

INTRODUCTION TO BUSINESS STATISTICS - OER - SPRING 2023



JoEllen Green
Fresno City College

Licensing

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

Fresno City College
Introduction to Business Statistics - OER -
Spring 2023

JoEllen Green

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 12/19/2023

TABLE OF CONTENTS

Licensing

Licensing

1: Sampling and Data

- 1.1: Introduction to Sampling and Data
- 1.2: Definitions of Statistics, Probability, and Key Terms
- 1.3: Data, Sampling, and Variation in Data and Sampling
- 1.4: Levels of Measurement
- 1.5: Experimental Design and Ethics
- 1.6: Chapter Key Terms
- 1.7: Chapter References
- 1.H: Sampling and Data (Homework)
- 1.R: Sampling and Data (Review)
- 1.S: Sampling and Data (Solutions)

2: Descriptive Statistics

- 2.1: introduction to Descriptive Statistics
- 2.2: Display Data
- 2.3: Measures of the Location of the Data
- 2.4: Measures of the Center of the Data
- 2.5: Sigma Notation and Calculating the Arithmetic Mean
- 2.6: Geometric Mean
- 2.7: Skewness and the Mean, Median, and Mode
- 2.8: Measures of the Spread of the Data
- 2.9: Chapter Key Terms
- 2.10: Chapter Formula Review
- 2.11: Chapter Homework
- 2.12: Chapter Practice
- 2.13: Homework
- 2.14: Chapter References
- 2.15: Chapter Homework Solutions
- 2.R: Descriptive Statistics (Review)

3: Probability Topics

- 3.1: Introduction to Probability
- 3.2: Probability Terminology
- 3.3: Independent and Mutually Exclusive Events
- 3.4: Two Basic Rules of Probability
- 3.5: Contingency Tables and Probability Trees
- 3.6: Venn Diagrams
- 3.7: Chapter Key Terms
- 3.8: Chapter Review
- 3.9: Chapter Formula Review
- 3.10: Chapter Practice
- 3.11: Chapter More Practice
- 3.12: Chapter Homework

- 3.13: Chapter Reference
- 3.14: Chapter Solution (Practice + Homework)

4: Discrete Random Variables

- 4.1: Introduction to Discrete Random Variables
- 4.2: Hypergeometric Distribution
- 4.3: Binomial Distribution
- 4.4: Geometric Distribution
- 4.5: Poisson Distribution
- 4.6: Chapter Key Items
- 4.7: Chapter Review
- 4.8: Chapter Formula Review
- 4.9: Chapter Practice
- 4.10: Chapter Homework
- 4.11: Chapter References
- 4.12: Chapter Solution (Practice + Homework)

5: Continuous Random Variables

- 5.1: Prelude to Continuous Random Variables
- 5.2: Properties of Continuous Probability Density Functions
- 5.3: The Uniform Distribution
- 5.4: The Exponential Distribution
- 5.5: Chapter Key Terms
- 5.6: Chapter Review
- 5.7: Chapter Formula Review
- 5.8: Chapter Practice
- 5.9: Chapter Homework
- 5.10: Chapter References
- 5.11: Chapter Solution (Practice + Homework)

6: The Normal Distribution

- 6.1: Introduction to Normal Distribution
- 6.2: The Standard Normal Distribution
- 6.3: Using the Normal Distribution
- 6.4: Estimating the Binomial with the Normal Distribution
- 6.5: Chapter Key Items
- 6.6: Chapter Review
- 6.7: Chapter Formula Review
- 6.8: Chapter Practice
- 6.9: Chapter Homework
- 6.10: Chapter References
- 6.11: Chapter Solution (Practice + Homework)

7: The Central Limit Theorem

- 7.1: Introduction to the Central Limit Theorem
- 7.2: The Central Limit Theorem for Sample Means
- 7.3: Using the Central Limit Theorem
- 7.4: The Central Limit Theorem for Proportions
- 7.5: Finite Population Correction Factor
- 7.6: Chapter Key Terms

- 7.7: Chapter Review
- 7.8: Chapter Formula Review
- 7.9: Chapter Practice
- 7.10: Chapter Homework
- 7.11: Chapter References
- 7.12: Chapter Solution (Practice + Homework)

8: Confidence Intervals

- 8.1: Introduction to Confidence Intervals
- 8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size
- 8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case
- 8.4: A Confidence Interval for A Population Proportion
- 8.5: Calculating the Sample Size n- Continuous and Binary Random Variables
- 8.6: Chapter Key Terms
- 8.7: Chapter Review
- 8.8: Chapter Formula Review
- 8.9: Chapter Practice
- 8.10: Chapter Homework
- 8.11: Chapter References
- 8.12: Chapter Solution (Practice + Homework)

9: Hypothesis Testing with One Sample

- 9.1: Introduction to Hypothesis Testing
- 9.2: Null and Alternative Hypotheses
- 9.3: Outcomes and the Type I and Type II Errors
- 9.4: Distribution Needed for Hypothesis Testing
- 9.5: Full Hypothesis Test Examples
- 9.6: Chapter Key Terms
- 9.7: Chapter Review
- 9.8: Chapter Formula Review
- 9.9: Chapter Practice
- 9.10: Chapter Homework
- 9.11: Chapter References
- 9.12: Chapter Solution (Practice + Homework)

10: Hypothesis Testing with Two Samples

- 10.1: Introduction
- 10.2: Comparing Two Independent Population Means
- 10.3: Cohen's Standards for Small, Medium, and Large Effect Sizes
- 10.4: Test for Differences in Means- Assuming Equal Population Variances
- 10.5: Comparing Two Independent Population Proportions
- 10.6: Two Population Means with Known Standard Deviations
- 10.7: Matched or Paired Samples
- 10.8: Chapter Key Terms
- 10.9: Chapter Review
- 10.10: Chapter Formula Review
- 10.11: Chapter Practice
- 10.12: Homework
- 10.13: Chapter Homework
- 10.14: Chapter References
- 10.15: Chapter Solution (Practice + Homework)

