hh\}) = 0.25 \end{aligned}$$

Note that all the values of p are positive (second property of pmf's) and $p(0) + p(1) + p(2) = 1$ (first property of pmf's). Also, we can demonstrate the third property of pmf's (Equation 3.2.2) by computing the probability that there is at least one heads, i.e., $X \geq 1$, which we could represent by setting $A = \{1, 2\}$ so that we want the probability that X takes a value in A :

$$P(X \geq 1) = P(X \in A) = \sum_{x_i \in A} p(x_i) = p(1) + p(2) = 0.5 + 0.25 = 0.75$$

We can represent probability mass functions *numerically* with a table, *graphically* with a **histogram**, or *analytically* with a formula. The following example demonstrates the numerical and graphical representations. In the next three sections, we will see examples of pmf's defined analytically with a formula.

Example 3.2.3

We represent the pmf we found in [Example 3.2.2](#) in two ways below, numerically with a table on the left and graphically with a histogram on the right.

x	$p(x)$
0	0.25
1	0.5
2	0.25

Table 1: Frequency function of X




Figure 1: Histogram of X

In the histogram in Figure 1, note that we represent probabilities as *areas of rectangles*. More specifically, each rectangle in the histogram has width 1 and height equal to the probability of the value of the random variable X that the rectangle is centered over. For example, the leftmost rectangle in the histogram is centered at 0 and has height equal to $p(0) = 0.25$, which is also the area of the rectangle since the width is equal to 1. In this way, histograms provides a visualization of the *distribution* of the probabilities assigned to the possible values of the random variable X . This helps to explain where the common terminology of "probability distribution" comes from when talking about random variables.

Cumulative Distribution Functions (CDFs)

There is one more important function related to random variables that we define next. This function is again related to the probabilities of the random variable equalling specific values. It provides a shortcut for calculating many probabilities at once.

Definition 3.2.2

The **cumulative distribution function (cdf)** of a random variable X is a function on the real numbers that is denoted as F and is given by

$$F(x) = P(X \leq x), \quad \text{for any } x \in \mathbb{R}. \quad (3.2.3)$$

Before looking at an example of a cdf, we note a few things about the definition.

First of all, note that we did not specify the random variable X to be discrete. CDFs are also defined for continuous random variables (see [Chapter 4](#)) in exactly the same way.

Second, the cdf of a random variable is defined for *all* real numbers, unlike the pmf of a discrete random variable, which we only define for the possible values of the random variable. Implicit in the definition of a pmf is the assumption that it equals 0 for all real numbers that are not possible values of the discrete random variable, which should make sense since the random variable will never equal that value. However, cdf's, for *both* discrete and continuous random variables, are defined for all real numbers. In looking more closely at Equation 3.2.3, we see that a cdf F considers an upper bound, $x \in \mathbb{R}$, on the random variable X , and assigns that value x to the probability that the random variable X is less than or equal to that upper bound x . This type of probability is referred to as a *cumulative probability*, since it could be thought of as the probability accumulated by the random variable up to the specified upper bound. With this interpretation, we can represent Equation 3.2.3 as follows:

$$F : \underbrace{\mathbb{R}}_{\text{upper bounds on RV } X} \longrightarrow \underbrace{\mathbb{R}}_{\text{cumulative probabilities}} \quad (3.2.4)$$

In the case that X is a discrete random variable, with possible values denoted $x_1, x_2, \dots, x_i, \dots$, the cdf of X can be calculated using the third property of pmf's (Equation 3.2.2), since, for a fixed $x \in \mathbb{R}$, if we let the set A contain the possible values of X that are less than or equal to x , i.e., $A = \{x_i \mid x_i \leq x\}$, then the cdf of X evaluated at x is given by

$$F(x) = P(X \leq x) = P(X \in A) = \sum_{x_i \leq x} p(x_i).$$

Example 3.2.4

Continuing with [Examples 3.2.2](#) and [3.2.3](#), we find the cdf for X . First, we find $F(x)$ for the possible values of the random variable, $x = 0, 1, 2$:

$$\begin{aligned} F(0) &= P(X \leq 0) = P(X = 0) = 0.25 \\ F(1) &= P(X \leq 1) = P(X = 0 \text{ or } 1) = p(0) + p(1) = 0.75 \\ F(2) &= P(X \leq 2) = P(X = 0 \text{ or } 1 \text{ or } 2) = p(0) + p(1) + p(2) = 1 \end{aligned}$$

Now, if $x < 0$, then the cdf $F(x) = 0$, since the random variable X will never be negative.

If $0 < x < 1$, then the cdf $F(x) = 0.25$, since the only value of the random variable X that is less than or equal to such a value x is 0. For example, consider $x = 0.5$. The probability that X is less than or equal to 0.5 is the same as the probability that $X = 0$, since 0 is the only possible value of X less than 0.5:

$$F(0.5) = P(X \leq 0.5) = P(X = 0) = 0.25.$$

Similarly, we have the following:

$$\begin{aligned} F(x) &= F(1) = 0.75, & \text{for } 1 < x < 2 \\ F(x) &= F(2) = 1, & \text{for } x > 2 \end{aligned}$$

Exercise 3.2.1

For this random variable X , compute the following values of the cdf:

- $F(-3)$
- $F(0.1)$
- $F(0.9)$
- $F(1.4)$
- $F(2.3)$
- $F(18)$

Answer

- $F(-3) = P(X \leq -3) = 0$

- b. $F(0.1) = P(X \leq 0.1) = P(X = 0) = 0.25$
- c. $F(0.9) = P(X \leq 0.9) = P(X = 0) = 0.25$
- d. $F(1.4) = P(X \leq 1.4) = \sum_{x_i \leq 1.4} p(x_i) = p(0) + p(1) = 0.25 + 0.5 = 0.75$
- e. $F(2.3) = P(X \leq 2.3) = \sum_{x_i \leq 2.3} p(x_i) = p(0) + p(1) + p(2) = 0.25 + 0.5 + 0.25 = 1$
- f. $F(18) = P(X \leq 18) = P(X \leq 2) = 1$

To summarize Example 3.2.4, we write the cdf F as a *piecewise function* and Figure 2 gives its graph:

$$F(x) = \begin{cases} 0, & \text{for } x < 0 \\ 0.25 & \text{for } 0 \leq x < 1 \\ 0.75 & \text{for } 1 \leq x < 2 \\ 1 & \text{for } x \geq 2. \end{cases}$$




Figure 2: Graph of cdf in Example 3.2.4

Note that the cdf we found in Example 3.2.4 is a "step function", since its graph resembles a series of steps. This is the case for *all* discrete random variables. Additionally, the value of the cdf for a discrete random variable will always "jump" at the possible values of the random variable, and the size of the "jump" is given by the value of the pmf at that possible value of the random variable. For example, the graph in Figure 2 "jumps" from 0.25 to 0.75 at $x = 1$, so the size of the "jump" is $0.75 - 0.25 = 0.5$ and note that $p(1) = P(X = 1) = 0.5$. The pmf for any discrete random variable can be obtained from the cdf in this manner.

We end this section with a statement of the properties of cdf's. The reader is encouraged to verify these properties hold for the cdf derived in Example 3.2.4 and to provide an intuitive explanation (or formal explanation using the axioms of probability and the properties of pmf's) for why these properties hold for cdf's in general.

Properties of Cumulative Distribution Functions

Let X be a random variable with cdf F . Then F satisfies the following:

1. F is non-decreasing, i.e., F may be constant, but otherwise it is increasing.
2. $\lim_{x \rightarrow -\infty} F(x) = 0$ and $\lim_{x \rightarrow \infty} F(x) = 1$

This page titled [3.2: Probability Mass Functions \(PMFs\) and Cumulative Distribution Functions \(CDFs\) for Discrete Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.3: Bernoulli and Binomial Distributions

In this section and the next two, we introduce *families* of common discrete probability distributions, i.e., probability distributions for discrete random variables. We refer to these as "families" of distributions because in each case we will define a probability mass function by specifying an explicit formula, and that formula will incorporate a constant (or set of constants) that are referred to as **parameters**. By specifying values for the parameter(s) in the pmf, we define a specific probability distribution for a specific random variable. For each family of distributions introduced, we will list a set of defining characteristics that will help determine when to use a certain distribution in a given context.

Bernoulli Distribution

Consider the following example.

Example 3.3.1

Let A be an event in a sample space S . Suppose we are only interested in whether or not the outcome of the underlying probability experiment is in the specified event A . To track this we can define an **indicator random variable**, denoted I_A , given by

$$I_A(s) = \begin{cases} 1, & \text{if } s \in A, \\ 0, & \text{if } s \in A^c. \end{cases}$$

In other words, the random variable I_A will equal 1 if the resulting outcome is in event A , and I_A equals 0 if the outcome is not in A . Thus, I_A is a discrete random variable. We can state the probability mass function of I_A in terms of the probability that the resulting outcome is in event A , i.e., the probability that event A occurs, $P(A)$:

$$\begin{aligned} p(0) &= P(I_A = 0) = P(A^c) = 1 - P(A) \\ p(1) &= P(I_A = 1) = P(A) \end{aligned}$$

In Example 3.3.1, the random variable I_A is a *Bernoulli random variable* because its pmf has the form of the *Bernoulli probability distribution*, which we define next.

Definition 3.3.1

A random variable X has a **Bernoulli distribution** with parameter p , where $0 \leq p \leq 1$, if it has only two possible values, typically denoted 0 and 1. The probability mass function (pmf) of X is given by

$$\begin{aligned} p(0) &= P(X = 0) = 1 - p, \\ p(1) &= P(X = 1) = p. \end{aligned}$$

The cumulative distribution function (cdf) of X is given by

$$F(x) = \begin{cases} 0, & x < 0 \\ 1 - p, & 0 \leq x < 1, \\ 1, & x \geq 1. \end{cases} \quad (3.3.1)$$

In Definition 3.3.1, note that the defining characteristic of the Bernoulli distribution is that it models random variables that have only two possible values. As noted in the definition, the two possible values of a Bernoulli random variable are usually 0 and 1. In the typical application of the Bernoulli distribution, a value of 1 indicates a "success" and a value of 0 indicates a "failure", where "success" refers that the event or outcome of interest. The parameter p in the Bernoulli distribution is given by the probability of a "success". In Example 3.3.1, we were interested in tracking whether or not event A occurred, and so that is what a "success" would be, which occurs with probability given by the probability of A . Thus, the value of the parameter p for the Bernoulli distribution in Example 3.3.1 is given by $p = P(A)$.

Exercise 3.3.1

Derive the general formula for the cdf of the Bernoulli distribution given in Equation 3.3.1.

Hint

First find $F(0)$ and $F(1)$.

Answer

Recall that the only two values of a Bernoulli random variable X are 0 and 1. So, first, we find the cdf at those two values:

$$\begin{aligned} F(0) &= P(X \leq 0) = P(X = 0) = p(0) = 1 - p \\ F(1) &= P(X \leq 1) = P(X = 0 \text{ or } 1) = p(0) + p(1) = (1 - p) + p = 1 \end{aligned}$$

Now for the other values, a Bernoulli random variable will never be negative, so $F(x) = 0$, for $x < 0$. Also, a Bernoulli random variable will always be less than or equal to 1, so $F(x) = 1$, for $x \geq 1$. Lastly, if x is in between 0 and 1, then the cdf is given by

$$F(x) = P(X \leq x) = P(X = 0) = p(0) = 1 - p, \text{ for } 0 \leq x < 1.$$

Binomial Distribution

To introduce the next family of distributions, we use our continuing example of tossing a coin, adding another toss.

Example 3.3.2

Suppose we toss a coin three times and record the sequence of heads (h) and tails (t). Supposing that the coin is fair, each toss results in heads with probability 0.5, and tails with the same probability of 0.5. Since the three tosses are mutually independent, the probability assigned to any outcome is 0.5^3 . More specifically, consider the outcome hth . We could write the probability of this outcome as $(0.5)^2(0.5)^1$ to emphasize the fact that two heads and one tails occurred. Note that there are two other outcomes with two heads and one tails: hht and thh . Recall from [Example 2.1.2](#) in [Section 2.1](#), that we can count the number of outcomes with two heads and one tails by counting the number of ways to select positions for the two heads to occur in a sequence of three tosses, which is given by $\binom{3}{2}$. In general, note that $\binom{3}{x}$ counts the number of possible sequences with exactly x heads, for $x = 0, 1, 2, 3$.

We generalize the above by defining the discrete random variable X to be the number of heads in an outcome. The possible values of X are $x = 0, 1, 2, 3$. Using the above facts, the pmf of X is given as follows:

$$\begin{aligned} p(0) &= P(X = 0) = P(\{ttt\}) = \frac{1}{8} = \binom{3}{0}(0.5)^0(0.5)^3 \\ p(1) &= P(X = 1) = P(\{htt, tht, tth\}) = \frac{3}{8} = \binom{3}{1}(0.5)^1(0.5)^2 \\ p(2) &= P(X = 2) = P(\{hht, hth, thh\}) = \frac{3}{8} = \binom{3}{2}(0.5)^2(0.5)^1 \\ p(3) &= P(X = 3) = P(\{hhh\}) = \frac{1}{8} = \binom{3}{3}(0.5)^3(0.5)^0 \end{aligned} \tag{3.3.2}$$

In the above, the fractions in orange are found by calculating the probabilities directly using equally likely outcomes (note that the sample space S has 8 outcomes, see [Example 2.1.1](#)). In each line, the value of x is highlighted in red so that we can see the pattern forming. For example, when $x = 2$, we see in the expression on the right-hand side of Equation 3.3.2 that "2" appears in the binomial coefficient $\binom{3}{2}$, which gives the number of outcomes resulting in the random variable equaling 2, and "2" also appears in the exponent on the first 0.5, which gives the probability of two heads occurring.

The pattern exhibited by the random variable X in Example 3.3.2 is referred to as the *binomial distribution*, which we formalize in the next definition.

Definition 3.3.2

Suppose that n independent trials of the same probability experiment are performed, where each trial results in either a "success" (with probability p), or a "failure" (with probability $1 - p$). If the random variable X denotes the total number of successes in the n trials, then X has a **binomial distribution** with parameters n and p , which we write $X \sim \text{binomial}(n, p)$. The probability mass function of X is given by

$$p(x) = P(X = x) = \binom{n}{x} p^x (1 - p)^{n-x}, \quad \text{for } x = 0, 1, \dots, n. \tag{3.3.3}$$

In Example 3.3.2, the independent trials are the three tosses of the coin, so in this case we have parameter $n = 3$. Furthermore, we were interested in counting the number of heads occurring in the three tosses, so a "success" is getting a heads on a toss, which occurs with probability 0.5 and so parameter $p = 0.5$. Thus, the random variable X in this example has a binomial(3, 0.5) distribution and applying the formula for the binomial pmf given in Equation 3.3.3 when $x = 2$ we get the same expression on the right-hand side of Equation 3.3.2:

$$p(x) = \binom{n}{x} p^x (1-p)^{n-x} \Rightarrow p(2) = \binom{3}{2} 0.5^2 (1-0.5)^{3-2} = \binom{3}{2} 0.5^2 0.5^1$$

In general, we can connect binomial random variables to Bernoulli random variables. If X is a binomial random variable, with parameters n and p , then it can be written as the sum of n independent *Bernoulli* random variables, X_1, \dots, X_n . (Note: We will formally define *independence* for random variables later, in Chapter 5.) Specifically, if we define the random variable X_i , for $i = 1, \dots, n$, to be 1 when the i^{th} trial is a "success", and 0 when it is a "failure", then the sum

$$X = X_1 + X_2 + \dots + X_n$$

gives the total number of success in n trials. This connection between the binomial and Bernoulli distribution will be useful in a later section.

One of the main applications of the binomial distribution is to model population characteristics as in the following example.

Example 3.3.3

Consider a group of 100 voters. If p denotes the probability that a voter will vote for a specific candidate, and we let random variable X denote the number of voters in the group that will vote for that candidate, then X follows a binomial distribution with parameters $n = 100$ and p .

This page titled 3.3: Bernoulli and Binomial Distributions is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.4: Hypergeometric, Geometric, and Negative Binomial Distributions

In this section, we consider three more families of discrete probability distributions. There are some similarities between the three, which can make them hard to distinguish at times. So throughout this section we will compare the three to each other and the binomial distribution, and point out their differences.

Hypergeometric Distribution

Consider the following example.

Example 3.4.1

An urn contains a total of N balls, where some number m of the balls are **orange** and the remaining $N - m$ are **grey**. Suppose we draw n balls from the urn without replacement, meaning once we select a ball we do not place it back in the urn before drawing out the next one. Then some of the balls in our selection may be **orange** and some may be **grey**. We can define the discrete random variable X to give the number of **orange** balls in our selection. The probability distribution of X is referred to as the *hypergeometric distribution*, which we define next.

Definition 3.4.1

Suppose in a collection of N objects, m are of type 1 and $N - m$ are of another type 2. Furthermore, suppose that n objects are randomly selected from the collection without replacement. Define the discrete random variable X to give the number of selected objects that are of type 1. Then X has a **hypergeometric distribution** with parameters N, m, n . The probability mass function of X is given by

$$\begin{aligned} p(x) &= P(X = x) = P(x \text{ type 1 objects \& } n - x \text{ type 2}) \\ &= \frac{(\# \text{ of ways to select } x \text{ type 1 objects from } m) \times (\# \text{ of ways to select } n - x \text{ type 2 objects from } N - m)}{\text{total \# of ways to select } n \text{ objects of any type from } N} \\ &= \frac{\binom{m}{x} \binom{N - m}{n - x}}{\binom{N}{n}} \end{aligned}$$

In some sense, the hypergeometric distribution is similar to the binomial, except that the method of sampling is *crucially* different. In each case, we are interested in the number of times a specific outcome occurs in a set number of repeated trials, where we could consider each selection of an object in the hypergeometric case as a trial. In the binomial case we are interested in the number of "successes" in the trials, and in the hypergeometric case we are interested in the number of a certain type of object being selected, which could be considered a "success". However, the trials in a binomial distribution are independent, while the trials in a hypergeometric distribution are not because the objects are selected without replacement. If, in [Example 3.4.1](#), the balls were drawn *with* replacement, then each draw would be an independent Bernoulli trial and the distribution of X would be binomial, since the same number of balls in the urn would be the same each time another ball is drawn. However, when the balls are drawn *without* replacement, each draw is not independent, since the number of balls in the urn decreases after each draw as well as the number of balls of a given type.

Exercise 3.4.1

Suppose your friend has 10 cookies, 3 of which are chocolate chip. Your friend randomly divides the cookies equally between herself and you. What is the probability that you get all the chocolate chip cookies?

Answer

Let random variable X = number of chocolate chip cookies you get. Then X is hypergeometric with $N = 10$ total cookies, $m = 3$ chocolate chip cookies, and $n = 5$ cookies selected by your friend to give to you. We want the probability that you get all the chocolate chip cookies, i.e., $P(X = 3)$, which is

$$P(X = 3) = \frac{\binom{3}{3} \binom{7}{2}}{\binom{10}{5}} = 0.083$$

Note that X has a hypergeometric distribution and not binomial because the cookies are being selected (or divided) without replacement.