11: The Chi-Square Distribution

- 11.1: Prelude to the Chi-Square Distribution
- 11.2: Facts About the Chi-Square Distribution
- 11.3: Test of a Single Variance
- 11.4: Goodness-of-Fit Test
- 11.5: Test of Independence
- 11.6: Test for Homogeneity
- 11.7: Comparison of the Chi-Square Tests
- 11.8: Chapter Key Terms
- 11.9: Chapter Review
- 11.10: Chapter Formula Review
- 11.11: Chapter Practice
- 11.12: Homework
- 11.13: Chapter Homework
- 11.14: Chapter References
- 11.15: Chapter Solution (Practice + Homework)

12: F Distribution and One-Way ANOVA

- 12.1: Introduction to F Distribution and One-Way ANOVA
- 12.2: Test of Two Variances
- 12.3: One-Way ANOVA
- 12.4: The F Distribution and the F-Ratio
- 12.5: Facts About the F Distribution
- 12.6: Chapter Key Terms
- 12.7: Chapter Review
- 12.8: Chapter Formula Review
- 12.9: Chapter Practice
- 12.10: Chapter Homework
- 12.11: Chapter Reference
- 12.12: Chapter Solution (Practice + Homework)

13: Linear Regression and Correlation

- 13.1: Introduction to Linear Regression and Correlation
- 13.2: The Correlation Coefficient r
- 13.3: Testing the Significance of the Correlation Coefficient
- 13.4: Linear Equations
- 13.5: The Regression Equation
- 13.6: Interpretation of Regression Coefficients- Elasticity and Logarithmic Transformation
- 13.7: Predicting with a Regression Equation
- 13.8: How to Use Microsoft Excel® for Regression Analysis
- 13.9: Chapter Key Terms
- 13.10: Chapter Review
- 13.11: Chapter Practice
- 13.12: Chapter Solution

[Index](#)

[Glossary](#)

Licensing

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

CHAPTER OVERVIEW

1: Sampling and Data

- 1.1: Introduction to Sampling and Data
- 1.2: Definitions of Statistics, Probability, and Key Terms
- 1.3: Data, Sampling, and Variation in Data and Sampling
- 1.4: Levels of Measurement
- 1.5: Experimental Design and Ethics
- 1.6: Chapter Key Terms
- 1.7: Chapter References
- 1.H: Sampling and Data (Homework)
- 1.R: Sampling and Data (Review)
- 1.S: Sampling and Data (Solutions)

This page titled [1: Sampling and Data](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

1.1: Introduction to Sampling and Data

You are probably asking yourself the question, "When and where will I use statistics?" If you read any newspaper, watch television, or use the Internet, you will see statistical information. There are statistics about crime, sports, education, politics, and real estate. Typically, when you read a newspaper article or watch a television news program, you are given sample information. With this information, you may make a decision about the correctness of a statement, claim, or "fact." Statistical methods can help you make the "best educated guess."



Figure 1.1.1 We encounter statistics in our daily lives more often than we probably realize and from many different sources, like the news. (credit: David Sim)

Since you will undoubtedly be given statistical information at some point in your life, you need to know some techniques for analyzing the information thoughtfully. Think about buying a house or managing a budget. Think about your chosen profession. The fields of economics, business, psychology, education, biology, law, computer science, police science, and early childhood development require at least one course in statistics.

Included in this chapter are the basic ideas and words of probability and statistics. You will soon understand that statistics and probability work together. You will also learn how data are gathered and what "good" data can be distinguished from "bad."

This page titled [1.1: Introduction to Sampling and Data](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [1.0: Introduction to Sampling and Data](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.2: Definitions of Statistics, Probability, and Key Terms

The science of **statistics** deals with the collection, analysis, interpretation, and presentation of **data**. We see and use data in our everyday lives.

In this course, you will learn how to organize and summarize data. Organizing and summarizing data is called **descriptive statistics**. Two ways to summarize data are by graphing and by using numbers (for example, finding an average). After you have studied probability and probability distributions, you will use formal methods for drawing conclusions from "good" data. The formal methods are called **inferential statistics**. Statistical inference uses probability to determine how confident we can be that our conclusions are correct.

Effective interpretation of data (inference) is based on good procedures for producing data and thoughtful examination of the data. You will encounter what will seem to be too many mathematical formulas for interpreting data. The goal of statistics is not to perform numerous calculations using the formulas, but to gain an understanding of your data. The calculations can be done using a calculator or a computer. The understanding must come from you. If you can thoroughly grasp the basics of statistics, you can be more confident in the decisions you make in life.

Probability

Probability is a mathematical tool used to study randomness. It deals with the chance (the likelihood) of an event occurring. For example, if you toss a **fair** coin four times, the outcomes may not be two heads and two tails. However, if you toss the same coin 4,000 times, the outcomes will be close to half heads and half tails. The expected theoretical probability of heads in any one toss is $\frac{1}{2}$ or 0.5. Even though the outcomes of a few repetitions are uncertain, there is a regular pattern of outcomes when there are many repetitions. After reading about the English statistician Karl **Pearson** who tossed a coin 24,000 times with a result of 12,012 heads, one of the authors tossed a coin 2,000 times. The results were 996 heads. The fraction $\frac{996}{2000}$ is equal to 0.498 which is very close to 0.5, the expected probability.

The theory of probability began with the study of games of chance such as poker. Predictions take the form of probabilities. To predict the likelihood of an earthquake, of rain, or whether you will get an A in this course, we use probabilities. Doctors use probability to determine the chance of a vaccination causing the disease the vaccination is supposed to prevent. A stockbroker uses probability to determine the rate of return on a client's investments. You might use probability to decide to buy a lottery ticket or not. In your study of statistics, you will use the power of mathematics through probability calculations to analyze and interpret your data.

Key Terms

In statistics, we generally want to study a **population**. You can think of a population as a collection of persons, things, or objects under study. To study the population, we select a **sample**. The idea of **sampling** is to select a portion (or subset) of the larger population and study that portion (the sample) to gain information about the population. Data are the result of sampling from a population.

Because it takes a lot of time and money to examine an entire population, sampling is a very practical technique. If you wished to compute the overall grade point average at your school, it would make sense to select a sample of students who attend the school. The data collected from the sample would be the students' grade point averages. In presidential elections, opinion poll samples of 1,000–2,000 people are taken. The opinion poll is supposed to represent the views of the people in the entire country. Manufacturers of canned carbonated drinks take samples to determine if a 16 ounce can contains 16 ounces of carbonated drink.

From the sample data, we can calculate a statistic. A **statistic** is a number that represents a property of the sample. For example, if we consider one math class to be a sample of the population of all math classes, then the average number of points earned by students in that one math class at the end of the term is an example of a statistic. The statistic is an estimate of a population parameter, in this case the mean. A **parameter** is a numerical characteristic of the whole population that can be estimated by a statistic. Since we considered all math classes to be the population, then the average number of points earned per student over all the math classes is an example of a parameter.

One of the main concerns in the field of statistics is how accurately a statistic estimates a parameter. The accuracy really depends on how well the sample represents the population. The sample must contain the characteristics of the population in order to be a

representative sample. We are interested in both the sample statistic and the population parameter in inferential statistics. In a later chapter, we will use the sample statistic to test the validity of the established population parameter.

A **variable**, or random variable, usually notated by capital letters such as X and Y , is a characteristic or measurement that can be determined for each member of a population. Variables may be **numerical** or **categorical**. **Numerical variables** take on values with equal units such as weight in pounds and time in hours. **Categorical variables** place the person or thing into a category. If we let X equal the number of points earned by one math student at the end of a term, then X is a numerical variable. If we let Y be a person's party affiliation, then some examples of Y include Republican, Democrat, and Independent. Y is a categorical variable. We could do some math with values of X (calculate the average number of points earned, for example), but it makes no sense to do math with values of Y (calculating an average party affiliation makes no sense).

Data are the actual values of the variable. They may be numbers or they may be words. **Datum** is a single value.

Two words that come up often in statistics are **mean** and **proportion**. If you were to take three exams in your math classes and obtain scores of 86, 75, and 92, you would calculate your mean score by adding the three exam scores and dividing by three (your mean score would be 84.3 to one decimal place). If, in your math class, there are 40 students and 22 are men and 18 are women, then the proportion of men students is $\frac{22}{40}$ and the proportion of women students is $\frac{18}{40}$. Mean and proportion are discussed in more detail in later chapters.

NOTE

The words "**mean**" and "**average**" are often used interchangeably. The substitution of one word for the other is common practice. The technical term is "arithmetic mean," and "average" is technically a center location. However, in practice among non-statisticians, "average" is commonly accepted for "arithmetic mean."

Example 1.1

Determine what the key terms refer to in the following study. We want to know the average (mean) amount of money first year college students spend at ABC College on school supplies that do not include books. We randomly surveyed 100 first year students at the college. Three of those students spent \$150, \$200, and \$225, respectively.

Answer

Solution 1.1

The **population** is all first year students attending ABC College this term.

The **sample** could be all students enrolled in one section of a beginning statistics course at ABC College (although this sample may not represent the entire population).

The **parameter** is the average (mean) amount of money spent (excluding books) by first year college students at ABC College this term: the population mean.

The **statistic** is the average (mean) amount of money spent (excluding books) by first year college students in the sample.

The **variable** could be the amount of money spent (excluding books) by one first year student. Let X = the amount of money spent (excluding books) by one first year student attending ABC College.

The **data** are the dollar amounts spent by the first year students. Examples of the data are \$150, \$200, and \$225.

Exercise 1.1

Determine what the key terms refer to in the following study. We want to know the average (mean) amount of money spent on school uniforms each year by families with children at Knoll Academy. We randomly survey 100 families with children in the school. Three of the families spent \$65, \$75, and \$95, respectively.

Example 1.2

Determine what the key terms refer to in the following study.

A study was conducted at a local college to analyze the average cumulative GPA's of students who graduated last year. Fill in the letter of the phrase that best describes each of the items below.

1. Population ____ 2. Statistic ____ 3. Parameter ____ 4. Sample ____ 5. Variable ____ 6. Data ____
- all students who attended the college last year
 - the cumulative GPA of one student who graduated from the college last year
 - 3.65, 2.80, 1.50, 3.90
 - a group of students who graduated from the college last year, randomly selected
 - the average cumulative GPA of students who graduated from the college last year
 - all students who graduated from the college last year
 - the average cumulative GPA of students in the study who graduated from the college last year

Answer

Solution 1.2

1. f; 2. g; 3. e; 4. d; 5. b; 6. c

Example 1.3

Determine what the key terms refer to in the following study.

As part of a study designed to test the safety of automobiles, the National Transportation Safety Board collected and reviewed data about the effects of an automobile crash on test dummies. Here is the criterion they used:

Table 1.1

Speed at which cars crashed	Location of "drive" (i.e. dummies)
35 miles/hour	Front Seat

Cars with dummies in the front seats were crashed into a wall at a speed of 35 miles per hour. We want to know the proportion of dummies in the driver's seat that would have had head injuries, if they had been actual drivers. We start with a simple random sample of 75 cars.

Answer

Solution 1.3

The **population** is all cars containing dummies in the front seat.

The **sample** is the 75 cars, selected by a simple random sample.

The **parameter** is the proportion of driver dummies (if they had been real people) who would have suffered head injuries in the population.

The **statistic** is proportion of driver dummies (if they had been real people) who would have suffered head injuries in the sample.

The **variable** X = the number of driver dummies (if they had been real people) who would have suffered head injuries.

The **data** are either: yes, had head injury, or no, did not.

Example 1.4

Determine what the key terms refer to in the following study.

An insurance company would like to determine the proportion of all medical doctors who have been involved in one or more malpractice lawsuits. The company selects 500 doctors at random from a professional directory and determines the number in the sample who have been involved in a malpractice lawsuit.

Answer

Solution 1.4

The **population** is all medical doctors listed in the professional directory.

The **parameter** is the proportion of medical doctors who have been involved in one or more malpractice suits in the population.

The **sample** is the 500 doctors selected at random from the professional directory.

The **statistic** is the proportion of medical doctors who have been involved in one or more malpractice suits in the sample.

The **variable** X = the number of medical doctors who have been involved in one or more malpractice suits.

The **data** are either: yes, was involved in one or more malpractice lawsuits, or no, was not.