Geometric Distribution & Negative Binomial Distribution

The geometric and negative binomial distributions are related to the binomial distribution in that the underlying probability experiment is the same, i.e., independent trials with two possible outcomes. However, the random variable defined in the geometric and negative binomial case highlights a different aspect of the experiment, namely the number of trials needed to obtain a specific number of "successes". We start with the geometric distribution.

Definition 3.4.2

Suppose that a sequence of independent Bernoulli trials is performed, with $p = P(\text{"success"})$ for each trial. Define the random variable X to give the number of trial at which the first success occurs. Then X has a **geometric distribution** with parameter p . The probability mass function of X is given by

$$\begin{aligned} p(x) &= P(X = x) = P(1^{\text{st}} \text{ success on } x^{\text{th}} \text{ trial}) \\ &= P(1^{\text{st}} (x-1) \text{ trials are failures \& } x^{\text{th}} \text{ trial is success}) \\ &= (1-p)^{x-1} p, \quad \text{for } x = 1, 2, 3, \dots \end{aligned} \quad (3.4.1)$$

Exercise 3.4.2

Verify that the pmf for a geometric distribution (Equation 3.4.1) satisfies the **two properties for pmf's**, i.e.,

1. $p(x) \geq 0$, for $x = 1, 2, 3, \dots$
2. $\sum_{x=1}^{\infty} p(x) = 1$ *Hint: It's called "geometric" for a reason!*

Answer

1. Note that $0 \leq p \leq 1$, so that we also have $0 \leq (1-p) \leq 1$ and $0 \leq (1-p)^{x-1} \leq 1$, for $x = 1, 2, \dots$. Thus, it follows that $p(x) = (1-p)^{x-1} p \geq 0$.
2. Recall the formula for the sum of a geometric series:

$$\sum_{x=1}^{\infty} ar^{x-1} = \frac{a}{1-r}, \quad \text{if } |r| < 1.$$

Note that the sum of the geometric pmf is a geometric series with $a = p$ and $r = 1-p < 1$. Thus, we have

$$\sum_{x=1}^{\infty} p(x) = \sum_{x=1}^{\infty} (1-p)^{x-1} p = \frac{p}{1-(1-p)} = \frac{p}{p} = 1 \quad \checkmark$$

Example 3.4.2

Each of the following is an example of a random variable with the geometric distribution.

1. Toss a fair coin until the first heads occurs. In this case, a "success" is getting a heads ("failure" is getting tails) and so the parameter $p = P(h) = 0.5$.
2. Buy lottery tickets until getting the first win. In this case, a "success" is getting a lottery ticket that wins money, and a "failure" is not winning. The parameter p will depend on the odds of winning for a specific lottery.

3. Roll a pair of fair dice until getting the first double 1's. In this case, a "success" is getting double 1's, and a "failure" is simply not getting double 1's (so anything else). To find the parameter p , note that the underlying sample space consists of all possible rolls of a pair of fair dice, of which there are $6 \times 6 = 36$ because each die has 6 possible sides. Each of these rolls is equally likely, so

$$p = P(\text{double 1's}) = \frac{\# \text{ of ways to roll double 1's}}{36} = \frac{1}{36}.$$

The negative binomial distribution generalizes the geometric distribution by considering any number of successes.

Definition 3.4.3

Suppose that a sequence of independent Bernoulli trials is performed, with $p = P(\text{"success"})$ for each trial. Fix an integer r to be greater than or equal to 2 and define the random variable X to give the number of trial at which the r^{th} success occurs. Then X has a **negative binomial distribution** with parameters r and p . The probability mass function of X is given by

$$\begin{aligned} p(x) &= P(X = x) = P(r^{\text{th}} \text{ success is on } x^{\text{th}} \text{ trial}) \\ &= \underbrace{P(1^{\text{st}} (r-1) \text{ successes in } 1^{\text{st}} (x-1) \text{ trials})}_{\text{binomial with } n=x-1} \times P(r^{\text{th}} \text{ success on } x^{\text{th}} \text{ trial}) \\ &= \binom{x-1}{r-1} p^{r-1} (1-p)^{(x-1)-(r-1)} \times p \\ &= \binom{x-1}{r-1} p^r (1-p)^{x-r}, \quad \text{for } x = r, r+1, r+2, \dots \end{aligned}$$

Example 3.4.3

For examples of the negative binomial distribution, we can alter the geometric examples given in [Example 3.4.2](#).

1. Toss a fair coin until get 8 heads. In this case, the parameter p is still given by $p = P(h) = 0.5$, but now we also have the parameter $r = 8$, the number of desired "successes", i.e., heads.
2. Buy lottery tickets until win 5 times. In this case, the parameter p is still given by the odds of winning the lottery, but now we also have the parameter $r = 5$, the number of desired wins.
3. Roll a pair of fair dice until get 100 double 1's. In this case, the parameter p is still given by $p = P(\text{double 1's}) = \frac{1}{36}$, but now we also have the parameter $r = 100$, the number of desired "successes".

In general, note that a geometric distribution can be thought of a negative binomial distribution with parameter $r = 1$.

Note that for both the geometric and negative binomial distributions the number of possible values the random variable can take is infinite. These are still *discrete* distributions though, since we can "list" the values. In other words, the possible values are **countable**. This is in contrast to the Bernoulli, binomial, and hypergeometric distributions, where the number of possible values is finite.

We again note the distinction between the binomial distribution and the geometric and negative binomial distributions. In the binomial distribution, the number of *trials* is fixed, and we count the number of "successes". Whereas, in the geometric and negative binomial distributions, the number of "*successes*" is fixed, and we count the number of trials needed to obtain the desired number of "successes".

This page titled [3.4: Hypergeometric, Geometric, and Negative Binomial Distributions](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.5: Poisson Distribution

In this section, we consider our final family of discrete probability distributions. We begin with the definition.

Definition 3.5.1

A random variable X has a **Poisson distribution**, with parameter $\lambda > 0$, if its probability mass function is given by

$$p(x) = P(X = x) = \frac{e^{-\lambda} \lambda^x}{x!}, \quad \text{for } x = 0, 1, 2, \dots \quad (3.5.1)$$

We write $X \sim \text{Poisson}(\lambda)$.

The main application of the Poisson distribution is to count the number of times some event occurs over a fixed interval of time or space. More specifically, if the random variable X denotes the number of times the event occurs during an interval of length T , and r denotes the average rate at which the event occurs per unit interval, then X has a Poisson distribution with parameter $\lambda = rT$. Consider the following examples:

- The number of customers arriving at McDonald's between 8 a.m. and 9 a.m.
- The number of calls made to 911 in South Bend on a Saturday.
- The number of accidents at a particular intersection during the month of June.

All of the examples above count the number of times something occurs over an interval of *time*. The next example gives an example where the interval is in *space*.

Example 3.5.1

Suppose typos occur at an average rate of $r = 0.01$ per page in the Friday edition of the *New York Times*, which is 45 pages long. Let X denote the number of typos on the front page. Then X has a Poisson distribution with parameter

$$\lambda = 0.01 \times 1 = 0.01,$$

since we are considering an interval of length one page ($T = 1$). Thus, the probability that there is at least one typo on the front page is given by

$$P(X \geq 1) = P(\{X = 0\}^c) = 1 - P(X = 0) = 1 - \frac{e^{-0.01}(0.01)^0}{0!} = 1 - e^{-0.01} \approx 0.00995.$$

Now, if we let random variable Y denote the number of typos in the entire paper, then Y has a Poisson distribution with parameter

$$\lambda = 0.01 \times 45 = 0.45,$$

since we are considering an interval of $T = 45$ pages. The probability that there are less than three typos in the entire paper is

$$P(Y < 3) = P(Y = 0, 1, \text{ or } 2) = \frac{e^{-0.45}(0.45)^0}{0!} + \frac{e^{-0.45}(0.45)^1}{1!} + \frac{e^{-0.45}(0.45)^2}{2!} \approx 0.98912.$$

The Poisson distribution is similar to all previously considered families of discrete probability distributions in that it counts the number of times something happens. However, the Poisson distribution is different in that there is not an act that is being repeatedly performed. In other words, there are no set trials, but rather a set window of time or space to observe.

This page titled [3.5: Poisson Distribution](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.6: Expected Value of Discrete Random Variables

In this section, and the next, we look at various *numerical* characteristics of discrete random variables. These give us a way of classifying and comparing random variables.

Expected Value of Discrete Random Variables

We begin with the formal definition.

Definition 3.6.1

If X is a discrete random variable with possible values $x_1, x_2, \dots, x_i, \dots$, and probability mass function $p(x)$, then the **expected value** (or **mean**) of X is denoted $E[X]$ and given by

$$E[X] = \sum_i x_i \cdot p(x_i). \quad (3.6.1)$$

The expected value of X may also be denoted as μ_X or simply μ if the context is clear.

The expected value of a random variable has many interpretations. First, looking at the formula in Definition 3.6.1 for computing expected value (Equation 3.6.1), note that it is essentially a **weighted average**. Specifically, for a discrete random variable, the expected value is computed by "weighting", or multiplying, each value of the random variable, x_i , by the probability that the random variable takes that value, $p(x_i)$, and then summing over all possible values. This interpretation of the expected value as a weighted average explains why it is also referred to as the **mean** of the random variable.

The expected value of a random variable is also interpreted as the **long-run value** of the random variable. In other words, if we repeat the underlying random experiment several times and take the average of the values of the random variable corresponding to the outcomes, we would get the expected value, approximately. (Note: This interpretation of expected value is similar to the [relative frequency approximation](#) for probability discussed in [Section 1.2](#).) Again, we see that the expected value is related to an average value of the random variable. Given the interpretation of the expected value as an average, either "weighted" or "long-run", the expected value is often referred to as a **measure of center** of the random variable.

Finally, the expected value of a random variable has a graphical interpretation. The expected value gives the **center of mass** of the probability mass function, which the following example demonstrates.

Example 3.6.1

Consider again the context of [Example 1.1.1](#), where we recorded the sequence of heads and tails in two tosses of a fair coin. In [Example 3.1.1](#) we defined the discrete random variable X to denote the number of heads obtained. In [Example 3.2.2](#) we found the pmf of X . We now apply Equation 3.6.1 from Definition 3.6.1 and compute the expected value of X :

$$\begin{aligned} E[X] &= 0 \cdot p(0) + 1 \cdot p(1) + 2 \cdot p(2) \\ &= 0 \cdot (0.25) + 1 \cdot (0.5) + 2 \cdot (0.25) \\ &= 0.5 + 0.5 = 1. \end{aligned}$$

Thus, we expect that the number of heads obtained in two tosses of a fair coin will be 1 in the long-run or on average. Figure 1 demonstrates the graphical representation of the expected value as the center of mass of the probability mass function.




Figure 1: Histogram of X . The red arrow represents the center of mass, or the expected value of X

Example 3.6.2

Suppose we toss a fair coin three times and define the random variable X to be our winnings on a single play of a game where

- we win $\$x$ if the first heads is on the x^{th} toss, for $x = 1, 2, 3$,
- and we lose $\$1$ if we get no heads in all three tosses.

Then X is a discrete random variable, with possible values $x = -1, 1, 2, 3$, and pmf given by the following table:

x	$p(x) = P(X = x)$
-1	$\frac{1}{8}$
1	$\frac{1}{2}$
2	$\frac{1}{4}$
3	$\frac{1}{8}$

Applying [Definition 3.6.1](#), we find

$$\begin{aligned} E[X] &= \sum_i x_i \cdot p(x_i) \\ &= (-1) \cdot \frac{1}{8} + 1 \cdot \frac{1}{2} + 2 \cdot \frac{1}{4} + 3 \cdot \frac{1}{8} = \frac{5}{4} = 1.25. \end{aligned}$$

Thus, the expected winnings for a single play of the game is \$1.25. In other words, if we played the game multiple times, we expect the average winnings to be \$1.25.

For many of the common probability distributions, the expected value is given by a parameter of the distribution as the next exercise shows for the Poisson distribution.

Exercise 3.6.1

Suppose the discrete random variable X has a Poisson distribution with parameter λ . Show that $E[X] = \lambda$.

Hint

Recall from calculus the series expansion: $e^y = \sum_{x=1}^{\infty} \frac{y^{x-1}}{(x-1)!}$

Answer

First, recall that the pmf for Poisson distribution is $p(x) = \frac{e^{-\lambda} \lambda^x}{x!}$, for $x = 0, 1, \dots$. Then, we apply [Definition 3.6.1](#), giving us

$$\begin{aligned} E[X] &= \sum_{x=0}^{\infty} x \cdot p(x) \\ &= \sum_{x=0}^{\infty} x \cdot \frac{e^{-\lambda} \lambda^x}{x!} \quad (\text{note that the } x = 0 \text{ term is 0, so we drop it from the sum}) \\ &= \sum_{x=1}^{\infty} x \cdot \frac{e^{-\lambda} \lambda^x}{x!} = \sum_{x=1}^{\infty} x \cdot \frac{e^{-\lambda} \lambda \cdot \lambda^{x-1}}{x \cdot (x-1)!} \\ &= e^{-\lambda} \cdot \lambda \sum_{x=1}^{\infty} \frac{\lambda^{x-1}}{(x-1)!} \\ &= e^{-\lambda} \cdot \lambda \cdot e^{\lambda} \quad (\text{here we used the series expansion of } e^y) \\ &= \lambda \end{aligned}$$

as needed.

The expected value may not be exactly equal to a parameter of the probability distribution, but rather it may be a function of the parameters. The following table gives the expected value for each of the common discrete distributions we considered earlier. We will see later ([Section 3.8](#)) how to derive these results using a technique involving moment-generating functions.

Expected Values for Discrete Distributions

Distribution	Expected Value
Bernoulli(p)	p
binomial(n, p)	np
hypergeometric(N, n, m)	$\frac{nm}{N}$
geometric(p)	$\frac{1}{p}$
negative binomial(r, p)	$\frac{r}{p}$
Poisson(λ)	λ

Expected Value of Functions of Random Variables

In many applications, we may not be interested in the value of a random variable itself, but rather in a function applied to the random variable or a collection of random variables. For example, we may be interested in the value of X^2 . The following theorems, which we state without proof, demonstrate how to calculate the expected value of *functions of random variables*.

Theorem 3.6.1

Let X be a random variable and let g be a real-valued function. Define the random variable $Y = g(X)$.

If X is a discrete random variable with possible values $x_1, x_2, \dots, x_i, \dots$, and probability mass function $p(x)$, then the expected value of Y is given by

$$E[Y] = \sum_i g(x_i) \cdot p(x_i).$$

To put it simply, Theorem 3.6.1 states that to find the expected value of a function of a random variable, just apply the function to the possible values of the random variable in the definition of expected value. Before stating an important special case of Theorem 3.6.1, a word of caution regarding order of operations. Note that, in general,

$$E[g(X)] \neq g(E[X])! \quad (3.6.2)$$

For example, $E[X^2] \neq (E[X])^2$, in general. However, as the next theorem states, there are exceptions to Equation 3.6.2.

Special Case of Theorem 3.6.1

Let X be a random variable. If g is a *linear* function, i.e., $g(x) = ax + b$, then

$$E[g(X)] = E[aX + b] = aE[X] + b.$$

The above special case is referred to as the **linearity** of expected value, which implies the following properties of the expected value.

Linearity of Expected Value

Let X be a random variable, c, c_1, c_2 constants, and g, g_1, g_2 real-valued functions. Then expectation $E[\cdot]$ satisfies the following:

1. The expected value of a constant is constant:

$$E[c] = c$$

2. Constants can be factored out of expected values:

$$E[cg(X)] = cE[g(X)]$$

3. The expected value of a sum is equal to the sum of expected values:

$$E[c_1g_1(X) + c_2g_2(X)] = c_1E[g_1(X)] + c_2E[g_2(X)]$$

This page titled [3.6: Expected Value of Discrete Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.7: Variance of Discrete Random Variables

We now look at our second numerical characteristic associated to random variables.

Definition 3.7.1

The **variance** of a random variable X is given by

$$\sigma^2 = \text{Var}(X) = E[(X - \mu)^2],$$

where μ denotes the expected value of X . The **standard deviation** of X is given by

$$\sigma = \text{SD}(X) = \sqrt{\text{Var}(X)}.$$

In words, the variance of a random variable is the average of the squared deviations of the random variable from its mean (expected value). Notice that the variance of a random variable will result in a number with units squared, but the standard deviation will have the same units as the random variable. Thus, the standard deviation is easier to interpret, which is why we make a point to define it.

The variance and standard deviation give us a **measure of spread** for random variables. The standard deviation is interpreted as a measure of how "spread out" the possible values of X are with respect to the mean of X , $\mu = E[X]$.

Example 3.7.1

Consider the two random variables X_1 and X_2 , whose probability mass functions are given by the histograms in Figure 1 below. Note that X_1 and X_2 have the same mean. However, in looking at the histograms, we see that the possible values of X_2 are more "spread out" from the mean, indicating that the variance (and standard deviation) of X_2 is larger.




Figure 1: Histograms for random variables X_1 and X_2 , both with same expected value different variance.

Theorem 3.6.1 actually tells us how to compute variance, since it is given by finding the expected value of a *function* applied to the random variable. First, if X is a discrete random variable with possible values $x_1, x_2, \dots, x_i, \dots$, and probability mass function $p(x)$, then the variance of X is given by

$$\text{Var}(X) = \sum_i (x_i - \mu)^2 \cdot p(x_i).$$

The above formula follows directly from [Definition 3.7.1](#). However, there is an alternate formula for calculating variance, given by the following theorem, that is often easier to use.

Theorem 3.7.1

Let X be any random variable, with mean μ . Then the variance of X is

$$\text{Var}(X) = E[X^2] - \mu^2. \quad (3.7.1)$$

Proof

By the definition of *variance* ([Definition 3.7.1](#)) and the *linearity of expectation*, we have the following:

$$\begin{aligned} \text{Var}(X) &= E[(X - \mu)^2] \\ &= E[X^2 + \mu^2 - 2X\mu] \\ &= E[X^2] + E[\mu^2] - E[2X\mu] \\ &= E[X^2] + \mu^2 - 2\mu E[X] \quad (\text{Note: since } \mu \text{ is constant, we can take it out from the expected value}) \\ &= E[X^2] + \mu^2 - 2\mu^2 \\ &= E[X^2] - \mu^2 \end{aligned}$$

Example 3.7.2

Continuing in the context of [Example 3.6.1](#), we calculate the variance and standard deviation of the random variable X denoting the number of heads obtained in two tosses of a fair coin. Using the alternate formula for variance, we need to first calculate $E[X^2]$, for which we use [Theorem 3.6.1](#):

$$E[X^2] = 0^2 \cdot p(0) + 1^2 \cdot p(1) + 2^2 \cdot p(2) = 0 + 0.5 + 1 = 1.5.$$

In [Example 3.6.1](#), we found that $\mu = E[X] = 1$. Thus, we find

$$\begin{aligned} \text{Var}(X) &= E[X^2] - \mu^2 = 1.5 - 1 = 0.5 \\ \Rightarrow \text{SD}(X) &= \sqrt{\text{Var}(X)} = \sqrt{0.5} \approx 0.707 \end{aligned}$$

Exercise 3.7.1

Consider the context of [Example 3.6.2](#), where we defined the random variable X to be our winnings on a single play of game involving flipping a fair coin three times. We found that $E[X] = 1.25$. Now find the variance and standard deviation of X .