This page titled [1.2: Definitions of Statistics, Probability, and Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [1.1: Definitions of Statistics, Probability, and Key Terms](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.3: Data, Sampling, and Variation in Data and Sampling

Data may come from a population or from a sample. Lowercase letters like x or y generally are used to represent data values. Most data can be put into the following categories:

- Qualitative
- Quantitative

Qualitative data are the result of categorizing or describing attributes of a population. **Qualitative data** are also often called categorical data. Hair color, blood type, ethnic group, the car a person drives, and the street a person lives on are examples of qualitative(categorical) data. Qualitative(categorical) data are generally described by words or letters. For instance, hair color might be black, dark brown, light brown, blonde, gray, or red. Blood type might be AB+, O-, or B+. Researchers often prefer to use quantitative data over qualitative(categorical) data because it lends itself more easily to mathematical analysis. For example, it does not make sense to find an average hair color or blood type.

Quantitative data are always numbers. Quantitative data are the result of **counting** or **measuring** attributes of a population. Amount of money, pulse rate, weight, number of people living in your town, and number of students who take statistics are examples of quantitative data. Quantitative data may be either **discrete** or **continuous**.

All data that are the result of counting are called **quantitative discrete data**. These data take on only certain numerical values. If you count the number of phone calls you receive for each day of the week, you might get values such as zero, one, two, or three.

Data that are not only made up of counting numbers, but that may include fractions, decimals, or irrational numbers, are called **quantitative continuous data**. Continuous data are often the results of measurements like lengths, weights, or times. A list of the lengths in minutes for all the phone calls that you make in a week, with numbers like 2.4, 7.5, or 11.0, would be quantitative continuous data.

Example 1.3.1: DATA SAMPLE OF QUANTITATIVE DISCRETE DATA

The data are the number of books students carry in their backpacks. You sample five students. Two students carry three books, one student carries four books, one student carries two books, and one student carries one book. The numbers of books (three, four, two, and one) are the quantitative discrete data.

Exercise 1.3.1

The data are the number of machines in a gym. You sample five gyms. One gym has 12 machines, one gym has 15 machines, one gym has ten machines, one gym has 22 machines, and the other gym has 20 machines. What type of data is this?

Example 1.3.2: DATA SAMPLE OF QUANTITATIVE CONTINUOUS DATA

The data are the weights of backpacks with books in them. You sample the same five students. The weights (in pounds) of their backpacks are 6.2, 7, 6.8, 9.1, 4.3. Notice that backpacks carrying three books can have different weights. Weights are quantitative continuous data.

Exercise 1.3.2

The data are the areas of lawns in square feet. You sample five houses. The areas of the lawns are 144 sq. feet, 160 sq. feet, 190 sq. feet, 180 sq. feet, and 210 sq. feet. What type of data is this?

Example 1.3.3

You go to the supermarket and purchase three cans of soup (19 ounces) tomato bisque, 14.1 ounces lentil, and 19 ounces Italian wedding), two packages of nuts (walnuts and peanuts), four different kinds of vegetable (broccoli, cauliflower, spinach, and carrots), and two desserts (16 ounces pistachio ice cream and 32 ounces chocolate chip cookies).

Name data sets that are quantitative discrete, quantitative continuous, and qualitative(categorical).

Answer

One Possible Solution:

- The three cans of soup, two packages of nuts, four kinds of vegetables and two desserts are quantitative discrete data because you count them.
- The weights of the soups (19 ounces, 14.1 ounces, 19 ounces) are quantitative continuous data because you measure weights as precisely as possible.
- Types of soups, nuts, vegetables and desserts are qualitative(categorical) data because they are categorical.

Try to identify additional data sets in this example.

Example 1.3.4

The data are the colors of backpacks. Again, you sample the same five students. One student has a red backpack, two students have black backpacks, one student has a green backpack, and one student has a gray backpack. The colors red, black, black, green, and gray are qualitative(categorical) data.

Exercise 1.3.4

The data are the colors of houses. You sample five houses. The colors of the houses are white, yellow, white, red, and white. What type of data is this?

You may collect data as numbers and report it categorically. For example, the quiz scores for each student are recorded throughout the term. At the end of the term, the quiz scores are reported as A, B, C, D, or F

Example 1.3.5

Work collaboratively to determine the correct data type (quantitative or qualitative). Indicate whether quantitative data are continuous or discrete. Hint: Data that are discrete often start with the words "the number of."

- a. the number of pairs of shoes you own
- b. the type of car you drive
- c. the distance from your home to the nearest grocery store
- d. the number of classes you take per school year
- e. the type of calculator you use
- f. weights of sumo wrestlers
- g. number of correct answers on a quiz
- h. IQ scores (This may cause some discussion.)

Answer

Items a, d, and g are quantitative discrete; items c, f, and h are quantitative continuous; items b and e are qualitative, or categorical.

Exercise 1.3.5

Determine the correct data type (quantitative or qualitative) for the number of cars in a parking lot. Indicate whether quantitative data are continuous or discrete.

Example 1.3.6

A statistics professor collects information about the classification of her students as freshmen, sophomores, juniors, or seniors. The data she collects are summarized in the pie chart Figure 1.2. What type of data does this graph show?

Classification of Statistics Students




Figure 1.2

Answer

This pie chart shows the students in each year, which is **qualitative (or categorical) data**.

Exercise 1.3.6

The registrar at State University keeps records of the number of credit hours students complete each semester. The data he collects are summarized in the histogram. The class boundaries are 10 to less than 13, 13 to less than 16, 16 to less than 19, 19 to less than 22, and 22 to less than 25.




Figure 1.3

What type of data does this graph show?

Qualitative Data Discussion

Below are tables comparing the number of part-time and full-time students at De Anza College and Foothill College enrolled for the spring 2010 quarter. The tables display counts (frequencies) and percentages or proportions (relative frequencies). The percent columns make comparing the same categories in the colleges easier. Displaying percentages along with the numbers is often helpful, but it is particularly important when comparing sets of data that do not have the same totals, such as the total enrollments for both colleges in this example. Notice how much larger the percentage for part-time students at Foothill College is compared to De Anza College.

Table 1.3.1: Fall Term 2007 (Census day)

De Anza College			Foothill College		
	Number	Percent		Number	Percent
Full-time	9,200	40.9%	Full-time	4,059	28.6%
Part-time	13,296	59.1%	Part-time	10,124	71.4%
Total	22,496	100%	Total	14,183	100%

Tables are a good way of organizing and displaying data. But graphs can be even more helpful in understanding the data. There are no strict rules concerning which graphs to use. Two graphs that are used to display qualitative(categorical) data are pie charts and

bar graphs.

- In a **pie chart**, categories of data are represented by wedges in a circle and are proportional in size to the percent of individuals in each category.
- In a **bar graph**, the length of the bar for each category is proportional to the number or percent of individuals in each category. Bars may be vertical or horizontal.
- A **Pareto chart** consists of bars that are sorted into order by category size (largest to smallest).

Look at Figure 1.5 and determine which graph (pie or bar) you think displays the comparisons better.

It is a good idea to look at a variety of graphs to see which is the most helpful in displaying the data. We might make different choices of what we think is the “best” graph depending on the data and the context. Our choice also depends on what we are using the data for.

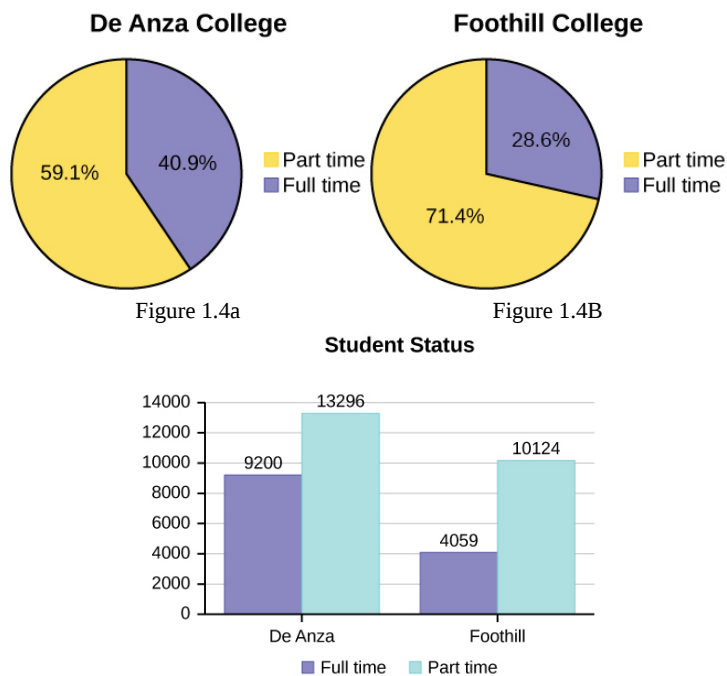


Figure 1.5

Percentages That Add to More (or Less) Than 100%

Sometimes percentages add up to be more than 100% (or less than 100%). In the graph, the percentages add to more than 100% because students can be in more than one category. A bar graph is appropriate to compare the relative size of the categories. A pie chart cannot be used. It also could not be used if the percentages added to less than 100%.

Table 1.3.2: De Anza College Spring 2010

Characteristic/category	Percent
Full-time students	40.9%
Students who intend to transfer to a 4-year educational institution	48.6%
Students under age 25	61.0%
TOTAL	150.5%




Figure 1.3.2

Omitting Categories/Missing Data

The table displays Ethnicity of Students but is missing the "Other/Unknown" category. This category contains people who did not feel they fit into any of the ethnicity categories or declined to respond. Notice that the frequencies do not add up to the total number of students. In this situation, create a bar graph and not a pie chart.

Table 1.3.3: Ethnicity of Students at De Anza College Fall Term 2007 (Census Day)

	Frequency	Percent
Asian	8,794	36.1%
Black	1,412	5.8%
Filipino	1,298	5.3%
Hispanic	4,180	17.1%
Native American	146	0.6%
Pacific Islander	236	1.0%
White	5,978	24.5%
TOTAL	22,044 out of 24,382	90.4% out of 100%

Figure 1.3.3

The following graph is the same as the previous graph but the "Other/Unknown" percent (9.6%) has been included. The "Other/Unknown" category is large compared to some of the other categories (Native American, 0.6%, Pacific Islander 1.0%). This is important to know when we think about what the data are telling us.

This particular bar graph in Figure 1.9 is a Pareto chart. The Pareto chart has the bars sorted from largest to smallest and is easier to read and interpret.




Figure 1.3.4: Bar Graph with Other/Unknown Category




Figure 1.3.4: Pareto Chart With Bars Sorted by Size

Pie Charts: No Missing Data

The following pie charts have the “Other/Unknown” category included (since the percentages must add to 100%). The chart in Figure 1.10.




Figure 1.3.5: Paste Caption Here

Sampling

Gathering information about an entire population often costs too much or is virtually impossible. Instead, we use a sample of the population. **A sample should have the same characteristics as the population it is representing.** Most statisticians use various methods of random sampling in an attempt to achieve this goal. This section will describe a few of the most common methods. There are several different methods of **random sampling**. In each form of random sampling, each member of a population initially has an equal chance of being selected for the sample. Each method has pros and cons. The easiest method to describe is called a **simple random sample**. Any group of n individuals is equally likely to be chosen as any other group of n individuals if the simple random sampling technique is used. In other words, each sample of the same size has an equal chance of being selected.

Besides simple random sampling, there are other forms of sampling that involve a chance process for getting the sample. **Other well-known random sampling methods are the stratified sample, the cluster sample, and the systematic sample.**

To choose a **stratified sample**, divide the population into groups called strata and then take a **proportionate** number from each stratum. For example, you could stratify (group) your college population by department and then choose a proportionate simple random sample from each stratum (each department) to get a stratified random sample. To choose a simple random sample from each department, number each member of the first department, number each member of the second department, and do the same for the remaining departments. Then use simple random sampling to choose proportionate numbers from the first department and do the same for each of the remaining departments. Those numbers picked from the first department, picked from the second department, and so on represent the members who make up the stratified sample.

To choose a **cluster sample**, divide the population into clusters (groups) and then randomly select some of the clusters. All the members from these clusters are in the cluster sample. For example, if you randomly sample four departments from your college population, the four departments make up the cluster sample. Divide your college faculty by department. The departments are the clusters. Number each department, and then choose four different numbers using simple random sampling. All members of the four departments with those numbers are the cluster sample.