Answer

First, find $E[X^2]$:

$$\begin{aligned} E[X^2] &= \sum_i x_i^2 \cdot p(x_i) \\ &= (-1)^2 \cdot \frac{1}{8} + 1^2 \cdot \frac{1}{2} + 2^2 \cdot \frac{1}{4} + 3^2 \cdot \frac{1}{8} = \frac{11}{4} = 2.75 \end{aligned}$$

Now, we use the alternate formula for calculating variance:

$$\begin{aligned} \text{Var}(X) &= E[X^2] - E[X]^2 = 2.75 - 1.25^2 = 1.1875 \\ \Rightarrow \text{SD}(X) &= \sqrt{1.1875} \approx 1.0897 \end{aligned}$$

Given that the variance of a random variable is defined to be the expected value of *squared* deviations from the mean, variance is not linear as expected value is. We do have the following useful property of variance though.

Theorem 3.7.2

Let X be a random variable, and a, b be constants. Then the following holds:

$$\text{Var}(aX + b) = a^2 \text{Var}(X).$$

Exercise 3.7.2

Prove Theorem 3.7.2.

Answer

First, let $\mu = E[X]$ and note that by the linearity of expectation we have

$$E[aX + b] = aE[X] + b = a\mu + b.$$

Now, we use the alternate formula for variance given in [Theorem 3.7.1](#) to prove the result:

$$\begin{aligned} \text{Var}(aX + b) &= E[(aX + b)^2] - (E[aX + b])^2 \\ &= E[a^2X^2 + 2abX + b^2] - (a\mu + b)^2 \\ &= a^2E[X^2] + 2abE[X] + b^2 - a^2\mu^2 - 2ab\mu - b^2 \\ &= a^2E[X^2] - a^2\mu^2 = a^2(E[X^2] - \mu^2) = a^2\text{Var}(X) \end{aligned}$$

Theorem 3.7.2 easily follows from a little algebraic modification. Note that the "+ b " disappears in the formula. There is an intuitive reason for this. Namely, the "+ b " corresponds to a *horizontal shift* of the probability mass function for the random variable. Such a transformation to this function is not going to affect the *spread*, i.e., the variance will not change.

As with expected values, for many of the common probability distributions, the variance is given by a parameter or a function of the parameters for the distribution.

Variance for Discrete Distributions

Distribution	Variance
Bernoulli(p)	$p(1 - p)$
binomial(n, p)	$np(1 - p)$
hypergeometric(N, n, m)	$\frac{n(m/N)(1-m/N)(N-n)}{N-1}$
geometric(p)	$\frac{1-p}{p^2}$
negative binomial(r, p)	$\frac{r(1-p)}{p^2}$
Poisson(λ)	λ

This page titled [3.7: Variance of Discrete Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

3.8: Moment-Generating Functions (MGFs) for Discrete Random Variables

The expected value and variance of a random variable are actually special cases of a more general class of numerical characteristics for random variables given by *moments*.

Definition 3.8.1

The r^{th} **moment** of a random variable X is given by

$$E[X^r].$$

The r^{th} **central moment** of a random variable X is given by

$$E[(X - \mu)^r],$$

where $\mu = E[X]$.

Note that the **expected value** of a random variable is given by the **first moment**, i.e., when $r = 1$. Also, the **variance** of a random variable is given the **second central moment**.

As with expected value and variance, the moments of a random variable are used to characterize the distribution of the random variable and to compare the distribution to that of other random variables. Moments can be calculated directly from the definition, but, even for moderate values of r , this approach becomes cumbersome. The next definition and theorem provide an easier way to generate moments.

Definition 3.8.2

The **moment-generating function (mgf)** of a random variable X is given by

$$M_X(t) = E[e^{tX}], \quad \text{for } t \in \mathbb{R}.$$

Theorem 3.8.1

If random variable X has mgf $M_X(t)$, then

$$M_X^{(r)}(0) = \frac{d^r}{dt^r} [M_X(t)]_{t=0} = E[X^r].$$

In other words, the r^{th} derivative of the mgf evaluated at $t = 0$ gives the value of the r^{th} moment.

Theorem 3.8.1 tells us how to derive the mgf of a random variable, since the mgf is given by taking the expected value of a function applied to the random variable:

$$M_X(t) = E[e^{tX}] = \sum_i e^{tx_i} \cdot p(x_i)$$

We can now derive the first moment of the Poisson distribution, i.e., derive the fact we mentioned in Section 3.6, but left as an [exercise](#), that the expected value is given by the parameter λ . We also find the variance.

Example 3.8.1

Let $X \sim \text{Poisson}(\lambda)$. Then, the pmf of X is given by

$$p(x) = \frac{e^{-\lambda} \lambda^x}{x!}, \quad \text{for } x = 0, 1, 2, \dots$$

Before we derive the mgf for X , we recall from calculus the Taylor series expansion of the exponential function e^y :

$$e^y = \sum_{x=0}^{\infty} \frac{y^x}{x!}.$$

Using this fact, we find

$$M_X(t) = E[e^{tX}] = \sum_{x=0}^{\infty} e^{tx} \cdot \frac{e^{-\lambda} \lambda^x}{x!} = e^{-\lambda} \sum_{x=0}^{\infty} \frac{(e^t \lambda)^x}{x!} = e^{-\lambda} e^{e^t \lambda} = e^{\lambda(e^t - 1)}.$$

Now we take the first and second derivatives of $M_X(t)$. Remember we are differentiating with respect to t :

$$\begin{aligned} M'_X(t) &= \frac{d}{dt} \left[e^{\lambda(e^t - 1)} \right] = \lambda e^t e^{\lambda(e^t - 1)} \\ M''_X(t) &= \frac{d}{dt} \left[\lambda e^t e^{\lambda(e^t - 1)} \right] = \lambda e^t e^{\lambda(e^t - 1)} + \lambda^2 e^{2t} e^{\lambda(e^t - 1)} \end{aligned}$$

Next we evaluate the derivatives at $t = 0$ to find the first and second moments of X :

$$\begin{aligned} E[X] &= M'_X(0) = \lambda e^0 e^{\lambda(e^0 - 1)} = \lambda \\ E[X^2] &= M''_X(0) = \lambda e^0 e^{\lambda(e^0 - 1)} + \lambda^2 e^{2 \cdot 0} e^{\lambda(e^0 - 1)} = \lambda + \lambda^2 \end{aligned}$$

Finally, in order to find the variance, we use the alternate formula:

$$\text{Var}(X) = E[X^2] - (E[X])^2 = \lambda + \lambda^2 - \lambda^2 = \lambda.$$

Thus, we have shown that both the mean and variance for the Poisson(λ) distribution is given by the parameter λ .

Note that the mgf of a random variable is a *function* of t . The main application of mgf's is to find the moments of a random variable, as the previous example demonstrated. There are more properties of mgf's that allow us to find moments for functions of random variables.

Theorem 3.8.2

Let X be a random variable with mgf $M_X(t)$, and let a, b be constants. If random variable $Y = aX + b$, then the mgf of Y is given by

$$M_Y(t) = e^{bt} M_X(at).$$

Theorem 3.8.3

If X_1, \dots, X_n are independent random variables with mgf's $M_{X_1}(t), \dots, M_{X_n}(t)$, respectively, then the mgf of random variable $Y = X_1 + \dots + X_n$ is given by

$$M_Y(t) = M_{X_1}(t) \cdots M_{X_n}(t).$$

Recall that a binomially distributed random variable can be written as a sum of independent Bernoulli random variables. We use this and Theorem 3.8.3 to derive the mean and variance for a binomial distribution. First, we find the mean and variance of a Bernoulli distribution.

Example 3.8.2

Recall that X has a Bernoulli(p) distribution if it is assigned the value of 1 with probability p and the value of 0 with probability $1 - p$. Thus, the pmf of X is given by

$$p(x) = \begin{cases} 1 - p, & \text{if } x = 0 \\ p, & \text{if } x = 1 \end{cases}$$

In order to find the mean and variance of X , we first derive the mgf:

$$M_X(t) = E[e^{tX}] = e^{t(0)}(1-p) + e^{t(1)}p = 1 - p + e^t p.$$

Now we differentiate $M_X(t)$ with respect to t :

$$\begin{aligned} M'_X(t) &= \frac{d}{dt} [1 - p + e^t p] = e^t p \\ M''_X(t) &= \frac{d}{dt} [e^t p] = e^t p \end{aligned}$$

Next we evaluate the derivatives at $t = 0$ to find the first and second moments:

$$M'_X(0) = M''_X(0) = e^0 p = p.$$

Thus, the expected value of X is $E[X] = p$. Finally, we use the alternate formula for calculating variance:

$$\text{Var}(X) = E[X^2] - (E[X])^2 = p - p^2 = p(1 - p).$$

Example 3.8.3

Let $X \sim \text{binomial}(n, p)$. If X_1, \dots, X_n denote n independent Bernoulli(p) random variables, then we can write

$$X = X_1 + \dots + X_n.$$

In Example 3.8.2, we found the mgf for a Bernoulli(p) random variable. Thus, we have

$$M_{X_i}(t) = 1 - p + e^t p, \quad \text{for } i = 1, \dots, n.$$

Using Theorem 3.8.3, we derive the mgf for X :

$$M_X(t) = M_{X_1}(t) \cdots M_{X_n}(t) = (1 - p + e^t p) \cdots (1 - p + e^t p) = (1 - p + e^t p)^n.$$

Now we can use the mgf of X to find the moments:

$$\begin{aligned} M'_X(t) &= \frac{d}{dt} [(1 - p + e^t p)^n] = n(1 - p + e^t p)^{n-1} e^t p \\ &\Rightarrow M'_X(0) = np \\ M''_X(t) &= \frac{d}{dt} [n(1 - p + e^t p)^{n-1} e^t p] = n(n-1)(1 - p + e^t p)^{n-2} (e^t p)^2 + n(1 - p + e^t p)^{n-1} e^t p \\ &\Rightarrow M''_X(0) = n(n-1)p^2 + np \end{aligned}$$

Thus, the expected value of X is $E[X] = np$, and the variance is

$$\text{Var}(X) = E[X^2] - (E[X])^2 = n(n-1)p^2 + np - (np)^2 = np(1 - p).$$

We end with a final property of mgf's that relates to the comparison of the distribution of random variables.

Theorem 3.8.4

The mgf $M_X(t)$ of random variable X uniquely determines the probability distribution of X . In other words, if random variables X and Y have the same mgf, $M_X(t) = M_Y(t)$, then X and Y have the same probability distribution.

Exercise 3.8.1

Suppose the random variable X has the following mgf:

$$M_X(t) = (0.85 + 0.15e^t)^{33}$$

What is the distribution of X ?

Hint

Use [Theorem 3.8.4](#) and look at [Example 3.8.3](#).

Answer

We found in [Example 3.8.3](#) that the mgf for a binomial distribution is

$$M_X(t) = (1 - p + e^t p)^n,$$

which is the mgf given with $p = 0.15$ and $n = 33$. Thus, $X \sim \text{binomial}(33, 0.15)$

This page titled [3.8: Moment-Generating Functions \(MGFs\) for Discrete Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

CHAPTER OVERVIEW

4: Continuous Random Variables

Topic hierarchy

- 4.1: Probability Density Functions (PDFs) and Cumulative Distribution Functions (CDFs) for Continuous Random Variables
- 4.2: Expected Value and Variance of Continuous Random Variables
- 4.3: Uniform Distributions
- 4.4: Normal Distributions
- 4.5: Exponential and Gamma Distributions
- 4.6: Weibull Distributions
- 4.7: Chi-Squared Distributions
- 4.8: Beta Distributions

This page titled [4: Continuous Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

4.1: Probability Density Functions (PDFs) and Cumulative Distribution Functions (CDFs) for Continuous Random Variables

Probability Density Functions (PDFs)

Recall that continuous random variables have uncountably many possible values (think of intervals of real numbers). Just as for discrete random variables, we can talk about probabilities for continuous random variables using *density functions*.

Definition 4.1.1

The **probability density function (pdf)**, denoted f , of a continuous random variable X satisfies the following:

1. $f(x) \geq 0$, for all $x \in \mathbb{R}$
2. f is piecewise continuous
3. $\int_{-\infty}^{\infty} f(x) dx = 1$
4. $P(a \leq X \leq b) = \int_a^b f(x) dx$

The first three conditions in the definition state the properties necessary for a function to be a **valid pdf** for a continuous random variable. The fourth condition tells us how to use a pdf to calculate probabilities for continuous random variables, which are given by *integrals* the continuous analog to sums.

Example 4.1.1

Let the random variable X denote the time a person waits for an elevator to arrive. Suppose the longest one would need to wait for the elevator is 2 minutes, so that the possible values of X (in minutes) are given by the interval $[0, 2]$. A possible pdf for X is given by

$$f(x) = \begin{cases} x, & \text{for } 0 \leq x \leq 1 \\ 2 - x, & \text{for } 1 < x \leq 2 \\ 0, & \text{otherwise} \end{cases}$$

The graph of f is given below, and we verify that f satisfies the first three conditions in Definition 4.1.1:

1. From the graph, it is clear that $f(x) \geq 0$, for all $x \in \mathbb{R}$.
2. Since there are no holes, jumps, asymptotes, we see that $f(x)$ is (piecewise) continuous.
3. Finally we compute:

$$\int_{-\infty}^{\infty} f(x) dx = \int_0^2 f(x) dx = \int_0^1 x dx + \int_1^2 (2 - x) dx = 1$$




Figure 1: Graph of pdf for X , $f(x)$

So, if we wish to calculate the probability that a person waits less than 30 seconds (or 0.5 minutes) for the elevator to arrive, then we calculate the following probability using the pdf and the fourth property in Definition 4.1.1:

$$P(0 \leq X \leq 0.5) = \int_0^{0.5} f(x) dx = \int_0^{0.5} x dx = 0.125$$

Note that, unlike discrete random variables, continuous random variables have **zero point probabilities**, i.e., the probability that a continuous random variable equals a single value is always given by 0. Formally, this follows from properties of integrals:

$$P(X = a) = P(a \leq X \leq a) = \int_a^a f(x) dx = 0.$$

Informally, if we realize that **probability for a continuous random variable is given by areas under pdf's**, then, since there is no area in a line, there is no probability assigned to a random variable taking on a single value. This does not mean that a continuous random variable will never equal a single value, only that we do not assign any probability to single values for the random variable. For this reason, we only talk about the probability of a continuous random variable taking a value in an INTERVAL, not at a point. And whether or not the endpoints of the interval are included does not affect the probability. In fact, the following probabilities are all equal:

$$P(a \leq X \leq b) = P(a < X < b) = P(a \leq X < b) = P(a < X \leq b) = \int_a^b f(x) dx$$

Cumulative Distribution Functions (CDFs)

Recall Definition 3.2.2, the definition of the cdf, which applies to both discrete and continuous random variables. For continuous random variables we can further specify how to calculate the cdf with a formula as follows. Let X have pdf f , then the cdf F is given by

$$F(x) = P(X \leq x) = \int_{-\infty}^x f(t) dt, \quad \text{for } x \in \mathbb{R}.$$

In other words, the cdf for a continuous random variable is found by *integrating* the pdf. Note that the **Fundamental Theorem of Calculus** implies that the pdf of a continuous random variable can be found by *differentiating* the cdf. This relationship between the pdf and cdf for a continuous random variable is incredibly useful.

Relationship between PDF and CDF for a Continuous Random Variable

Let X be a continuous random variable with pdf f and cdf F .

- By definition, the cdf is found by *integrating* the pdf:

$$F(x) = \int_{-\infty}^x f(t) dt$$

- By the Fundamental Theorem of Calculus, the pdf can be found by *differentiating* the cdf:

$$f(x) = \frac{d}{dx}[F(x)]$$

Example 4.1.2

Continuing in the context of [Example 4.1.1](#), we find the corresponding cdf. First, let's find the cdf at two possible values of X , $x = 0.5$ and $x = 1.5$:

$$F(0.5) = \int_{-\infty}^{0.5} f(t) dt = \int_0^{0.5} t dt = \left. \frac{t^2}{2} \right|_0^{0.5} = 0.125$$

$$F(1.5) = \int_{-\infty}^{1.5} f(t) dt = \int_0^1 t dt + \int_1^{1.5} (2-t) dt = \left. \frac{t^2}{2} \right|_0^1 + \left(2t - \frac{t^2}{2} \right) \Big|_1^{1.5} = 0.5 + (1.875 - 1.5) = 0.875$$

Now we find $F(x)$ more generally, working over the intervals that $f(x)$ has different formulas:

$$\text{for } x < 0: \quad F(x) = \int_{-\infty}^x 0 dt = 0$$

$$\text{for } 0 \leq x \leq 1: \quad F(x) = \int_0^x t dt = \left. \frac{t^2}{2} \right|_0^x = \frac{x^2}{2}$$

$$\text{for } 1 < x \leq 2: \quad F(x) = \int_0^1 t dt + \int_1^x (2-t) dt = \left. \frac{t^2}{2} \right|_0^1 + \left(2t - \frac{t^2}{2} \right) \Big|_1^x = 0.5 + \left(2x - \frac{x^2}{2} \right) - (2 - 0.5) = 2x - \frac{x^2}{2} - 1$$

$$\text{for } x > 2: \quad F(x) = \int_{-\infty}^x f(t) dt = 1$$

Putting this altogether, we write F as a piecewise function and Figure 2 gives its graph:

$$F(x) = \begin{cases} 0, & \text{for } x < 0 \\ \frac{x^2}{2}, & \text{for } 0 \leq x \leq 1 \\ 2x - \frac{x^2}{2} - 1, & \text{for } 1 < x \leq 2 \\ 1, & \text{for } x > 2 \end{cases}$$




Figure 2: Graph of cdf in Example 4.1.2

Recall that the graph of the cdf for a discrete random variable is always a step function. Looking at Figure 2 above, we note that the cdf for a continuous random variable is always a *continuous* function.

Percentiles of a Distribution

Definition 4.1.2

The **(100p)th percentile** ($0 \leq p \leq 1$) of a probability distribution with cdf F is the value π_p such that

$$F(\pi_p) = P(X \leq \pi_p) = p.$$

To find the percentile π_p of a continuous random variable, which is a possible value of the random variable, we are specifying a cumulative probability p and solving the following equation for π_p :

$$\int_{-\infty}^{\pi_p} f(t)dt = p$$

Special Cases: There are a few values of p for which the corresponding percentile has a special name.

- **Median** or 50th percentile: $\pi_{.5} = \mu = Q_2$, separates probability (area under pdf) into two equal halves
- **1st Quartile** or 25th percentile: $\pi_{.25} = Q_1$, separates 1st quarter (25%) of probability (area) from the rest
- **3rd Quartile** or 75th percentile: $\pi_{.75} = Q_3$, separates 3rd quarter (75%) of probability (area) from the rest

Example 4.1.3

Continuing in the context of Example 4.1.2, we find the median and quartiles.