To choose a **systematic sample**, randomly select a starting point and take every n^{th} piece of data from a listing of the population. For example, suppose you have to do a phone survey. Your phone book contains 20,000 residence listings. You must choose 400

names for the sample. Number the population 1–20,000 and then use a simple random sample to pick a number that represents the first name in the sample. Then choose every fiftieth name thereafter until you have a total of 400 names (you might have to go back to the beginning of your phone list). Systematic sampling is frequently chosen because it is a simple method.

A type of sampling that is non-random is convenience sampling. **Convenience sampling** involves using results that are readily available. For example, a computer software store conducts a marketing study by interviewing potential customers who happen to be in the store browsing through the available software. The results of convenience sampling may be very good in some cases and highly biased (favor certain outcomes) in others.

Sampling data should be done very carefully. Collecting data carelessly can have devastating results. Surveys mailed to households and then returned may be very biased (they may favor a certain group). It is better for the person conducting the survey to select the sample respondents.

True random sampling is done **with replacement**. That is, once a member is picked, that member goes back into the population and thus may be chosen more than once. However for practical reasons, in most populations, simple random sampling is done **without replacement**. Surveys are typically done without replacement. That is, a member of the population may be chosen only once. Most samples are taken from large populations and the sample tends to be small in comparison to the population. Since this is the case, sampling without replacement is approximately the same as sampling with replacement because the chance of picking the same individual more than once with replacement is very low.

In a college population of 10,000 people, suppose you want to pick a sample of 1,000 randomly for a survey. **For any particular sample of 1,000, if you are sampling with replacement,**

- the chance of picking the first person is 1,000 out of 10,000 (0.1000);
- the chance of picking a different second person for this sample is 999 out of 10,000 (0.0999);
- the chance of picking the same person again is 1 out of 10,000 (very low).

If you are sampling **without replacement,**

- the chance of picking the first person for any particular sample is 1000 out of 10,000 (0.1000);
- the chance of picking a different second person is 999 out of 9,999 (0.0999);
- you do not replace the first person before picking the next person.

Compare the fractions $999/10,000$ and $999/9,999$. For accuracy, carry the decimal answers to four decimal places. To four decimal places, these numbers are equivalent (0.0999).

Sampling without replacement instead of sampling with replacement becomes a mathematical issue only when the population is small. For example, if the population is 25 people, the sample is ten, and you are sampling **with replacement for any particular sample**, then the chance of picking the first person is ten out of 25, and the chance of picking a different second person is nine out of 25 (you replace the first person).

If you sample **without replacement**, then the chance of picking the first person is ten out of 25, and then the chance of picking the second person (who is different) is nine out of 24 (you do not replace the first person).

Compare the fractions $9/25$ and $9/24$. To four decimal places, $9/25 = 0.3600$ and $9/24 = 0.3750$. To four decimal places, these numbers are not equivalent.

When you analyze data, it is important to be aware of **sampling errors** and nonsampling errors. The actual process of sampling causes sampling errors. For example, the sample may not be large enough. Factors not related to the sampling process cause **nonsampling errors**. A defective counting device can cause a nonsampling error.

In reality, a sample will never be exactly representative of the population so there will always be some sampling error. As a rule, the larger the sample, the smaller the sampling error.

In statistics, a **sampling bias** is created when a sample is collected from a population and some members of the population are not as likely to be chosen as others (remember, each member of the population should have an equally likely chance of being chosen). When a sampling bias happens, there can be incorrect conclusions drawn about the population that is being studied.

Critical Evaluation

We need to evaluate the statistical studies we read about critically and analyze them before accepting the results of the studies. Common problems to be aware of include

- Problems with samples: A sample must be representative of the population. A sample that is not representative of the population is biased. Biased samples that are not representative of the population give results that are inaccurate and not valid.
- Self-selected samples: Responses only by people who choose to respond, such as call-in surveys, are often unreliable.
- Sample size issues: Samples that are too small may be unreliable. Larger samples are better, if possible. In some situations, having small samples is unavoidable and can still be used to draw conclusions. Examples: crash testing cars or medical testing for rare conditions
- Undue influence: collecting data or asking questions in a way that influences the response
- Non-response or refusal of subject to participate: The collected responses may no longer be representative of the population. Often, people with strong positive or negative opinions may answer surveys, which can affect the results.
- Causality: A relationship between two variables does not mean that one causes the other to occur. They may be related (correlated) because of their relationship through a different variable.
- Self-funded or self-interest studies: A study performed by a person or organization in order to support their claim. Is the study impartial? Read the study carefully to evaluate the work. Do not automatically assume that the study is good, but do not automatically assume the study is bad either. Evaluate it on its merits and the work done.
- Misleading use of data: improperly displayed graphs, incomplete data, or lack of context
- Confounding: When the effects of multiple factors on a response cannot be separated. Confounding makes it difficult or impossible to draw valid conclusions about the effect of each factor.

Example 1.3.7

A study is done to determine the average tuition that San Jose State undergraduate students pay per semester. Each student in the following samples is asked how much tuition he or she paid for the Fall semester. What is the type of sampling in each case?

- A sample of 100 undergraduate San Jose State students is taken by organizing the students' names by classification (freshman, sophomore, junior, or senior), and then selecting 25 students from each.
- A random number generator is used to select a student from the alphabetical listing of all undergraduate students in the Fall semester. Starting with that student, every 50th student is chosen until 75 students are included in the sample.
- A completely random method is used to select 75 students. Each undergraduate student in the fall semester has the same probability of being chosen at any stage of the sampling process.
- The freshman, sophomore, junior, and senior years are numbered one, two, three, and four, respectively. A random number generator is used to pick two of those years. All students in those two years are in the sample.
- An administrative assistant is asked to stand in front of the library one Wednesday and to ask the first 100 undergraduate students he encounters what they paid for tuition the Fall semester. Those 100 students are the sample.

Answer

- a. stratified; b. systematic; c. simple random; d. cluster; e. convenience

Example 1.3.8

Determine the type of sampling used (simple random, stratified, systematic, cluster, or convenience).

- A soccer coach selects six players from a group of boys aged eight to ten, seven players from a group of boys aged 11 to 12, and three players from a group of boys aged 13 to 14 to form a recreational soccer team.
- A pollster interviews all human resource personnel in five different high tech companies.
- A high school educational researcher interviews 50 high school female teachers and 50 high school male teachers.
- A medical researcher interviews every third cancer patient from a list of cancer patients at a local hospital.
- A high school counselor uses a computer to generate 50 random numbers and then picks students whose names correspond to the numbers.
- A student interviews classmates in his algebra class to determine how many pairs of jeans a student owns, on the average.

Answer

- a. stratified; b. cluster; c. stratified; d. systematic; e. simple random; f. convenience

If we were to examine two samples representing the same population, even if we used random sampling methods for the samples, they would not be exactly the same. Just as there is variation in data, there is variation in samples. As you become accustomed to sampling, the variability will begin to seem natural.

Example 1.3.8

Suppose ABC College has 10,000 part-time students (the population). We are interested in the average amount of money a part-time student spends on books in the fall term. Asking all 10,000 students is an almost impossible task.

Suppose we take two different samples.

First, we use convenience sampling and survey ten students from a first term organic chemistry class. Many of these students are taking first term calculus in addition to the organic chemistry class. The amount of money they spend on books is as follows:

\$128; \$87; \$173; \$116; \$130; \$204; \$147; \$189; \$93; \$153

The second sample is taken using a list of senior citizens who take P.E. classes and taking every fifth senior citizen on the list, for a total of ten senior citizens. They spend:

\$50; \$40; \$36; \$15; \$50; \$100; \$40; \$53; \$22; \$22

It is unlikely that any student is in both samples.

a. Do you think that either of these samples is representative of (or is characteristic of) the entire 10,000 part-time student population?

Answer

a. No. The first sample probably consists of science-oriented students. Besides the chemistry course, some of them are also taking first-term calculus. Books for these classes tend to be expensive. Most of these students are, more than likely, paying more than the average part-time student for their books. The second sample is a group of senior citizens who are, more than likely, taking courses for health and interest. The amount of money they spend on books is probably much less than the average part-time student. Both samples are biased. Also, in both cases, not all students have a chance to be in either sample.

b. Since these samples are not representative of the entire population, is it wise to use the results to describe the entire population?

Answer

Solution 1.13

b. No. For these samples, each member of the population did not have an equally likely chance of being chosen.

Now, suppose we take a third sample. We choose ten different part-time students from the disciplines of chemistry, math, English, psychology, sociology, history, nursing, physical education, art, and early childhood development. (We assume that these are the only disciplines in which part-time students at ABC College are enrolled and that an equal number of part-time students are enrolled in each of the disciplines.) Each student is chosen using simple random sampling. Using a calculator, random numbers are generated and a student from a particular discipline is selected if he or she has a corresponding number. The students spend the following amounts:

\$180; \$50; \$150; \$85; \$260; \$75; \$180; \$200; \$200; \$150

c. Is the sample biased?

Answer

Solution 1.13

c. The sample is unbiased, but a larger sample would be recommended to increase the likelihood that the sample will be close to representative of the population. However, for a biased sampling technique, even a large sample runs the risk of not being representative of the population.

Students often ask if it is "good enough" to take a sample, instead of surveying the entire population. If the survey is done well, the answer is yes.

Exercise 1.3.8

A local radio station has a fan base of 20,000 listeners. The station wants to know if its audience would prefer more music or more talk shows. Asking all 20,000 listeners is an almost impossible task.

The station uses convenience sampling and surveys the first 200 people they meet at one of the station's music concert events. 24 people said they'd prefer more talk shows, and 176 people said they'd prefer more music.

Do you think that this sample is representative of (or is characteristic of) the entire 20,000 listener population?

Variation in Data

Variation is present in any set of data. For example, 16-ounce cans of beverage may contain more or less than 16 ounces of liquid. In one study, eight 16 ounce cans were measured and produced the following amount (in ounces) of beverage:

15.8; 16.1; 15.2; 14.8; 15.8; 15.9; 16.0; 15.5

Measurements of the amount of beverage in a 16-ounce can may vary because different people make the measurements or because the exact amount, 16 ounces of liquid, was not put into the cans. Manufacturers regularly run tests to determine if the amount of beverage in a 16-ounce can falls within the desired range.

Be aware that as you take data, your data may vary somewhat from the data someone else is taking for the same purpose. This is completely natural. However, if two or more of you are taking the same data and get very different results, it is time for you and the others to reevaluate your data-taking methods and your accuracy.

Variation in Samples

It was mentioned previously that two or more **samples** from the same **population**, taken randomly, and having close to the same characteristics of the population will likely be different from each other. Suppose Doreen and Jung both decide to study the average amount of time students at their college sleep each night. Doreen and Jung each take samples of 500 students. Doreen uses systematic sampling and Jung uses cluster sampling. Doreen's sample will be different from Jung's sample. Even if Doreen and Jung used the same sampling method, in all likelihood their samples would be different. Neither would be wrong, however.

Think about what contributes to making Doreen's and Jung's samples different.

If Doreen and Jung took larger samples (i.e. the number of data values is increased), their sample results (the average amount of time a student sleeps) might be closer to the actual population average. But still, their samples would be, in all likelihood, different from each other. This **variability in samples** cannot be stressed enough.

Size of a Sample

The size of a sample (often called the number of observations, usually given the symbol n) is important. The examples you have seen in this book so far have been small. Samples of only a few hundred observations, or even smaller, are sufficient for many purposes. In polling, samples that are from 1,200 to 1,500 observations are considered large enough and good enough if the survey is random and is well done. Later we will find that even much smaller sample sizes will give very good results. You will learn why when you study confidence intervals.

Be aware that many large samples are biased. For example, call-in surveys are invariably biased, because people choose to respond or not.

This page titled [1.3: Data, Sampling, and Variation in Data and Sampling](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **1.2: Data, Sampling, and Variation in Data and Sampling** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.4: Levels of Measurement

Once you have a set of data, you will need to organize it so that you can analyze how frequently each datum occurs in the set. However, when calculating the frequency, you may need to round your answers so that they are as precise as possible.