- **median:** find $\pi_{.5}$, such that $F(\pi_{.5}) = 0.5 \Rightarrow \pi_{.5} = 1$ (from graph in Figure 1)
- **1st quartile:** find $Q_1 = \pi_{.25}$, such that $F(\pi_{.25}) = 0.25$. For this, we use the formula and the graph of the cdf in Figure 2:

$$\frac{\pi_{.25}^2}{2} = 0.25 \Rightarrow Q_1 = \pi_{.25} = \sqrt{0.5} \approx 0.707$$

- **3rd quartile:** find $Q_3 = \pi_{.75}$, such that $F(\pi_{.75}) = 0.75$. Again, use the graph of the cdf:

$$2\pi_{.75} - \frac{\pi_{.75}^2}{2} - 1 = 0.75 \Rightarrow \text{(using Quadratic Formula)} Q_3 = \pi_{.75} = \frac{4 - \sqrt{2}}{2} \approx 1.293$$

This page titled [4.1: Probability Density Functions \(PDFs\) and Cumulative Distribution Functions \(CDFs\) for Continuous Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

4.2: Expected Value and Variance of Continuous Random Variables

We now consider the expected value and variance for continuous random variables. Note that the interpretation of each is the same as in the discrete setting, but we now have a different method of calculating them in the continuous setting.

Definition 4.2.1

If X is a continuous random variable with pdf $f(x)$, then the **expected value** (or **mean**) of X is given by

$$\mu = \mu_X = E[X] = \int_{-\infty}^{\infty} x \cdot f(x) dx.$$

The formula for the expected value of a continuous random variable is the continuous analog of the expected value of a discrete random variable, where instead of *summing* over all possible values we *integrate* (recall [Sections 3.6 & 3.7](#)).

For the **variance** of a continuous random variable, the definition is the same and we can still use the alternative formula given by [Theorem 3.7.1](#), only we now integrate to calculate the value:

$$\text{Var}(X) = E[X^2] - \mu^2 = \left(\int_{-\infty}^{\infty} x^2 \cdot f(x) dx \right) - \mu^2$$

Example 4.2.1

Consider again the context of [Example 4.1.1](#), where we defined the continuous random variable X to denote the time a person waits for an elevator to arrive. The pdf of X was given by

$$f(x) = \begin{cases} x, & \text{for } 0 \leq x \leq 1 \\ 2 - x, & \text{for } 1 < x \leq 2 \\ 0, & \text{otherwise} \end{cases}$$

Applying Definition 4.2.1, we compute the expected value of X :

$$E[X] = \int_0^1 x \cdot x dx + \int_1^2 x \cdot (2 - x) dx = \int_0^1 x^2 dx + \int_1^2 (2x - x^2) dx = \frac{1}{3} + \frac{2}{3} = 1.$$

Thus, we expect a person will wait 1 minute for the elevator on average. Figure 1 demonstrates the graphical representation of the expected value as the center of mass of the pdf.




Figure 1: The red arrow represents the center of mass, or the expected value, of X .

Now we calculate the variance and standard deviation of X , by first finding the expected value of X^2 .

$$\mathbb{E}[X^2] = \int_0^1 x^2 \cdot x \, dx + \int_1^2 x^2 \cdot (2-x) \, dx = \int_0^1 x^3 \, dx + \int_1^2 (2x^2 - x^3) \, dx = \frac{1}{4} + \frac{11}{12} = \frac{7}{6}.$$

Thus, we have

$$\begin{aligned}\text{Var}(X) &= \mathbb{E}[X^2] - \mu^2 = \frac{7}{6} - 1 = \frac{1}{6} \\ \Rightarrow \text{SD}(X) &= \sqrt{\text{Var}(X)} = \frac{1}{\sqrt{6}} \approx 0.408\end{aligned}$$

This page titled [4.2: Expected Value and Variance of Continuous Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

4.3: Uniform Distributions

Definition 4.3.1

A random variable X has a **uniform distribution** on interval $[a, b]$, write $X \sim \text{uniform}[a, b]$, if it has pdf given by

$$f(x) = \begin{cases} \frac{1}{b-a}, & \text{for } a \leq x \leq b \\ 0, & \text{otherwise} \end{cases}$$

The uniform distribution is also sometimes referred to as the **box distribution**, since the graph of its pdf looks like a box. See Figure 1 below.




Figure 1: Graph of the pdf for a uniform distribution on interval $[a, b]$

Exercise 4.3.1

Verify that the uniform pdf is a valid pdf, i.e., show that it satisfies the first three conditions of [Definition 4.1.1](#).

Answer

1. In looking either at the formula in Definition 4.3.1 or the graph in Figure 1, we can see that the uniform pdf is always non-negative, i.e., $f(x) \geq 0$, for all $x \in \mathbb{R}$.
2. Given that the uniform pdf is a piecewise constant function, it is also piecewise continuous.
3. Finally, we need to verify that the area under the uniform pdf is equal to 1. This is quickly seen from the graph in Figure 1, since we calculate the area of rectangle with width $(b - a)$ and height $1/(b - a)$. Thus, the area is

$$(b - a) \times \frac{1}{(b - a)} = 1.$$

A typical application of the uniform distribution is to model randomly generated numbers. In other words, it provides the probability distribution for a random variable representing a randomly chosen number between numbers a and b .

The uniform distribution assigns equal probabilities to intervals of equal lengths, since it is a constant function, on the interval it is non-zero $[a, b]$. This is the continuous analog to equally likely outcomes in the discrete setting.

We close the section by finding the expected value of the uniform distribution.

Example 4.3.1

If X has a uniform distribution on the interval $[a, b]$, then we apply [Definition 4.2.1](#) and compute the expected value of X :

$$E[X] = \int_a^b x \cdot \frac{1}{b-a} dx = \frac{b^2 - a^2}{2} \cdot \frac{1}{b-a} = \frac{(b-a)(b+a)}{2} \cdot \frac{1}{b-a} = \frac{b+a}{2}.$$

Thus, the expected value of the uniform $[a, b]$ distribution is given by the average of the parameters a and b , or the midpoint of the interval $[a, b]$. This is readily apparent when looking at a graph of the pdf in Figure 1 and remembering the interpretation of expected value as the center of mass. Since the pdf is constant over $[a, b]$, the center of mass is simply given by the midpoint.

This page titled [4.3: Uniform Distributions](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

4.4: Normal Distributions

Definition 4.4.1

A random variable X has a **normal distribution**, with parameters μ and σ , write $X \sim \text{normal}(\mu, \sigma)$, if it has pdf given by

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}, \quad \text{for } x \in \mathbb{R},$$

where $\mu \in \mathbb{R}$ and $\sigma > 0$.

If a continuous random variable X has a normal distribution with parameters μ and σ , then $E[X] = \mu$ and $\text{Var}(X) = \sigma^2$. These facts can be derived using Definition 4.2.1; however, the integral calculations require many tricks. Note that the normal case is why the notation μ is often used for the expected value, and σ^2 is used for the variance. So, μ gives the center of the normal pdf, and its graph is symmetric about μ , while σ determines how spread out the graph is. Figure 1 below shows the graph of two different normal pdf's.

Example 4.4.1

Suppose $X_1 \sim \text{normal}(0, 2^2)$ and $X_2 \sim \text{normal}(0, 3^2)$. So, X_1 and X_2 are both normally distributed random variables with the same mean, but X_2 has a larger standard deviation. Given our interpretation of standard deviation, this implies that the possible values of X_2 are more "spread out" from the mean. This is easily seen by looking at the graphs of the pdf's corresponding to X_1 and X_2 given in Figure 1.




Figure 1: Graph of normal pdf's: $X_1 \sim \text{normal}(0, 2^2)$ in blue, $X_2 \sim \text{normal}(0, 3^2)$ in red

The normal distribution is arguably the most important probability distribution. It is used to model the distribution of population characteristics such as weight, height, and IQ. The pdf is terribly tricky to work with, in fact integrals involving the normal pdf cannot be solved exactly, but rather require numerical methods to approximate. Because of this, there is no closed form for the corresponding cdf of a normal distribution. Given the importance of the normal distribution though, many software programs have been built in normal probability calculators. There are also many useful properties of the normal distribution that make it easy to work with. We state these properties without proof below. Note that we also include the connection to expected value and variance given by the parameters.

Properties of the Normal Distribution

1. If $X \sim \text{normal}(\mu, \sigma)$, then $aX + b$ also follows a normal distribution with parameters $a\mu + b$ and $a\sigma$.
2. If $X \sim \text{normal}(\mu, \sigma)$, then $\frac{X - \mu}{\sigma}$ follows the **standard normal distribution**, i.e., the normal distribution with parameters $\mu = 0$ and $\sigma = 1$.

The first property says that any linear transformation of a normally distributed random variable is also normally distributed. The second property is a special case of the first, since we can re-write the transformation on X as

$$\frac{X - \mu}{\sigma} = \left(\frac{1}{\sigma} \right) X - \frac{\mu}{\sigma}.$$

This transformation, subtracting the mean and dividing by the standard deviation, is referred to as **standardizing** X , since the resulting random variable will *always* have the standard normal distribution with mean 0 and standard deviation 1. In this way, standardizing a normal random variable has the effect of removing the units. Before the prevalence of calculators and computer software capable of calculating normal probabilities, people would apply the standardizing transformation to the normal random variable and use a table of probabilities for the standard normal distribution.

This page titled [4.4: Normal Distributions](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

4.5: Exponential and Gamma Distributions

In this section, we introduce two families of continuous probability distributions that are commonly used.

Exponential Distributions

Definition 4.5.1

A random variable X has an **exponential distribution** with parameter $\lambda > 0$, write $X \sim \text{exponential}(\lambda)$, if X has pdf given by

$$f(x) = \begin{cases} \lambda e^{-\lambda x}, & \text{for } x \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$




Figure 1: Graph of pdf for exponential($\lambda = 5$) distribution.

Example 4.5.1

A typical application of exponential distributions is to model *waiting times* or *lifetimes*. For example, each of the following gives an application of an exponential distribution.

- X = lifetime of a radioactive particle
- X = how long you have to wait for an accident to occur at a given intersection
- X = length of interval between consecutive occurrences of Poisson distributed events

The parameter λ is referred to as the **rate parameter**, it represents how quickly events occur. For example, in the first case above where X denotes the lifetime of a radioactive particle, λ would give the rate at which such particles decay.

Properties of Exponential Distributions

If $X \sim \text{exponential}(\lambda)$, then the following hold.

1. The cdf of X is given by

$$F(x) = \begin{cases} 0, & \text{for } x < 0, \\ 1 - e^{-\lambda x}, & \text{for } x \geq 0. \end{cases}$$

2. For any $0 < p < 1$, the $(100p)^{\text{th}}$ percentile is $\pi_p = \frac{-\ln(1-p)}{\lambda}$.

3. The mean of X is $E[X] = \frac{1}{\lambda}$.

4. The variance of X is $\text{Var}(X) = \frac{1}{\lambda^2}$.

5. The mgf of X is

$$M_X(t) = \frac{1}{1 - (t/\lambda)}, \quad \text{for } t < \lambda.$$

6. X satisfies the **Memoryless Property**, i.e., $P(X > t + s \mid X > s) = P(X > t)$, for any $t, s \geq 0$.

Partial Proof

We prove Properties #1 & #3, the others are left as an exercise.

For the first property, we consider two cases based on the value of x . First, if $x < 0$, then the pdf is constant and equal to 0, which gives the following for the cdf:

$$F(x) = \int_{-\infty}^x f(t)dt = \int_{-\infty}^x 0dt = 0$$

Second, if $x \geq 0$, then the pdf is $\lambda e^{-\lambda x}$, and the cdf is given by

$$F(x) = \int_{-\infty}^x f(t)dt = \int_0^x \lambda e^{-\lambda t} dt = -e^{-\lambda t} \Big|_0^x = -e^{-\lambda x} - (-e^0) = 1 - e^{-\lambda x}.$$

For the third property, we use Definition 4.2.1 to calculate the expected value of a continuous random variable:

$$E[X] = \int_{-\infty}^{\infty} x \cdot f(x)dx = \int_0^{\infty} x \cdot \lambda e^{-\lambda x} dx = -x \cdot e^{-\lambda x} \Big|_0^{\infty} + \int_0^{\infty} e^{-\lambda x} dx = 0 + \frac{-e^{-\lambda x}}{\lambda} \Big|_0^{\infty} = \frac{1}{\lambda}.$$

In words, the **Memoryless Property** of exponential distributions states that, given that you have already waited more than s units of time ($X > s$), the *conditional* probability that you will have to wait t more ($X > t + s$) is equal to the *unconditional* probability you just have to wait more than t units of time. For example, suppose you are waiting for the bus and the amount of time you have to wait is exponentially distributed. If you have already been waiting 5 minutes at the bus stop, the probability that you have to wait 4 more minutes (so more than 9 minutes total) is equal to the probability that you only had to wait more than 4 minutes once arriving at the bus stop. In calculating the conditional probability, the exponential distribution "forgets" about the condition or the time already spent waiting and you can just calculate the unconditional probability that you have to wait longer. Note that we saw earlier that geometric distributions also have the Memoryless Property.

Gamma Distributions

Definition 4.5.2

A random variable X has a **gamma distribution** with parameters $\alpha, \lambda > 0$, write $X \sim \text{gamma}(\alpha, \lambda)$, if X has pdf given by

$$f(x) = \begin{cases} \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x}, & \text{for } x \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

where $\Gamma(\alpha)$ is a function (referred to as the gamma function) given by the following integral:

$$\Gamma(\alpha) = \int_0^{\infty} t^{\alpha-1} e^{-t} dt.$$

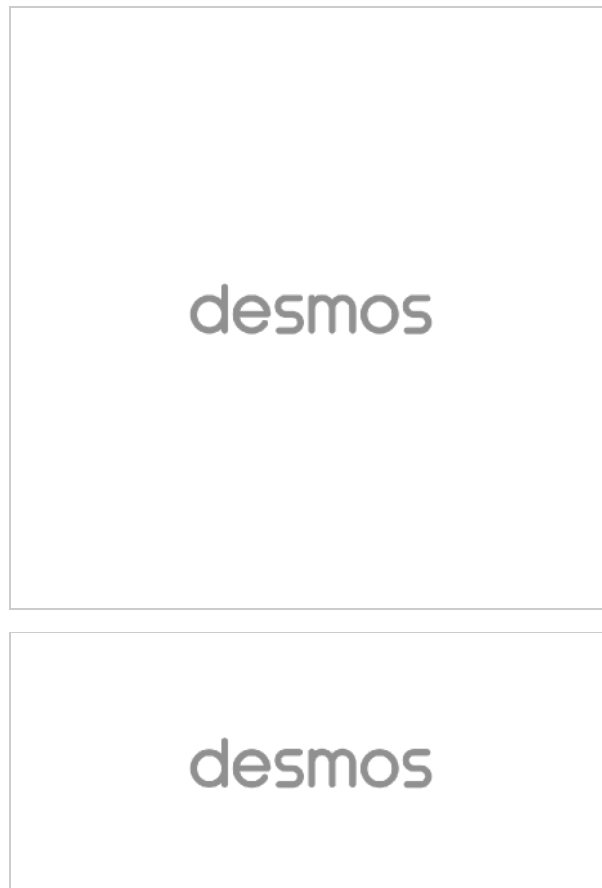


Figure 2: Graph of pdf's for various gamma distributions. On the left, for the purple pdf $\alpha = 0.5$ and for the green pdf $\alpha = 1.5$. On the right, for the blue pdf $\alpha = 4$ and for the orange pdf $\alpha = 8$. For all pdf's, $\lambda = 5$.

Note that the gamma function, $\Gamma(\alpha)$, ensures that the gamma pdf is valid, i.e., that it integrates to 1, which you are asked to show in the following exercise. The value of $\Gamma(\alpha)$ depends on the value of the parameter α , but for a given value of α it is just a number, i.e., it is a constant value in the gamma pdf, given specific parameter values. In this case, $\Gamma(\alpha)$ is referred to as a *scaling constant*, since it "scales" the rest of the pdf, $\lambda^\alpha x^{\alpha-1} e^{-\lambda x}$, which is referred to as the *kernel* of the distribution, so that the result integrates to 1.

Exercise 4.5.1

Show: $\int_0^\infty \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} dx = 1$

Answer

In the integral, we can make the substitution: $u = \lambda x \rightarrow du = \lambda dx$. Therefore, we have

$$\int_0^\infty \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} dx = \int_0^\infty \frac{\lambda \lambda^{\alpha-1}}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} dx = \frac{1}{\Gamma(\alpha)} \int_0^\infty u^{\alpha-1} e^{-u} du = \frac{1}{\Gamma(\alpha)} \Gamma(\alpha) = 1.$$

Properties of Gamma Distributions

If $X \sim \text{gamma}(\alpha, \lambda)$, then the following hold.

1. The mgf of X is

$$M_X(t) = \frac{1}{(1 - (t/\lambda))^\alpha}, \quad \text{for } t < \lambda.$$

2. The mean of X is $E[X] = \frac{\alpha}{\lambda}$.
3. The variance of X is $\text{Var}(X) = \frac{\alpha}{\lambda^2}$.

Notes about Gamma Distributions:

- If $\alpha = 1$, then the corresponding gamma distribution is given by the exponential distribution, i.e., $\text{gamma}(1, \lambda) = \text{exponential}(\lambda)$. This is left as an exercise for the reader.
- The parameter α is referred to as the **shape parameter**, and λ is the **rate parameter**. Varying the value of α changes the shape of the pdf, as is seen in Figure 2 above, whereas varying the value of λ corresponds to changing the units (e.g., from inches to centimeters) and does not alter the shape of the pdf.
- A closed form does not exist for the cdf of a gamma distribution, computer software must be used to calculate gamma probabilities. [Here is a link to a gamma calculator online](#). (Note that different notation is used on this online calculator, namely, λ is referred to as β instead.)

Example 4.5.2

A typical application of gamma distributions is to model the time it takes for a given number of events to occur. For example, each of the following gives an application of a gamma distribution.

- X = lifetime of 5 radioactive particles
- X = how long you have to wait for 3 accidents to occur at a given intersection

In these examples, the parameter λ represents the rate at which the event occurs, and the parameter α is the number of events desired. So, in the first example, $\alpha = 5$ and λ represents the rate at which particles decay.

This page titled [4.5: Exponential and Gamma Distributions](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

4.6: Weibull Distributions

In this section, we introduce the Weibull distributions, which are very useful in the field of actuarial science.

Weibull Distributions

Definition 4.6.1

A random variable X has a **Weibull distribution** with parameters $\alpha, \beta > 0$, write $X \sim \text{Weibull}(\alpha, \beta)$, if X has pdf given by

$$f(x) = \begin{cases} \frac{\alpha}{\beta^\alpha} x^{\alpha-1} e^{-(x/\beta)^\alpha}, & \text{for } x \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$




Figure 1: Graph of pdf for Weibull($\alpha = 2, \beta = 5$) distribution.

Example 4.6.1

A typical application of Weibull distributions is to model *lifetimes* that are not “memoryless”. For example, each of the following gives an application of the Weibull distribution.

- modeling the lifetime of a car battery
- modeling the probability that someone survives past the age of 80 years old

The parameter α is referred to as the **shape parameter**, and β is the **scale parameter**. When $\alpha = 1$, the Weibull distribution is an exponential distribution with $\lambda = 1/\beta$, so the exponential distribution is a special case of both the Weibull distributions and the gamma distributions. We can see the similarities between the Weibull and exponential distributions more readily when comparing the cdf's of each. The cdf of the Weibull distribution is given below, with proof, along with other important properties, stated without proof.

Properties of Weibull Distributions

If $X \sim \text{Weibull}(\alpha, \beta)$, then the following hold.

1. The cdf of X is given by

$$F(x) = \begin{cases} 0 & \text{for } x < 0, \\ 1 - e^{-(x/\beta)^\alpha}, & \text{for } x \geq 0. \end{cases}$$

2. For any $0 < p < 1$, the $(100p)^{\text{th}}$ percentile is $\pi_p = \beta(-\ln(1-p))^{1/\alpha}$.

3. The mean of X is $E[X] = \beta\Gamma\left(1 + \frac{1}{\alpha}\right)$.

4. The variance of X is $\text{Var}(X) = \beta^2 \left[\Gamma\left(1 + \frac{2}{\alpha}\right) - \left[\Gamma\left(1 + \frac{1}{\alpha}\right) \right]^2 \right]$.

Partial Proof

We prove Property #1, but leave #2 as an exercise. Properties #3 and #4 are rather tricky to prove, so we state them without proof.

For the first property, we consider two cases based on the value of x . First, if $x < 0$, then the pdf is constant and equal to 0, which gives the following for the cdf:

$$F(x) = \int_{-\infty}^x f(t)dt = \int_{-\infty}^x 0dt = 0$$

Second, if $x \geq 0$, then the pdf is $\frac{\alpha}{\beta^\alpha} x^{\alpha-1} e^{-(x/\beta)^\alpha}$, and the cdf is given by the following integral, which is solved by making the substitution $u = \left(\frac{t}{\beta}\right)^\alpha$:

$$F(x) = \int_{-\infty}^x f(t)dt = \int_0^x \frac{\alpha}{\beta^\alpha} t^{\alpha-1} e^{-(t/\beta)^\alpha} dt = \int_0^{(x/\beta)^\alpha} e^{-u} du = -e^{-u} \Big|_0^{(x/\beta)^\alpha} = -e^{-(x/\beta)^\alpha} - (-e^0) = 1 - e^{-(x/\beta)^\alpha}.$$

4.6: Weibull Distributions is shared under a [not declared](#) license and was authored, remixed, and/or curated by LibreTexts.

4.7: Chi-Squared Distributions

In this section, we introduce the chi-squared distributions, which are very useful in statistics.

Chi-Squared Distributions

Definition 4.7.1

A random variable X has a **chi-squared distribution** with k degrees of freedom, where k is a positive integer, write $X \sim \chi^2(k)$, if X has pdf given by

$$f(x) = \begin{cases} \frac{1}{\Gamma(k/2)2^{k/2}} x^{k/2-1} e^{-x/2}, & \text{for } x \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$




Figure 1: Graph of pdf for $\chi^2(1)$ distribution.

The chi-squared distributions are a special case of the gamma distributions with $\alpha = \frac{k}{2}$, $\lambda = \frac{1}{2}$, which can be used to establish the following properties of the chi-squared distribution.

Properties of Chi-Squared Distributions

If $X \sim \chi^2(k)$, then X has the following properties.

1. The mgf of X is given by

$$M_X(t) = \frac{1}{(1-2t)^{k/2}}, \quad \text{for } t < \frac{1}{2}$$

2. The mean of X is $E[X] = k$, i.e., the degrees of freedom.
3. The variance of X is $\text{Var}(X) = 2k$, i.e., twice the degrees of freedom.

Note that there is no closed form equation for the cdf of a chi-squared distribution in general. But most graphing calculators have a built-in function to compute chi-squared probabilities. On the TI-84 or 89, this function is named " χ^2 cdf".

The main applications of the chi-squared distributions relate to their importance in the field of statistics, which result from the following relationships between the chi-squared distributions and the normal distributions.

Relationships of Chi-Squared Distributions

1. If Z is a standard normal random variable, i.e., $Z \sim N(0, 1)$, then the distribution of Z^2 is chi-squared with $k = 1$ degree of freedom.
2. If X_1, \dots, X_n is a collection of independent, chi-squared random variables each with 1 degree of freedom, i.e., $X_i \sim \chi^2(1)$, for each $i = 1, \dots, n$, then the sum $X_1 + \dots + X_n$ is also chi-squared with $k = n$ degrees of freedom.
3. If $X \sim \chi^2(k_1)$ and $Y \sim \chi^2(k_2)$ are independent random variables, then $X + Y \sim \chi^2(k_1 + k_2)$.

4.7: Chi-Squared Distributions is shared under a [not declared](#) license and was authored, remixed, and/or curated by LibreTexts.

4.8: Beta Distributions

In this section, we introduce beta distributions, which are very useful in a branch of statistics known as Bayesian Statistics.

Beta Distributions

Definition 4.8.1

A random variable X has a **beta distribution** with parameters $\alpha, \beta > 0$, write $X \sim \text{beta}(\alpha, \beta)$, if X has pdf given by

$$f(x) = \begin{cases} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1} (1-x)^{\beta-1}, & \text{for } 0 \leq x \leq 1, \\ 0 & \text{otherwise,} \end{cases} \quad (4.8.1)$$

Note that the gamma function, $\Gamma(\alpha)$, is defined in [Definition 4.5.2](#).

In the formula for the pdf of the beta distribution given in Equation 4.8.1, note that the term with the gamma functions, i.e., $\frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)}$ is the *scaling constant* so that the pdf is valid, i.e., integrates to 1. This is similar to the role the gamma function plays for the gamma distribution introduced in [Section 4.5](#). Ignoring the scaling constant for the beta distribution, we can focus on what is referred to as the *kernel* of the distribution, which is given by

$$x^{\alpha-1} (1-x)^{\beta-1}, \quad \text{for } x \in [0, 1]. \quad (4.8.2)$$

The parameters, α and β , are both **shape parameters** for the beta distribution, varying their values changes the shape of the pdf.

As is the case for the normal, gamma, and chi-squared distributions, there is no closed form equation for the cdf of the beta distribution and computer software must be used to calculate beta probabilities. [Here is a link to a beta calculator online.](#)

Beta distributions are useful for modeling random variables that only take values on the unit interval $[0, 1]$. In fact, if both parameters are equal to one, i.e., $\alpha = \beta = 1$, the corresponding beta distribution is equal to the uniform $[0, 1]$ distribution. In statistics, beta distributions are used to model proportions of random samples taken from a population that have a certain characteristic of interest. For example, the proportion of surface area in a randomly selected urban neighborhood that is green space, i.e., parks or garden area.

We state the following important properties of beta distributions without proof.

Properties of Beta Distributions

If $X \sim \text{beta}(\alpha, \beta)$, then:

1. the mean of X is $E[X] = \frac{\alpha}{\alpha + \beta}$,
2. the variance of X is $\text{Var}(X) = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}$.

4.8: Beta Distributions is shared under a [not declared](#) license and was authored, remixed, and/or curated by LibreTexts.

CHAPTER OVERVIEW

5: Probability Distributions for Combinations of Random Variables

[5.1: Joint Distributions of Discrete Random Variables](#)

[5.2: Joint Distributions of Continuous Random Variables](#)

[5.3: Conditional Probability Distributions](#)

[5.4: Finding Distributions of Functions of Continuous Random Variables](#)

[5.5: Sample Mean](#)

This page titled [5: Probability Distributions for Combinations of Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

5.1: Joint Distributions of Discrete Random Variables

In this chapter we consider two or more random variables defined on the same sample space and discuss how to model the probability distribution of the random variables *jointly*. We will begin with the discrete case by looking at the joint probability mass function for two discrete random variables. In the following section, we will consider continuous random variables.

Definition 5.1.1

If discrete random variables X and Y are defined on the same sample space S , then their **joint probability mass function (joint pmf)** is given by

$$p(x, y) = P(X = x \text{ and } Y = y),$$

where (x, y) is a pair of possible values for the pair of random variables (X, Y) , and $p(x, y)$ satisfies the following conditions:

- $0 \leq p(x, y) \leq 1$
- $\sum_{(x,y)} p(x, y) = 1$
- $P((X, Y) \in A) = \sum_{(x,y) \in A} p(x, y)$

Note that conditions #1 and #2 in Definition 5.1.1 are required for $p(x, y)$ to be a valid joint pmf, while the third condition tells us how to use the joint pmf to find probabilities for the pair of random variables (X, Y) .

In the discrete case, we can obtain the **joint cumulative distribution function (joint cdf)** of X and Y by *summing* the joint pmf:

$$F(x, y) = P(X \leq x \text{ and } Y \leq y) = \sum_{x_i \leq x} \sum_{y_j \leq y} p(x_i, y_j),$$

where x_i denotes possible values of X and y_j denotes possible values of Y . From the joint pmf, we can also obtain the individual probability distributions of X and Y *separately* as shown in the next definition.

Definition 5.1.2

Suppose that discrete random variables X and Y have joint pmf $p(x, y)$. Let $x_1, x_2, \dots, x_i, \dots$ denote the possible values of X , and let $y_1, y_2, \dots, y_j, \dots$ denote the possible values of Y . The **marginal probability mass functions (marginal pmf's)** of X and Y are respectively given by the following:

$$p_X(x) = \sum_j p(x, y_j) \quad (\text{fix a value of } X \text{ and sum over possible values of } Y)$$

$$p_Y(y) = \sum_i p(x_i, y) \quad (\text{fix a value of } Y \text{ and sum over possible values of } X)$$

[Link to Video: Overview of Definitions 5.1.1 & 5.1.2](#)

Example 5.1.1

Consider again the probability experiment of [Example 3.3.2](#), where we toss a fair coin three times and record the sequence of heads (h) and tails (t). Again, we let random variable X denote the number of heads obtained. We also let random variable Y denote the winnings earned in a single play of a game with the following rules, based on the outcomes of the probability experiment (this is the same as [Example 3.6.2](#)):

- player wins \$1 if first h occurs on the first toss
- player wins \$2 if first h occurs on the second toss

- player wins \$3 if first h occurs on the third toss
- player loses \$1 if no h occur

Note that the possible values of X are $x = 0, 1, 2, 3$, and the possible values of Y are $y = -1, 1, 2, 3$. We represent the joint pmf using a table:

Table 1: joint pmf of X and Y

$p(x, y)$	X			
Y	0	1	2	3
-1	1/8	0	0	0
1	0	1/8	2/8	1/8
2	0	1/8	1/8	0
3	0	1/8	0	0

The values in Table 1 give the values of $p(x, y)$. For example, consider $p(0, -1)$:

$$p(0, -1) = P(X = 0 \text{ and } Y = -1) = P(ttt) = \frac{1}{8}.$$

Since the outcomes are equally likely, the values of $p(x, y)$ are found by counting the number of outcomes in the sample space S that result in the specified values of the random variables, and then dividing by 8, the total number of outcomes in S . The sample space is given below, color coded to help explain the values of $p(x, y)$:

$$S = \{ttt, htt, tht, tth, hht, hth, thh, hhh\}$$

Given the joint pmf, we can now find the marginal pmf's. Note that the marginal pmf for X is found by computing sums of the *columns* in Table 1, and the marginal pmf for Y corresponds to the *row* sums. (Note that we found the pmf for X in [Example 3.3.2](#) as well, it is a binomial random variable. We also found the pmf for Y in [Example 3.6.2](#).)

Table 2: marginal pmf's for X and Y

x	$p_X(x)$	y	$p_Y(y)$
0	1/8	-1	1/8
1	3/8	1	1/2
2	3/8	2	1/4
3	1/8	3	1/8

Finally, we can find the joint cdf for X and Y by summing over values of the joint frequency function. For example, consider $F(1, 1)$:

$$F(1, 1) = P(X \leq 1 \text{ and } Y \leq 1) = \sum_{x \leq 1} \sum_{y \leq 1} p(x, y) = p(0, -1) + p(0, 1) + p(1, -1) + p(1, 1) = \frac{1}{4}$$

Again, we can represent the joint cdf using a table:

Table 3: joint cdf of X and Y

$F(x, y)$	X			
Y	0	1	2	3
-1	1/8	1/8	1/8	1/8
1	1/8	1/4	1/2	5/8
2	1/8	3/8	3/4	7/8
3	1/8	1/2	7/8	1

[Link to Video: Walkthrough of Example 5.1.1](#)

Expectations of Functions of Jointly Distributed Discrete Random Variables

We now look at taking the expectation of jointly distributed discrete random variables. Because expected values are defined for a single quantity, we will actually define the expected value of a combination of the pair of random variables, i.e., we look at the expected value of a function applied to (X, Y) .

Theorem 5.1.1

Suppose that X and Y are jointly distributed discrete random variables with joint pmf $p(x, y)$.

If $g(X, Y)$ is a function of these two random variables, then its expected value is given by the following:

$$E[g(X, Y)] = \sum_{(x, y)} g(x, y) p(x, y).$$

Example 5.1.2

Consider again the discrete random variables we defined in [Example 5.1.1](#) with joint pmf given in [Table 1](#). We will find the expected value of three different functions applied to (X, Y) .

1. First, we define $g(x, y) = xy$, and compute the expected value of XY :

$$\begin{aligned}
 E[XY] &= \sum_{(x, y)} xy \cdot p(x, y) = (0)(-1) \left(\frac{1}{8} \right) \\
 &\quad + (1)(1) \left(\frac{1}{8} \right) + (2)(1) \left(\frac{2}{8} \right) + (3)(1) \left(\frac{1}{8} \right) \\
 &\quad + (1)(2) \left(\frac{1}{8} \right) + (2)(2) \left(\frac{1}{8} \right) \\
 &\quad + (1)(3) \left(\frac{1}{8} \right) \\
 &= \frac{17}{8} = 2.125
 \end{aligned}$$

2. Next, we define $g(x) = x$, and compute the expected value of X :

$$\begin{aligned}
 E[X] &= \sum_{(x,y)} \sum x \cdot p(x,y) = (0) \left(\frac{1}{8}\right) \\
 &\quad + (1) \left(\frac{1}{8}\right) + (2) \left(\frac{2}{8}\right) + (3) \left(\frac{1}{8}\right) \\
 &\quad + (1) \left(\frac{1}{8}\right) + (2) \left(\frac{1}{8}\right) \\
 &\quad + (1) \left(\frac{1}{8}\right) \\
 &= \frac{12}{8} = 1.5
 \end{aligned}$$

Recall that $X \sim \text{binomial}(n = 3, p = 0.5)$, and that the expected value of a binomial random variable is given by np . Thus, we can verify the expected value of X that we calculated above using Theorem 5.1.1 using this fact for binomial distributions: $E[X] = np = 3(0.5) = 1.5$.

3. Lastly, we define $g(x, y) = y$, and calculate the expected value of Y :

$$\begin{aligned}
 E[Y] &= \sum_{(x,y)} \sum y \cdot p(x,y) = (-1) \left(\frac{1}{8}\right) \\
 &\quad + (1) \left(\frac{1}{8}\right) + (1) \left(\frac{2}{8}\right) + (1) \left(\frac{1}{8}\right) \\
 &\quad + (2) \left(\frac{1}{8}\right) + (2) \left(\frac{1}{8}\right) \\
 &\quad + (3) \left(\frac{1}{8}\right) \\
 &= \frac{10}{8} = 1.25
 \end{aligned}$$

Again, we can verify this result by reviewing the calculations done in [Example 3.6.2](#).

[Link to Video: Walkthrough of Example 5.1.2](#)

Independent Random Variables

In some cases, the probability distribution of one random variable will not be affected by the distribution of another random variable defined on the same sample space. In those cases, the joint distribution functions have a very simple form, and we refer to the random variables as independent.

Definition 5.1.3

Discrete random variables X_1, X_2, \dots, X_n are **independent** if the joint pmf factors into a product of the marginal pmf's:

$$p(x_1, x_2, \dots, x_n) = p_{X_1}(x_1) \cdot p_{X_2}(x_2) \cdots p_{X_n}(x_n). \quad (5.1.1)$$

It is equivalent to check that this condition holds for the cumulative distribution functions.

Recall the definition of independent *events* ([Definition 2.3.2](#)): A and B are independent events if $P(A \cap B) = P(A) P(B)$. This is the basis for the definition of independent random variables because we can write the pmf's in Equation 5.1.1 in terms of events as follows:

$$p(x, y) = P(X = x \text{ and } Y = y) = P(\{X = x\} \cap \{Y = y\}) = P(X = x)P(Y = y) = p_X(x)p_Y(y)$$

In the above, we use the idea that if X and Y are independent, then the event that X takes on a given value x is independent of the event that Y takes the value y .

Example 5.1.3

Consider yet again the discrete random variables defined in [Example 5.1.1](#). According to the definition, X and Y are independent if

$$p(x, y) = p_X(x) \cdot p_Y(y),$$

for all pairs (x, y) . Recall that the joint pmf for (X, Y) is given in [Table 1](#) and that the marginal pmf's for X and Y are given in [Table 2](#). Note that, for $(x, y) = (0, -1)$, we have the following

$$p(0, -1) = \frac{1}{8}, \quad p_X(0) = \frac{1}{8}, \quad p_Y(-1) = \frac{1}{8} \quad \Rightarrow \quad p(0, -1) \neq p_X(0) \cdot p_Y(-1).$$

Thus, X and Y are **not** independent, or in other words, X and Y are **dependent**. This should make sense given the definition of X and Y . The winnings earned depend on the number of heads obtained. So the probabilities assigned to the values of Y will be affected by the values of X .

We also have the following very useful theorem about the expected value of a product of independent random variables, which is simply given by the product of the expected values for the individual random variables.

Theorem 5.1.2

If X and Y are independent random variables, then $E[XY] = E[X] E[Y]$.

Proof

Assume X and Y are independent random variables. If we let $p(x, y)$ denote the joint pmf of (X, Y) , then, by [Definition 5.1.3](#), $p(x, y) = p_X(x)p_Y(y)$, for all pairs (x, y) . Using this fact and [Theorem 5.1.1](#), we have

$$\begin{aligned} E[XY] &= \sum_{(x,y)} xy \cdot p(x, y) = \sum_{(x,y)} xy \cdot p_X(x)p_Y(y) \\ &= \sum_x \sum_y xyp(x)p_Y(y) = \sum_x xp_X(x) \left(\sum_y p_Y(y) \right) = \sum_x xp_X(x)E[Y] \\ &= E[Y] \sum_x xp_X(x) = E[Y] E[X]. \end{aligned}$$

Theorem 5.1.2 can be used to show that two random variables are **not** independent: if $E[XY] \neq E[X] E[Y]$, then X and Y **cannot** be independent. However, beware using Theorem 5.1.2 to show that random variables are independent. Note that Theorem 5.1.2 **assumes** that X and Y are independent and then the property about the expected value follows. The other direction does not hold. In other words, if $E[XY] = E[X] E[Y]$, then X and Y **may or may not** be independent.

[Link to Video: Independent Random Variables](#)

This page titled [5.1: Joint Distributions of Discrete Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

5.2: Joint Distributions of Continuous Random Variables

Having considered the discrete case, we now look at joint distributions for continuous random variables.

Definition 5.2.1

If continuous random variables X and Y are defined on the same sample space S , then their **joint probability density function (joint pdf)** is a piecewise continuous function, denoted $f(x, y)$, that satisfies the following.

1. $f(x, y) \geq 0$, for all $(x, y) \in \mathbb{R}^2$
2. $\iint_{\mathbb{R}^2} f(x, y) dx dy = 1$
3. $P((X, Y) \in A) = \iint_A f(x, y) dx dy$, for any $A \subseteq \mathbb{R}^2$

The first two conditions in Definition 5.2.1 provide the requirements for a function to be a valid joint pdf. The third condition indicates how to use a joint pdf to calculate probabilities. As an example of applying the third condition in Definition 5.2.1, the **joint cdf** for continuous random variables X and Y is obtained by integrating the joint density function over a set A of the form

$$A = \{(x, y) \in \mathbb{R}^2 \mid X \leq a \text{ and } Y \leq b\},$$

where a and b are constants. Specifically, if A is given as above, then the joint cdf of X and Y , at the point (a, b) , is given by

$$F(a, b) = P(X \leq a \text{ and } Y \leq b) = \int_{-\infty}^b \int_{-\infty}^a f(x, y) dx dy.$$

Note that probabilities for continuous jointly distributed random variables are now *volumes* instead of areas as in the case of a single continuous random variable.

As in the discrete case, we can also obtain the individual, *marginal* pdf's of X and Y from the joint pdf.

Definition 5.2.2

Suppose that continuous random variables X and Y have joint density function $f(x, y)$. The **marginal pdf's** of X and Y are respectively given by the following.

$$f_X(x) = \int_{-\infty}^{\infty} f(x, y) dy \quad (\text{fix a value of } X, \text{ and integrate over all possible values of } Y)$$

$$f_Y(y) = \int_{-\infty}^{\infty} f(x, y) dx \quad (\text{fix a value of } Y, \text{ and integrate over all possible values of } X)$$

[Link to Video: Overview of Definitions 5.2.1 & 5.2.2](#)

Example 5.2.1

Suppose a radioactive particle is contained in a unit square. We can define random variables X and Y to denote the x - and y -coordinates of the particle's location in the unit square, with the bottom left corner placed at the origin. Radioactive particles follow completely random behavior, meaning that the particle's location should be uniformly distributed over the unit square. This implies that the joint density function of X and Y should be constant over the unit square, which we can write as

$$f(x, y) = \begin{cases} c, & \text{if } 0 \leq x \leq 1 \text{ and } 0 \leq y \leq 1 \\ 0, & \text{otherwise,} \end{cases}$$

where c is some unknown constant. We can find the value of c by using the first condition in [Definition 5.2.1](#) and solving the following:

$$\iint_{\mathbb{R}^2} f(x, y) dx dy = 1 \Rightarrow \int_0^1 \int_0^1 c dx dy = 1 \Rightarrow c \int_0^1 \int_0^1 1 dx dy = 1 \Rightarrow c = 1$$

We can now use the joint pdf of X and Y to compute probabilities that the particle is in some specific region of the unit square. For example, consider the region

$$A = \{(x, y) \mid x - y > 0.5\},$$

which is graphed in Figure 1 below.




Figure 1: Graph of region A , shaded in blue.

If we want the probability that the particle's location is in the lower right corner of the unit square that intersects with the region A , then we integrate the joint density function over that portion of A in the unit square, which gives the following probability:

$$P(X - Y > 0.5) = \iint_A f(x, y) dx dy = \int_0^{0.5} \int_{y+0.5}^1 1 dx dy = 0.125$$

Lastly, we apply [Definition 5.2.2](#) to find the marginal pdf's of X and Y .

$$f_X(x) = \int_0^1 1 dy = 1, \quad \text{for } 0 \leq x \leq 1$$

$$f_Y(y) = \int_0^1 1 dx = 1, \quad \text{for } 0 \leq y \leq 1$$

Note that both X and Y are individually uniform random variables, each over the interval $[0, 1]$. This should not be too surprising. Given that the particle's location was uniformly distributed over the unit square, we should expect that the individual coordinates would also be uniformly distributed over the unit intervals.

Example 5.2.2

At a particular gas station, gasoline is stocked in a bulk tank each week. Let random variable X denote the proportion of the tank's capacity that is *stocked* in a given week, and let Y denote the proportion of the tank's capacity that is *sold* in the same week. Note that the gas station cannot sell more than what was stocked in a given week, which implies that the value of Y cannot exceed the value of X . A possible joint pdf of X and Y is given by

$$f(x, y) = \begin{cases} 3x, & \text{if } 0 \leq y \leq x \leq 1 \\ 0, & \text{otherwise.} \end{cases}$$

Note that this function is only nonzero over the triangular region given by $\{(x, y) \mid 0 \leq y \leq x \leq 1\}$, which is graphed in Figure 2 below:




Figure 2: Region over which joint pdf $f(x, y)$ is nonzero.

[Link to Video: Marginal PDFs for Example 5.2.2](#)

We find the joint cdf of X and Y at the point $(x, y) = (1/2, 1/3)$:

$$\begin{aligned} F\left(\frac{1}{2}, \frac{1}{3}\right) &= P\left(X \leq \frac{1}{2} \text{ and } Y \leq \frac{1}{3}\right) = \int_0^{1/3} \int_y^{0.5} 3x \, dx \, dy \\ &= \int_0^{1/3} \left(\frac{3}{2}x^2 \Big|_y^{0.5}\right) dy = \int_0^{1/3} \left(\frac{3}{8} - \frac{3}{2}y^2\right) dy \\ &= \frac{3}{8}y - \frac{1}{2}y^3 \Big|_0^{1/3} \approx 0.1065 \end{aligned}$$

Thus, there is a 10.65% chance that less than half the tank is stocked and less than a third of the tank is sold in a given week. Note that in finding the above integral, we look at where the region given by $\{(x, y) \mid x \leq 1/2, y \leq 1/3\}$ intersects the region over which the joint pdf is nonzero, i.e., the region graphed in Figure 2. This tells us what the limits of integration are in the double integral. Figure 3 below is a graph of the intersection made on [desmos.com](https://www.desmos.com):




Figure 3: Intersection of $\{(x, y) \mid x \leq 1/2, y \leq 1/3\}$ with the region over which joint pdf $f(x, y)$ is nonzero.

Next, we find the probability that the amount of gas sold is less than half the amount that is stocked in a given week. In other words, we find $P(Y < 0.5X)$. In order to find this probability, we need to find the region over which we will integrate the joint pdf. To do this, look for the intersection of the region given by $\{(x, y) \mid y < 0.5x\}$ with the region in Figure 2, which is graphed in Figure 4 below:




Figure 4: Intersection of $\{(x, y) \mid y < 0.5x\}$ with the region over which joint pdf $f(x, y)$ is nonzero.

The calculation is as follows:

$$\begin{aligned}
 P(Y < 0.5X) &= \int_0^1 \int_0^{0.5x} 3x \, dy \, dx \\
 &= \int_0^1 \left(3xy \Big|_0^{0.5x} \right) dx \\
 &= \int_0^1 \left(\frac{3}{2}x^2 - 0 \right) dx = \frac{1}{2}x^3 \Big|_0^1 \\
 &= \frac{1}{2}
 \end{aligned}$$

Thus, there is a 50% chance that the amount of gas sold in a given week is less than half of the gas stocked.

Expectations of Functions of Jointly Distributed Continuous Random Variables

As we did in the discrete case of jointly distributed random variables, we can also look at the expected value of jointly distributed continuous random variables. Again we focus on the expected value of functions applied to the pair (X, Y) , since expected value is defined for a single quantity. At this point, it should not surprise you that the following theorem is similar to [Theorem 5.1.1](#), the result in the discrete setting, except the sums have been replaced by integrals.

Theorem 5.2.1

Suppose that X and Y are jointly distributed continuous random variables with joint pdf $f(x, y)$.

If $g(X, Y)$ is a function of these two random variables, then its expected value is given by the following:

$$E[g(X, Y)] = \iint_{\mathbb{R}^2} g(x, y) f(x, y) \, dx \, dy$$

We will give an example applying Theorem 5.2.1 in an example below.

Independent Random Variables

We can also define independent random variables in the continuous case, just as we did for discrete random variables.

Definition 5.2.3

Continuous random variables X_1, X_2, \dots, X_n are **independent** if the joint pdf factors into a product of the marginal pdf's:

$$f(x_1, x_2, \dots, x_n) = f_{X_1}(x_1) \cdot f_{X_2}(x_2) \cdots f_{X_n}(x_n).$$

It is equivalent to check that this condition holds for the cumulative distribution functions.

Example 5.2.3

Consider the continuous random variables defined in [Example 5.2.1](#), where the X and Y gave the location of a radioactive particle. We will show that X and Y are independent and then verify that [Theorem 5.1.2](#) also applies in the continuous setting.

Recall that we found the marginal pdf's to be the following:

$$\begin{aligned}
 f_X(x) &= 1, \text{ for } 0 \leq x \leq 1 \\
 f_Y(y) &= 1, \text{ for } 0 \leq y \leq 1
 \end{aligned}$$

So, for (x, y) in the unit square, i.e., $0 \leq x \leq 1$ and $0 \leq y \leq 1$, we have

$$f(x, y) = 1 = 1 \cdot 1 = f_X(x) f_Y(y),$$

and outside the unit square, at least one of marginal pdf's will be 0, so

$$f(x, y) = 0 = f_X(x)f_Y(y).$$

We have thus shown that $f(x, y) = f_X(x) f_Y(y)$, for all $(x, y) \in \mathbb{R}^2$, and so by Definition 5.2.3, X and Y are independent.

Now let's look at the expected value of the product of X and Y . To compute this we apply Theorem 5.2.1:

$$\mathbb{E}[XY] = \iint_{\mathbb{R}^2} xy \cdot f(x, y) \, dx dy = \int_0^1 \int_0^1 xy \cdot 1 \, dx dy = \int_0^1 \left(\frac{x^2}{2} y \Big|_0^1 \right) dy = \frac{1}{4}$$

Note that both X and Y are uniform on the interval $[0, 1]$. Therefore, their expected values are both $1/2$, the midpoint of $[0, 1]$. Putting this all together, we have

$$\mathbb{E}[XY] = \frac{1}{4} = \frac{1}{2} \cdot \frac{1}{2} = \mathbb{E}[X] \mathbb{E}[Y],$$

which is the conclusion to [Theorem 5.1.2](#).

[Link to Video: Independent Continuous Random Variables](#)

This page titled [5.2: Joint Distributions of Continuous Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

5.3: Conditional Probability Distributions

In this section, we consider the probability distribution of one random variable *given* information about the value of another random variable. As we will see in the formal definition, this kind of conditional distribution will involve the joint distribution of the two random variables under consideration, which we introduced in the previous two sections. We begin with discrete random variables, and then consider the continuous case.

Conditional Distributions of Discrete Random Variables

Recall the definition of conditional probability for events ([Definition 2.2.1](#)): the conditional probability of A given B is equal to

$$P(A | B) = \frac{P(A \cap B)}{P(B)}.$$

We use this same concept for events to define conditional probabilities for *random variables*.

Definition 5.3.1

If X and Y are discrete random variables with joint pmf given by $p(x, y)$, then the **conditional probability mass function of X , given that $Y = y$** , is denoted $p_{X|Y}(x|y)$ and given by

$$p_{X|Y}(x|y) = \frac{P(\{X = x\} \cap \{Y = y\})}{P(Y = y)} = \frac{p(x, y)}{p_Y(y)}, \quad \text{provided that } p_Y(y) > 0.$$

Note that if $p_Y(y) = 0$, then for that value of Y the conditional pmf of X does not exist.

Similarly, the **conditional probability mass function of Y , given that $X = x$** , is denoted $p_{Y|X}(y|x)$ and given by

$$p_{Y|X}(y|x) = \frac{P(\{Y = y\} \cap \{X = x\})}{P(X = x)} = \frac{p(x, y)}{p_X(x)}, \quad \text{provided that } p_X(x) > 0.$$

[Link to Video: Overview of Definition 5.3.1](#)

Example 5.3.1

For an example of conditional distributions for discrete random variables, we return to the context of [Example 5.1.1](#), where the underlying probability experiment was to flip a fair coin three times, and the random variable X denoted the number of heads obtained and the random variable Y denoted the winnings when betting on the placement of the first heads obtained in the three flips. We found the joint pmf for X and Y in [Table 1](#) of Section 5.1, and the marginal pmf's are given in [Table 2](#). We now find the conditional distributions of X and Y .

First, to find the conditional distribution of X given a value of Y , we can think of fixing a row in [Table 1](#) and dividing the values of the joint pmf in that row by the marginal pmf of Y for the corresponding value. For example, to find $p_{X|Y}(x|1)$, we divide each entry in the $Y = 1$ row by $p_Y(1) = 1/2$. Doing this for each row in [Table 1](#), results in the conditional distributions of X given a value of Y , which we represent in the following table.

Conditional pmf's of X given Y

x	$p_{X Y}(x -1)$	$p_{X Y}(x 1)$	$p_{X Y}(x 2)$	$p_{X Y}(x 3)$
0	1	0	0	0
1	0	1/4	1/2	1
2	0	1/2	1/2	0

x	$p_{X Y}(x -1)$	$p_{X Y}(x 1)$	$p_{X Y}(x 2)$	$p_{X Y}(x 3)$
3	0	1/4	0	0

Note that every column in the above table sums to 1. This is because each row is a different probability mass function for X given a value of Y . Specifically, if we look at the column for the conditional distribution of X given that $Y = 1$, $p_{X|Y}(x|1)$, this is the distribution of probability for the number of heads obtained, knowing that the winnings of the game are \$1. Recall that we win \$1 if the first heads is on the first toss. In that case, we know the outcome of the probability experiment must be one of $\{hhh, hht, hth, htt\}$. Working with this reduced sample space, we can see how the corresponding probabilities for the values of X arise.

Similar to the process for the conditional pmf's of X given Y , we can find the conditional pmf's of Y given X . This time, though, we take a column in Table 1, giving the joint pmf for a fixed value of X , and divide by the marginal pmf of X for the corresponding value. The following table gives the results.

Conditional pmf's of Y given X

y	$p_{Y X}(y 0)$	$p_{Y X}(y 1)$	$p_{Y X}(y 2)$	$p_{Y X}(y 3)$
-1	1	0	0	0
1	0	1/3	2/3	1
2	0	1/3	1/3	0
3	0	1/3	0	0

Note again that every column in the above table sums to 1. A conditional pmf is a pmf, just found in a specific way.

[Link to Video: Walkthrough of Example 5.3.1](#)

Informally, we can think of a conditional probability distribution as a probability distribution for a *sub-population*. In other words, a conditional probability distribution describes the probability that a randomly selected person from a sub-population has a given characteristic of interest. In this context, the joint probability distribution is the probability that a randomly selected person from the *entire* population has *both* characteristics of interest. The following example demonstrates these interpretations in a specific context.

Example 5.3.2

Suppose we are interested in the relationship between an individual's hair and eye color. Based on a random sample of Saint Mary's students, we have the following joint pmf, with marginal pmf's given in the margins:

$p(x, y)$	Hair Color (X)				
Eye Color (Y)	blonde (1)	red (2)	brown (3)	black (4)	$p_Y(y)$
blue (1)	0.12	0.05	0.12	0.01	0.30
green (2)	0.12	0.07	0.09	0	0.28
brown (3)	0.16	0.07	0.16	0.03	0.42
$p_X(x)$	0.40	0.19	0.37	0.04	1.00

The probabilities in the last row and column (orange cells) give the marginal pmf's for X and Y , while the probabilities in the interior (white and grey cells) give the joint pmf for pairs (X, Y) . For instance, $p(3, 2) = 0.09$ indicates the joint probability that a randomly selected SMC student has brown hair ($X = 3$) and green eyes ($Y = 2$) is 9%, $p_X(3) = 0.37$ indicates the marginal probability that a randomly selected SMC student has brown hair is 37%, and $p_Y(2) = 0.28$ indicates the marginal probability that a randomly selected SMC student has green eyes is 28%.

Given this table of probabilities, we can calculate conditional pmf values:

- First, let's find $p_{X|Y}(2|1)$:

$$p_{X|Y}(2|1) = \frac{p(2,1)}{p_Y(1)} = \frac{0.05}{0.3} = \frac{1}{6} \approx 0.167$$

Note that we can write $p_{X|Y}(2|1) = P(X=2 | Y=1) = P(\text{red hair} | \text{blue eyes})$. Here we are finding the probability that an individual in the sub-population of individuals with blue eyes has red hair. Specifically, we found that 1/6 (or approximately 16.7%) of SMC students with blue eyes have red hair.

- Now, let's reverse the order of X and Y , and find $p_{Y|X}(2|1)$:

$$p_{Y|X}(2|1) = \frac{p(1,2)}{p_X(1)} = \frac{0.12}{0.40} = 0.3$$

Now the sub-population is individuals with blonde hair, and we find the probability that an individual in this sub-population has green eyes. Specifically, we found that 30% of SMC students with blonde hair have green eyes.

[Link to Video: Interpretation of Joint & Conditional Probabilities \(Example 5.3.2\)](#)

Properties of Conditional PMF's

1. Conditional pmf's are valid pmf's. In other words, the conditional pmf for X , given $Y = y$, for a fixed y , is a valid pmf satisfying the following:

$$0 \leq p_{X|Y}(x|y) \leq 1 \quad \text{and} \quad \sum_x p_{X|Y}(x|y) = 1$$

Similarly, for a fixed x , we also have the following for the conditional pmf of Y , given $X = x$:

$$0 \leq p_{Y|X}(y|x) \leq 1 \quad \text{and} \quad \sum_y p_{Y|X}(y|x) = 1$$

2. In general, the conditional distribution of X given Y does not equal the conditional distribution of Y given X , i.e.,

$$p_{X|Y}(x|y) \neq p_{Y|X}(y|x).$$

3. If X and Y are independent, discrete random variables, then the following are true:

$$p_{X|Y}(x|y) = p_X(x)$$

$$p_{Y|X}(y|x) = p_Y(y)$$

In other words, if X and Y are independent, then knowing the value of one random variable does not affect the probability of the other one.

Now that we have defined conditional distributions, we define conditional expectation.

Definition 5.3.2

For discrete random variables X and Y , the **conditional expected value of X , given $Y = y$** , is given by

$$\mu_{X|Y=y} = E[X|Y = y] = \sum_x x p_{X|Y}(x|y),$$

and the **conditional expected value of Y , given $X = x$** , is given by

$$\mu_{Y|X=x} = E[Y|X = x] = \sum_y y p_{Y|X}(y|x).$$

Similarly, we can define conditional variances. The **conditional variance of X , given $Y = y$** , is given by

$$\begin{aligned} \sigma_{X|Y=y}^2 &= \text{Var}(X|Y = y) = E[(X - \mu_{X|Y=y})^2 | Y = y] = \sum_x (x - \mu_{X|Y=y})^2 p_{X|Y}(x|y) \\ &= E[X^2 | Y = y] - \mu_{X|Y=y}^2 = \left(\sum_x x^2 p_{X|Y}(x|y) \right) - \mu_{X|Y=y}^2 \end{aligned}$$

Example 5.3.3

Continuing in the context of [Example 5.3.2](#), we calculate the conditional mean and variance of hair color (X), given that a randomly selected student has brown eyes ($Y = 3$).

First, we derive the conditional pmf of X , given $Y = 3$, by taking the row for brown eyes in the joint pmf table and dividing each by the marginal pmf for Y at 3, i.e., $p_Y(3) = 0.42$. This gives the following:

x	$p_{X Y}(x 3) = p(x, 3)/p_Y(3)$
1	$0.16/0.42 = 8/21$
2	$0.07/0.42 = 1/6$
3	$0.16/0.42 = 8/21$
4	$0.03/0.42 = 1/14$

Using the conditional probabilities in the table above, we calculate the following:

$$\begin{aligned} E[X|Y = 3] &= \sum_{x=1}^4 x \cdot p_{X|Y}(x|3) = (1) \left(\frac{8}{21} \right) + (2) \left(\frac{1}{6} \right) + (3) \left(\frac{8}{21} \right) + (4) \left(\frac{1}{14} \right) = \frac{15}{7} \approx 2 \\ E[X^2|Y = 3] &= \sum_{x=1}^4 x^2 \cdot p_{X|Y}(x|3) = (1^2) \left(\frac{8}{21} \right) + (2^2) \left(\frac{1}{6} \right) + (3^2) \left(\frac{8}{21} \right) + (4^2) \left(\frac{1}{14} \right) = \frac{118}{21} \\ \Rightarrow \text{Var}(X|Y = 3) &= E[X^2|Y = 3] - (E[X|Y = 3])^2 = \frac{118}{21} - \left(\frac{15}{7} \right)^2 \approx 1 \end{aligned}$$

Thus, the expected hair color of a student with brown eyes is red.

Conditional Distributions of Continuous Random Variables

We now turn to the continuous setting. Note that definitions and results in the discrete setting transfer to the continuous setting by simply replacing sums with integrals and pmf's with pdf's. The following definition gives the formulas for conditional distributions and expectations of continuous random variables.

Definition 5.3.3

If X and Y are continuous random variables with joint pdf given by $f(x, y)$, then the **conditional probability density function (pdf) of X , given that $Y = y$** , is denoted $f_{X|Y}(x|y)$ and given by

$$f_{X|Y}(x|y) = \frac{f(x, y)}{f_Y(y)}.$$

The **conditional expected value of X , given $Y = y$** , is

$$E[X|Y = y] = \int_{\mathbb{R}} x f_{X|Y}(x|y) dx$$

and the **conditional variance of X , given $Y = y$** , is

$$\text{Var}(X|Y = y) = E[X^2 | Y = y] - (E[X|Y = y])^2.$$

Similarly, we can define the conditional pdf, expected value, and variance of Y , given $X = x$, by swapping the roles of X and Y in the above.

Properties of Conditional PDF's

1. Conditional pdf's are valid pdf's. In other words, the conditional pdf for X , given $Y = y$, for a fixed y , is a valid pdf satisfying the following:

$$0 \leq f_{X|Y}(x|y) \quad \text{and} \quad \int_{\mathbb{R}} f_{X|Y}(x|y) dx = 1$$

2. In general, the conditional distribution of X given Y does not equal the conditional distribution of Y given X , i.e.,

$$f_{X|Y}(x|y) \neq f_{Y|X}(y|x).$$

3. If X and Y are independent, discrete random variables, then the following are true:

$$\begin{aligned} f_{X|Y}(x|y) &= f_X(x) \\ f_{Y|X}(y|x) &= f_Y(y) \end{aligned}$$

In other words, if X and Y are independent, then knowing the value of one random variable does not affect the probability of the other one.

[Link to Video: Overview of Definition 5.3.3 & Properties of Conditional PDF's](#)

Example 5.3.4

We verify the third property of conditional pdf's for the radioactive particle example ([Example 5.2.1](#)). Recall that we had the following joint pdf and marginal pdf's for X and Y :

$$\begin{aligned}f(x, y) &= 1, & \text{for } 0 \leq x \leq 1 \text{ and } 0 \leq y \leq 1 \\f_X(x) &= 1, & \text{for } 0 \leq x \leq 1 \\f_Y(y) &= 1, & \text{for } 0 \leq y \leq 1\end{aligned}$$

We showed in [Example 5.2.3](#) that X and Y are independent. So, if we fix $y \in [0, 1]$, the following shows that the conditional pdf of X , given $Y = y$, is equal to the marginal pdf of X , as stated in the third property of conditional pdf's above:

$$f_{X|Y}(x|y) = \frac{f(x, y)}{f_Y(y)} = \frac{1}{1} = f_X(x)$$

Example 5.3.5

Continuing in the context of [Example 5.2.2](#), where X gave the amount of gas stocked and Y gave the amount of gas sold at a given gas station in a given week, we find the conditional pdf of Y given that $X = 0.5$. In other words, we find the conditional probability distribution for the amount of gas sold in a given week, when only half of the tank was stocked.

First, we find the marginal pdf for X :

$$f_X(x) = \int_{\mathbb{R}} f(x, y) dy = \int_0^x 3x dy = 3xy \Big|_0^x = 3x^2, \quad \text{for } 0 \leq x \leq 1.$$

So, if $X = 0.5$, then $f_X(0.5) = 3(0.5)^2 = 0.75$, and the conditional pdf of Y in this case is

$$f_{Y|X}(y|0.5) = \frac{f(0.5, y)}{f_X(0.5)} = \frac{3(0.5)}{0.75} = 2, \quad \text{for } 0 \leq y \leq 0.5.$$

Note the $f_{Y|X}(y|0.5)$ is the pdf for a uniform distribution on the interval $[0, 0.5]$. Thus, the conditional distribution of the amount of gas sold in a week, given that only half of the tank is stocked, is uniformly distributed between 0 and 0.5. Recognizing this, we can easily compute the conditional expected value of Y , given that $X = 0.5$:

$$E[Y|X = 0.5] = \frac{0.5}{2} = 0.25.$$

In other words, given that 50% of tank is stocked, we expect that 25% will be sold.

[Link to Video: Walkthrough of Example 5.3.5](#)

This page titled [5.3: Conditional Probability Distributions](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

5.4: Finding Distributions of Functions of Continuous Random Variables

Single Variable Case

In addition to considering the probability distributions of random variables simultaneously using joint distribution functions, there is also occasion to consider the probability distribution of *functions* applied to random variables as the following example demonstrates.

Example 5.4.1

Consider again the context of [Example 5.2.2](#), where the random variable X represented the amount of gas stocked in a given week at a specific gas station. In [Example 5.3.5](#), we showed that the pdf of X is given by

$$f_X(x) = \begin{cases} 3x^2, & \text{for } 0 \leq x \leq 1 \\ 0, & \text{otherwise.} \end{cases}$$

Suppose that the storage tank at the gas station holds 10,000 gallons, then $10,000X$ gives the number of gallons stocked in a week. If the current price of gas is \$3/gallon and there are fixed delivery costs of \$2000, then the total cost to stock 10,000 X gallons in a given week is given by the following

$$3(10,000X) + 2000 \Rightarrow C = 30,000X + 2000, \quad (5.4.1)$$

where the random variable C denotes the total cost of delivery.

One approach to finding the probability distribution of a function of a random variable relies on the relationship between the pdf and cdf for a continuous random variable:

$$\frac{d}{dx}[F(x)] = f(x) \quad \text{''derivative of cdf = pdf''}$$

As we will see in the following examples, it is often easier to find the cdf of a function of a continuous random variable, and then use the above relationship to derive the pdf.

Example 5.4.2

Continuing with [Example 5.4.1](#), we find the pdf of the total cost, C . First, we derive the cdf for X . If we let $0 \leq x \leq 1$, i.e., select a value of x where the pdf of X is nonzero, then we have

$$F_X(x) = P(X \leq x) = \int_{-\infty}^x f_X(t) dt = \int_0^x 3t^2 dt = t^3 \Big|_0^x = x^3.$$

For any $x < 0$, the cdf of X is necessarily 0, since X cannot be negative (we cannot stock a negative proportion of the tank). And for any $x > 1$, the cdf of X is necessarily equal to 1, since the proportion of gas stocked will always be less than or equal to 100% of the tank's capacity. Putting it all together, we have the cdf of X given as follows:

$$F_X(x) = \begin{cases} 0, & x < 0 \\ x^3, & 0 \leq x \leq 1 \\ 1, & 1 < x \end{cases} \quad (5.4.2)$$

We can now use the cdf of X to find the **cdf of C** . Let c denote a possible value for the random variable C . We relate the cdf of C to the cdf of X by substituting the expression of C in terms of X given in Equation 5.4.1, and then solving for X , as follows:

$$F_C(c) = P(C \leq c) = P(30000X + 2000 \leq c) = P\left(X \leq \frac{c-2000}{30000}\right) = F_X\left(\frac{c-2000}{30000}\right)$$

If $0 \leq (c-2000)/30000 \leq 1 \Rightarrow 2000 \leq c \leq 32000$, then, using the formula for the cdf of X we derived in Equation 5.4.2, we find that the cdf of C is

$$F_C(c) = \left(\frac{c-2000}{30000}\right)^3, \quad \text{for } 2000 \leq c \leq 32000. \quad (5.4.3)$$

Note that $2000 \leq c \leq 32000$ gives the possible values of C . C cannot be less than \$2,000 because of the fixed delivery costs, and C cannot be more than \$32,000, which is the cost of stocking the entire tank.

We now find the **pdf of C** by taking the derivative of the cdf in Equation 5.4.3:

$$f_C(c) = \frac{d}{dc}[F_C(c)] = \frac{d}{dc}\left[\left(\frac{c-2000}{30000}\right)^3\right] = 3\left(\frac{c-2000}{30000}\right)^2 \times \left(\frac{1}{30000}\right), \quad \text{for } 2000 \leq c \leq 32000$$

Video: Motivating Example (Walkthrough of Examples 5.4.1 & 5.4.2)

Example 5.4.2 demonstrates the general strategy to finding the probability distribution of a function of a random variable: we first find the cdf of the random variable *in terms of* the random variable it is a function of (assuming we know the cdf of that random variable), then we differentiate to find the pdf. Let's look at another example before formalizing the strategy.

Example 5.4.3

Let X be uniform on $[0, 1]$. We find the pdf of $Y = X^2$.

Recall that the pdf for a uniform $[0, 1]$ random variable is $f_X(x) = 1$, for $0 \leq x \leq 1$, and that the cdf is

$$F_X(x) = \begin{cases} 0, & x < 0 \\ x, & 0 \leq x \leq 1 \\ 1, & x > 1 \end{cases}$$

Note that the possible values of Y are $0 \leq y \leq 1$.

First, we find the **cdf of Y** , for $0 \leq y \leq 1$:

$$F_Y(y) = P(Y \leq y) = P(X^2 \leq y) = P(X \leq \sqrt{y}) = F_X(\sqrt{y}) = \sqrt{y}, \quad \text{for } 0 \leq y \leq 1$$

Now we differentiate the cdf to get the **pdf of Y** :

$$f_Y(y) = \frac{d}{dy}[F_Y(y)] = \frac{d}{dy}[\sqrt{y}] = \frac{1}{2}y^{-\frac{1}{2}}, \quad \text{for } 0 \leq y \leq 1$$

The following formalizes the strategy we took to find the cdf's in the previous examples.

Change-of-Variable Technique

Suppose that X is a continuous random variable with pdf $f_X(x)$, which is nonzero on interval I .

Further suppose that g is a differentiable function that is strictly monotonic on I .

Then the pdf of $Y = g(X)$ is given by

$$f_Y(y) = f_X(g^{-1}(y)) \times \left| \frac{d}{dy} [g^{-1}(y)] \right|, \quad \text{for } x \in I \text{ and } y = g(x). \quad (5.4.4)$$

Video: Change-of-Variable Technique & Walkthrough of Example 5.4.5

Returning to [Example 5.4.2](#), we demonstrate the Change-of-Variable technique.

Example 5.4.4

Recall that the random variable X denotes the amount of gas stocked. We now let Y denote the total cost of delivery for the gas stocked, i.e., $Y = 30000X + 2000$. We know that the pdf of X is given by

$$f_X(x) = 3x^2, \text{ for } 0 \leq x \leq 1,$$

so that $I = [0, 1]$, using the notation of the Change-of-Variable technique. We also have that $y = g(x) = 30000x + 2000$, and note that g is increasing on the interval I . Thus, the Change-of-Variable technique can be applied, and so we find the inverse of g and its derivative:

$$x = g^{-1}(y) = \frac{y - 2000}{30000} \quad \text{and} \quad \frac{d}{dy} [g^{-1}(y)] = \frac{d}{dy} \left[\frac{y - 2000}{30000} \right] = \frac{1}{30000}$$

So, applying Change-of-Variable formula given in Equation 5.4.4, we get

$$f_Y(y) = f_X(g^{-1}(y)) \times \frac{d}{dy} [g^{-1}(y)] = 3(g^{-1}(y))^2 \times \frac{1}{30000} = 3 \left(\frac{y - 2000}{30000} \right)^2 \times \frac{1}{30000},$$

which matches the result found in [Example 5.4.2](#).

Informally, the Change-of-Variable technique can be restated as follows.

1. Find the cdf of $Y = g(X)$ in terms of the cdf for X using the inverse of g , i.e., isolate X .
2. Take the derivative of the cdf of Y to get the pdf of Y using the chain rule. An absolute value is needed if g is decreasing.

The advantage of the Change-of-Variable technique is that we do not have to find the cdf of X in order to find the pdf of Y , as the next example demonstrates.

Example 5.4.5

Let the random variable X have pdf given by $f_X(x) = 5x^4$, for $0 \leq x \leq 1$, i.e., $I = [0, 1]$. Also, let $g(x) = 1 - x^2$. To find the pdf for $Y = g(X) = 1 - X^2$, we first find g^{-1} :

$$\begin{aligned}y &= g(x) \\y &= 1 - x^2 \\y - 1 &= -x^2 \\ \Rightarrow g^{-1}(y) &= \sqrt{1 - y} = x\end{aligned}$$

Now we can find the derivative of g^{-1} :

$$\frac{d}{dy}[g^{-1}(y)] = \frac{d}{dy}[\sqrt{1-y}] = -\frac{1}{2}(1-y)^{-\frac{1}{2}}$$

Applying the Change-of-Variable formula we find the pdf of Y :

$$f_Y(y) = f_X(g^{-1}(y)) \times \left| \frac{d}{dy}[g^{-1}(y)] \right| = 5(\sqrt{1-y})^4 \times \left| -\frac{1}{2}(1-y)^{-\frac{1}{2}} \right| = \frac{5}{2}(\sqrt{1-y})^3, \quad \text{for } 0 \leq y \leq 1.$$

The Change-of-Variable technique requires that a monotonic function g is applied. However, if that is not the case, we can just consider the monotonic pieces separately, as in the next example.

Example 5.4.6

Let X be uniform on $[-1, 1]$. Then the pdf of X is

$$f_X(x) = \frac{1}{2}, \text{ for } -1 \leq x \leq 1,$$

and the cdf of X is

$$F_X(x) = \begin{cases} 0, & \text{for } x < -1 \\ \frac{x+1}{2}, & \text{for } -1 \leq x \leq 1 \\ 1, & \text{for } x > 1 \end{cases}$$

We find the pdf for $Y = X^2$.

Following the general strategy, we first find the cdf of Y in terms of X :

$$\begin{aligned}F_Y(y) &= P(Y \leq y) = P(X^2 \leq y) = 0, \quad \text{if } y < 0 \\ \text{and } F_Y(y) &= P(-\sqrt{y} \leq X \leq \sqrt{y}), \quad \text{if } y \geq 0 \\ \Rightarrow F_Y(y) &= F_X(\sqrt{y}) - F_X(-\sqrt{y}) = \frac{\sqrt{y}+1}{2} - \frac{-\sqrt{y}+1}{2} = \sqrt{y}, \quad \text{if } 0 \leq \sqrt{y} \leq 1,\end{aligned}$$

which gives

$$F_Y(y) = \begin{cases} 0 & \text{if } y < 0 \\ \sqrt{y}, & \text{if } 0 \leq y \leq 1 \\ 1, & \text{if } y > 1 \end{cases}$$

Now we take the derivative of the cdf to get the pdf of Y :

$$f_Y(y) = \frac{d}{dy}[F_Y(y)] = \begin{cases} \frac{d}{dy}[0] = 0, & \text{if } y < 0 \\ \frac{d}{dy}[\sqrt{y}] = \frac{1}{2\sqrt{y}}, & \text{if } 0 \leq y \leq 1 \\ \frac{d}{dy}[1] = 0, & \text{if } y > 1 \end{cases}$$

In summary, the pdf of $Y = X^2$ is given by

$$f_Y(y) = \begin{cases} \frac{1}{2\sqrt{y}}, & \text{if } 0 \leq y \leq 1 \\ 0, & \text{otherwise} \end{cases}$$

[Video: Bonus Example](#)

Special Case: Normal Distributions

Functions of normally distributed random variables are of particular importance in statistics. In the next example, we derive the probability distribution of the square of a standard normal random variable.

Example 5.4.7

Let Z be a standard normal random variable, i.e., $Z \sim N(0, 1)$. We find the pdf of $Y = Z^2$.

Let Φ denote the cdf of Z , i.e., $\Phi(z) = P(Z \leq z) = F_Z(z)$. We first find the cdf of $Y = Z^2$ in terms of Φ (recall that there is no closed form expression for Φ):

$$\begin{aligned} F_Y(y) &= P(Y \leq y) = P(Z^2 \leq y) \\ &= P(-\sqrt{y} \leq Z \leq \sqrt{y}), \text{ for } y \geq 0 \\ &= \Phi(\sqrt{y}) - \Phi(-\sqrt{y}) \end{aligned}$$

Note that if $y < 0$, then $F_Y(y) = 0$, since it is not possible for $Y = Z^2$ to be negative. In other words, the possible values of $Y = Z^2$ are $y \geq 0$.

Next, we take the derivative of the cdf of Y to find its pdf. Before doing so, we note that if Φ is the cdf for Z , then its derivative is the pdf for Z , which is denoted φ . Since Z is a standard normal random variable, we know that

$$\varphi(z) = f_Z(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}, \quad \text{for } z \in \mathbb{R}.$$

Using this, we now find the pdf of Y :

$$\begin{aligned}
 f_Y(y) &= \frac{d}{dy} [F_Y(y)] = \frac{d}{dy} [\Phi(\sqrt{y}) - \Phi(-\sqrt{y})] \\
 &= \frac{d}{dy} [\Phi(\sqrt{y})] - \frac{d}{dy} [\Phi(-\sqrt{y})] \\
 &= \varphi(\sqrt{y}) \cdot \frac{1}{2\sqrt{y}} + \varphi(-\sqrt{y}) \cdot \frac{1}{2\sqrt{y}} \\
 &= \frac{1}{\sqrt{2\pi}} e^{-(\sqrt{y})^2/2} \cdot \frac{1}{2\sqrt{y}} + \frac{1}{\sqrt{2\pi}} e^{-(-\sqrt{y})^2/2} \cdot \frac{1}{2\sqrt{y}} \\
 &= \frac{1}{\sqrt{2\pi}} e^{-y/2} \cdot \frac{1}{\sqrt{y}}.
 \end{aligned}$$

In summary, if $Y = Z^2$, where $Z \sim N(0, 1)$, then the pdf for Y is given by

$$f_Y(y) = \frac{y^{-1/2}}{\sqrt{2\pi}} e^{-y/2}, \text{ for } y \geq 0.$$

Note that the pdf for Y is a *gamma* pdf with $\alpha = \lambda = \frac{1}{2}$. This is also referred to as the **chi-square distribution**, denoted χ^2 . See [below](#) for a video walkthrough of this example.

Multivariate Case


We can extend the Change-of-Variable technique to the multivariate case as well.

Example 5.4.8

Suppose that X and Y have joint pdf given by

$$f(x, y) = e^{-x}, \quad \text{for } 0 \leq y \leq x < \infty.$$

The following figure shows the region (shaded in blue) over which the joint pdf of X and Y is nonzero.



Let's define the random variable $W = g(X, Y) = X - Y$, and find the pdf of W .

First, we find the cdf of W ,

$$F_W(w) = P(W \leq w) = P(X - Y \leq w),$$

for which we need to find where the region given by $\{(x, y) | x - y \leq w\}$ intersects the region over which the joint pdf $f(x, y)$ is nonzero. Note that if $w < 0$, then $P(X - Y \leq w) = 0$, since there is no intersection with where the joint pdf $f(x, y)$ is nonzero. If $w \geq 0$, then

$$\begin{aligned} P(X - 1 \leq w) &= \int_0^\infty \int_y^{y+w} e^{-x} dx dy = \int_0^\infty -e^{-x} \Big|_y^{y+w} dy \\ &= \int_0^\infty (-e^{-(y+w)} + e^{-y}) dy \\ &= \int_0^\infty (-e^{-y} e^{-w} + e^{-1}) dy \\ &= \int_0^\infty e^{-y} (-e^{-w} + 1) dy \\ &= (-e^{-w})(-e^{-y}) \Big|_0^\infty \\ &= 1 - e^{-w} \end{aligned}$$

In summary, we have found the following

$$F_W(w) = \begin{cases} 0, & \text{if } w < 0 \\ 1 - e^{-w}, & \text{if } w \geq 0 \end{cases}$$

We then take derivative of cdf to find pdf:

$$f_W(w) = \frac{d}{dw} [F_W(w)] = \begin{cases} 0, & \text{if } w < 0 \\ \frac{d}{dw} [1 - e^{-w}] = e^{-w}, & \text{if } w \geq 0 \end{cases}$$

Note: W is exponential with $\lambda = 1$.

[Video: Walkthrough of Example 5.4.8](#)

Using Moment-Generating Functions

There is another approach to finding the probability distribution of functions of random variables, which involves moment-generating functions. Recall the following properties of mgf's.

Theorem 3.8.4

The mgf $M_X(t)$ of random variable X uniquely determines the probability distribution of X . In other words, if random variables X and Y have the same mgf, $M_X(t) = M_Y(t)$, then X and Y have the same probability distribution.

Theorem 3.8.2

Let X be a random variable with mgf $M_X(t)$, and let a, b be constants. If random variable $Y = aX + b$, then the mgf of Y is given by

$$M_Y(t) = e^{bt} M_X(at).$$

Theorem 3.8.3

If X_1, \dots, X_n are independent random variables with mgf's $M_{X_1}(t), \dots, M_{X_n}(t)$, respectively, then the mgf of random variable $Y = X_1 + \dots + X_n$ is given by

$$M_Y(t) = M_{X_1}(t) \cdots M_{X_n}(t).$$

[Theorem 3.8.4](#) states that mgf's are unique, and [Theorems 3.8.2 & 3.8.3](#) combined provide a process for finding the mgf of a linear combination of random variables. All three theorems provide a *Moment-Generating-Function technique* for finding the probability distribution of a function of random variable(s), which we demonstrate with the following examples involving the normal distribution.

Example 5.4.9

Suppose that $X \sim N(\mu, \sigma)$. It can be shown that the mgf of X is given by

$$M_X(t) = e^{\mu t + (\sigma^2 t^2 / 2)}, \quad \text{for } t \in \mathbb{R}.$$

Using this mgf formula, we can show that $Z = \frac{X - \mu}{\sigma}$ has the standard normal distribution.

1. Note that if $Z \sim N(0, 1)$, then the mgf is $M_Z(t) = e^{0t + (1^2 t^2 / 2)} = e^{t^2 / 2}$
2. Also note that $\frac{X - \mu}{\sigma} = \left(\frac{1}{\sigma}\right) X + \left(\frac{-\mu}{\sigma}\right)$, so by [Theorem 3.8.2](#),

$$M_{\frac{1}{\sigma} X - \frac{\mu}{\sigma}}(t) = e^{-\frac{\mu t}{\sigma}} M_X\left(\frac{t}{\sigma}\right) = e^{t^2 / 2}.$$

Thus, we have shown that Z and $\frac{X - \mu}{\sigma}$ have the same mgf, which by [Theorem 3.8.4](#), says that they have the same distribution.

Now suppose X_1, \dots, X_n are each independent normally distributed with means μ_1, \dots, μ_n and sd's $\sigma_1, \dots, \sigma_n$, respectively.

Let's find the probability distribution of the sum $Y = a_1 X_1 + \dots + a_n X_n$ (a_1, \dots, a_n constants) using the mgf technique:

By [Theorem 3.8.2](#), we have

$$M_{a_i X_i}(t) = M_{X_i}(a_i t) = e^{\mu_i a_i t + (\sigma_i^2 (a_i)^2 t^2 / 2)}, \quad \text{for } i = 1, \dots, n,$$

and then by [Theorem 3.8.3](#) we get the following:

$$\begin{aligned} M_Y(t) &= M_{a_1 X_1}(t) \cdot M_{a_2 X_2}(t) \cdots M_{a_n X_n}(t) \\ &= e^{\mu_1 a_1 t + \sigma_1^2 a_1^2 t^2 / 2} e^{\mu_2 a_2 t + \sigma_2^2 a_2^2 t^2 / 2} \cdots e^{\mu_n a_n t + \sigma_n^2 a_n^2 t^2 / 2} \\ &= e^{(\mu_1 a_1 + \mu_2 a_2 + \cdots + \mu_n a_n) t + (\sigma_1^2 a_1^2 + \sigma_2^2 a_2^2 + \cdots + \sigma_n^2 a_n^2) \frac{t^2}{2}} \\ &\Rightarrow M_Y(t) = e^{\mu_y t + \sigma_y^2 t^2 / 2} \end{aligned}$$

Thus, by [Theorem 3.8.4](#), $Y \sim N(\mu_y, \sigma_y)$.

The second part of Example 5.4.9 proved the following.

Sums of Independent Normal Random Variables

If X_1, \dots, X_n are mutually independent normal random variables with means μ_1, \dots, μ_n and standard deviations $\sigma_1, \dots, \sigma_n$, respectively, then the linear combination

$$Y = a_1X_1 + \dots + c_nX_n = \sum_{i=1}^n a_iX_i,$$

is normally distributed with the following mean and variance:

$$\mu_Y = a_1\mu_1 + \dots + a_n\mu_n = \sum_{i=1}^n a_i\mu_i \quad \sigma_Y^2 = a_1^2\sigma_1^2 + \dots + a_n^2\sigma_n^2 = \sum_{i=1}^n a_i^2\sigma_i^2$$

Video: Functions of Normal Random Variables

This page titled [5.4: Finding Distributions of Functions of Continuous Random Variables](#) is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

5.5: Sample Mean

Suppose we are interested in understanding the mean of some population of values, but do not have full information about the entire population. One approach to solving this problem is to obtain a random sample of a subset of values from the population and consider the mean of the sample. The mean of the sample is referred to as the *sample mean*. Since the sample is randomly selected, the sample mean may be thought of as a function applied to a collection of random variables.

Example 5.5.1

Suppose we want to know the average SAT math score for girls in Indiana. We could randomly select seniors from high schools in the South Bend School Corporation as a *sample* from all IN girls, and use the mean SAT math score for the South Bend girls as an estimate of the overall mean for IN girls.

The mean of SB girls depends on which sample we randomly select, therefore the *sample mean is a random variable*.

The probability distribution of the sample mean is referred to as the **sampling distribution** of the sample mean. The following result, which is a corollary to [Sums of Independent Normal Random Variables](#), indicates how to find the sampling distribution when the population of values follows a normal distribution.

Corollary 5.5.1

If X_1, \dots, X_n represent the values of a random sample from a $N(\mu, \sigma)$ population, then the *sample mean*

$$\bar{X} = \frac{X_1 + \dots + X_n}{n} = \sum_{i=1}^n \frac{1}{n} X_i,$$

is normally distributed with mean μ and standard deviation σ/\sqrt{n} . In other words, we can write

$$\bar{X} \sim N(\mu, \sigma/\sqrt{n}).$$

Proof

1. Sample observations are independent when randomly selected. Furthermore, each observation has same distribution as population. X_1, \dots, X_n represent the observations in the random sample $\implies X_1, \dots, X_n$ are independent and each $X_i \sim N(\mu, \sigma)$
2. \bar{X} is the sum of independent normally distributed random variables:

$$\bar{X} = \frac{X_1 + \dots + X_n}{n} = \frac{X_1}{n} + \frac{X_2}{n} + \dots + \frac{X_n}{n} = \sum_{i=1}^n \frac{1}{n} X_i \implies a_i = \frac{1}{n}, \text{ for } i = 1, \dots, n$$

3. By [Sums of Independent Normal Random Variables](#): $\bar{X} \sim N(\mu_{\bar{X}}, \sigma_{\bar{X}})$, where

$$\begin{aligned} \mu_{\bar{X}} &= \sum_{i=1}^n \frac{1}{n} \mu = \frac{1}{n} \mu + \dots + \frac{1}{n} \mu = n \frac{1}{n} \mu = \mu \\ \sigma_{\bar{X}}^2 &= \sum_{i=1}^n \left(\frac{1}{n}\right)^2 \sigma^2 = n \frac{1}{n^2} \sigma^2 = \frac{\sigma^2}{n} \\ \implies \sigma_{\bar{X}} &= \sqrt{\sigma_{\bar{X}}^2} = \sqrt{\frac{\sigma^2}{n}} = \frac{\sigma}{\sqrt{n}} \end{aligned}$$

So $\bar{X} \sim N(\mu, \frac{\sigma}{\sqrt{n}})$.

Example 5.5.2

Suppose that SAT math scores for girls in Indiana are assumed to be $N(549, 24)$.

Find and compare the sampling distributions for the sample means from a sample of size $n = 16$ and a sample of size $n = 36$.

For $n = 16$: sample mean $\bar{X} \sim N\left(549, \frac{24}{\sqrt{16}} = 6\right)$


For $n = 36$: sample mean $\bar{Y} \sim N\left(549, \frac{24}{\sqrt{36}} = 4\right)$

Let's find the probability that each sample mean will be within 10 points of actual population mean ($\mu = 549$):

$\bar{X} : P(|\bar{X} - \mu| \leq 10) = P(539 \leq \bar{X} \leq 559) = \text{normalcdf}(539, 559, 549, 6) = 0.9044$

$\bar{Y} : P(539 \leq \bar{Y} \leq 559) = \text{normalcdf}(539, 559, 549, 4) = 0.9876$

The following figure gives the plot of the pdf's for the sampling distributions of \bar{X} (blue) and \bar{Y} (yellow). Note that the spread of the pdf for \bar{X} is larger than for \bar{Y} .




The Central Limit Theorem

We saw that when "sampling" from a normally distributed population, the sampling distribution of the sample mean is also normal. But what if the population does not follow a normal distribution? What if it is *skewed* or *uniform*?

Example 5.5.3

Suppose we are interested in the lifetime of a radioactive particle. We saw in [Section 4.5](#), that the probability distribution of such lifetimes can be modeled with an exponential distribution. If $\lambda = 3$, for example, then the pdf is *skewed right*, because there is a tail of values with very low probabilities off to the right.



Central Limit Theorem

Let X_1, \dots, X_n be a random sample from *any* probability distribution with mean μ and sd σ . Then as the sample size $n \rightarrow \infty$, the probability distribution of the sample mean approaches the normal distribution. We write:

$$\bar{X} \xrightarrow{d} N(\mu, \sigma/\sqrt{n}), \quad \text{as } n \rightarrow \infty \quad (5.5.1)$$

In other words, if n is *sufficiently large*, we can *approximate* the sampling distribution of the sample mean as $N(\mu, \sigma/\sqrt{n})$.

Furthermore,

$$\frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \approx N(0, 1) \text{ as } n \rightarrow \infty \quad (5.5.2)$$

The d above the arrow in Equation 5.5.1 above stands for *distribution* and indicates that, as the sample size increases without bound, the limit of the probability distribution of \bar{X} is given by the $N(\mu, \sigma/\sqrt{n})$ distribution. This is referred to as *convergence in distribution*.

What's "sufficiently large"?

- If the distribution of the X_i is *symmetric, unimodal or continuous*, then a sample size n as small as 4 or 5 yields an adequate approximation.
- If the distribution of the X_i is *skewed*, then a sample size n of at least 25 or 30 yields an adequate approximation.
- If the distribution of the X_i is *extremely skewed*, then you may need an even larger n .

The following website provides a simulation of sampling distributions and demonstrates the Central Limit Theorem ([link available](#)).

Example 5.5.4

Continuing in the context of [Example 5.5.3](#), suppose we sample $n = 25$ such radioactive particles. Then the sampling distribution of the mean of the sample is approximated as follows.

Letting X_1, \dots, X_{25} denote the random sample, we have that each $X_i \sim \text{exponential}(\lambda = 3)$. By the [Properties of Exponential Distributions](#), we know that the mean of an exponential(3) distribution is given by $\mu = \frac{1}{\lambda} = \frac{1}{3}$ and the sd is also $\sigma = \frac{1}{\lambda} = \frac{1}{3}$. Thus, the sampling distribution of the sample mean is

$$\bar{X} \sim N\left(\frac{1}{3}, \frac{1/3}{\sqrt{25}}\right) \Rightarrow N\left(\frac{1}{3}, \frac{1}{15}\right).$$

What is the use of the Central Limit Theorem if we don't know μ , the mean of the population? We can use the CLT to approximate **estimation error probabilities**:

$$P(|\bar{x} - \mu| \leq \varepsilon), \quad (5.5.3)$$

the probability that \bar{X} is within ε units of μ . By the [Central Limit Theorem](#) and Equation 5.5.2, we know

$$\frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \approx N(0, 1). \quad (5.5.4)$$

From this fact, we can isolate μ in the inequality in Equation 5.5.3 as follows:

$$P(|\bar{X} - \mu| \leq \varepsilon) = P\left(\frac{|\bar{X} - \mu|}{\sigma/\sqrt{n}} \leq \frac{\varepsilon}{\sigma/\sqrt{n}}\right) \approx P\left(|Z| \leq \frac{\varepsilon}{\sigma/\sqrt{n}}\right) = P\left(-\frac{\varepsilon}{\sigma/\sqrt{n}} \leq Z \leq \frac{\varepsilon}{\sigma/\sqrt{n}}\right) \quad (5.5.5)$$

Example 5.5.5

Now suppose that we do not know the rate at which the radioactive particle of interest decays, i.e., we do not know the mean lifetime of such particles. We can develop a method for *approximating* the probability that the mean of a sample of size $n = 25$ is within 1 unit of the mean lifetime.

In other words, we want $P(|\bar{X} - \mu| \leq 1)$.

By the [Central Limit Theorem](#) and Equation 5.5.2, we know that

$$\frac{\bar{X} - \mu}{\sigma/\sqrt{25}} = \frac{\bar{X} - \mu}{\sigma/5} \approx N(0, 1).$$

From this we derive a formula for the desired probability:

$$P\left(\frac{\bar{X} - \mu}{\sigma/5} \leq \frac{1}{\sigma/5}\right) \approx P\left(|Z| \leq \frac{5}{\sigma}\right)$$

This page titled 5.5: Sample Mean is shared under a [not declared](#) license and was authored, remixed, and/or curated by [Kristin Kuter](#).

Index

W

Weibull Distributions

[4.6: Weibull Distributions](#)

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: MATH 345 - Probability (Kuter)

Webpages: 41

All licenses found:

- [Undeclared](#): 100% (41 pages)

By Page

- [MATH 345 - Probability \(Kuter\) - Undeclared](#)
 - [Front Matter - Undeclared](#)
 - [TitlePage - Undeclared](#)
 - [InfoPage - Undeclared](#)
 - [Table of Contents - Undeclared](#)
 - [Licensing - Undeclared](#)
 - [1: What is Probability? - Undeclared](#)
 - [1.1: Sample Spaces and Events - Undeclared](#)
 - [1.2: Probability Measures - Undeclared](#)
 - [2: Computing Probabilities - Undeclared](#)
 - [2.1: Equally Likely Outcomes and Counting Techniques \(Combinatorics\) - Undeclared](#)
 - [2.2: Conditional Probability and Bayes' Rule - Undeclared](#)
 - [2.3: Independent Events - Undeclared](#)
 - [3: Discrete Random Variables - Undeclared](#)
 - [3.1: Introduction to Random Variables - Undeclared](#)
 - [3.2: Probability Mass Functions \(PMFs\) and Cumulative Distribution Functions \(CDFs\) for Discrete Random Variables - Undeclared](#)
 - [3.3: Bernoulli and Binomial Distributions - Undeclared](#)
 - [3.4: Hypergeometric, Geometric, and Negative Binomial Distributions - Undeclared](#)
 - [3.5: Poisson Distribution - Undeclared](#)
 - [3.6: Expected Value of Discrete Random Variables - Undeclared](#)
 - [3.7: Variance of Discrete Random Variables - Undeclared](#)
 - [3.8: Moment-Generating Functions \(MGFs\) for Discrete Random Variables - Undeclared](#)
 - [4: Continuous Random Variables - Undeclared](#)
 - [4.1: Probability Density Functions \(PDFs\) and Cumulative Distribution Functions \(CDFs\) for Continuous Random Variables - Undeclared](#)
 - [4.2: Expected Value and Variance of Continuous Random Variables - Undeclared](#)
 - [4.3: Uniform Distributions - Undeclared](#)
 - [4.4: Normal Distributions - Undeclared](#)
 - [4.5: Exponential and Gamma Distributions - Undeclared](#)
 - [4.6: Weibull Distributions - Undeclared](#)
 - [4.7: Chi-Squared Distributions - Undeclared](#)
 - [4.8: Beta Distributions - Undeclared](#)
 - [5: Probability Distributions for Combinations of Random Variables - Undeclared](#)
 - [5.1: Joint Distributions of Discrete Random Variables - Undeclared](#)
 - [5.2: Joint Distributions of Continuous Random Variables - Undeclared](#)
 - [5.3: Conditional Probability Distributions - Undeclared](#)
 - [5.4: Finding Distributions of Functions of Continuous Random Variables - Undeclared](#)
 - [5.5: Sample Mean - Undeclared](#)
 - [Back Matter - Undeclared](#)
 - [Index - Undeclared](#)
 - [Glossary - Undeclared](#)
 - [Detailed Licensing - Undeclared](#)