Levels of Measurement

The way a set of data is measured is called its **level of measurement**. Correct statistical procedures depend on a researcher being familiar with levels of measurement. Not every statistical operation can be used with every set of data. Data can be classified into four levels of measurement. They are (from lowest to highest level):

- **Nominal scale level**
- **Ordinal scale level**
- **Interval scale level**
- **Ratio scale level**

Data that is measured using a **nominal scale** is **qualitative (categorical)**. Categories, colors, names, labels and favorite foods along with yes or no responses are examples of nominal level data. Nominal scale data are not ordered. For example, trying to classify people according to their favorite food does not make any sense. Putting pizza first and sushi second is not meaningful.

Smartphone companies are another example of nominal scale data. The data are the names of the companies that make smartphones, but there is no agreed upon order of these brands, even though people may have personal preferences. Nominal scale data cannot be used in calculations.

Data that is measured using an **ordinal scale** is similar to nominal scale data but there is a big difference. The ordinal scale data can be ordered. An example of ordinal scale data is a list of the top five national parks in the United States. The top five national parks in the United States can be ranked from one to five but we cannot measure differences between the data.

Another example of using the ordinal scale is a cruise survey where the responses to questions about the cruise are “excellent,” “good,” “satisfactory,” and “unsatisfactory.” These responses are ordered from the most desired response to the least desired. But the differences between two pieces of data cannot be measured. Like the nominal scale data, ordinal scale data cannot be used in calculations.

Data that is measured using the **interval scale** is similar to ordinal level data because it has a definite ordering but there is a difference between data. The differences between interval scale data can be measured though the data does not have a starting point.

Temperature scales like Celsius (C) and Fahrenheit (F) are measured by using the interval scale. In both temperature measurements, 40° is equal to 100° minus 60° . Differences make sense. But 0 degrees does not because, in both scales, 0 is not the absolute lowest temperature. Temperatures like -10° F and -15° C exist and are colder than 0.

Interval level data can be used in calculations, but one type of comparison cannot be done. 80° C is not four times as hot as 20° C (nor is 80° F four times as hot as 20° F). There is no meaning to the ratio of 80 to 20 (or four to one).

Data that is measured using the **ratio scale** takes care of the ratio problem and gives you the most information. Ratio scale data is like interval scale data, but it has a 0 point and ratios can be calculated. For example, four multiple choice statistics final exam scores are 80, 68, 20 and 92 (out of a possible 100 points). The exams are machine-graded.

The data can be put in order from lowest to highest: 20, 68, 80, 92.

The differences between the data have meaning. The score 92 is more than the score 68 by 24 points. Ratios can be calculated. The smallest score is 0. So 80 is four times 20. The score of 80 is four times better than the score of 20.

Frequency

Twenty students were asked how many hours they worked per day. Their responses, in hours, are as follows: 5; 6; 3; 3; 2; 4; 7; 5; 2; 3; 5; 6; 5; 4; 4; 3; 5; 2; 5; 3.

Table 1.4.5 lists the different data values in ascending order and their frequencies.

Data value	Frequency

Data value	Frequency
2	3
3	5
4	3
5	6
6	2
7	1

Table 1.5 Frequency Table of Student Work Hours

A **frequency** is the number of times a value of the data occurs. According to [Table 1.4.5](#), there are three students who work two hours, five students who work three hours, and so on. The sum of the values in the frequency column, 20, represents the total number of students included in the sample.

A **relative frequency** is the ratio (fraction or proportion) of the number of times a value of the data occurs in the set of all outcomes to the total number of outcomes. To find the relative frequencies, divide each frequency by the total number of students in the sample—in this case, 20. Relative frequencies can be written as fractions, percents, or decimals.

Data value	Frequency	Relative frequency
2	3	$\frac{3}{20}$ or 0.15
3	5	$\frac{5}{20}$ or 0.25
4	3	$\frac{3}{20}$ or 0.15
5	6	$\frac{6}{20}$ or 0.30
6	2	$\frac{2}{20}$ or 0.10
7	1	$\frac{1}{20}$ or 0.05

Table 1.6 Frequency Table of Student Work Hours with Relative Frequencies

The sum of the values in the relative frequency column of [Table 1.4.6](#) is $\frac{20}{20}$, or 1.

Cumulative relative frequency is the accumulation of the previous relative frequencies. To find the cumulative relative frequencies, add all the previous relative frequencies to the relative frequency for the current row, as shown in [Table 1.4.7](#).

Data value	Frequency	Relative frequency	Cumulative relative frequency
2	3	$\frac{3}{20}$ or 0.15	0.15
3	5	$\frac{5}{20}$ or 0.25	$0.15 + 0.25 = 0.40$
4	3	$\frac{3}{20}$ or 0.15	$0.40 + 0.15 = 0.55$
5	6	$\frac{6}{20}$ or 0.30	$0.55 + 0.30 = 0.85$
6	2	$\frac{2}{20}$ or 0.10	$0.85 + 0.10 = 0.95$
7	1	$\frac{1}{20}$ or 0.05	$0.95 + 0.05 = 1.00$

Table 1.7 Frequency Table of Student Work Hours with Relative and Cumulative Relative Frequencies

The last entry of the cumulative relative frequency column is one, indicating that one hundred percent of the data has been accumulated.

NOTE

Because of rounding, the relative frequency column may not always sum to one, and the last entry in the cumulative relative frequency column may not be one. However, they each should be close to one.

Table 1.4.8 represents the heights, in inches, of a sample of 100 male semiprofessional soccer players.

Heights (inches)	Frequency	Relative frequency	Cumulative relative frequency
59.95–61.95	5	$\frac{5}{100} = 0.05$	0.05
61.95–63.95	3	$\frac{3}{100} = 0.03$	$0.05 + 0.03 = 0.08$
63.95–65.95	15	$\frac{15}{100} = 0.15$	$0.08 + 0.15 = 0.23$
65.95–67.95	40	$\frac{40}{100} = 0.40$	$0.23 + 0.40 = 0.63$
67.95–69.95	17	$\frac{17}{100} = 0.17$	$0.63 + 0.17 = 0.80$
69.95–71.95	12	$\frac{12}{100} = 0.12$	$0.80 + 0.12 = 0.92$
71.95–73.95	7	$\frac{7}{100} = 0.07$	$0.92 + 0.07 = 0.99$
73.95–75.95	1	$\frac{1}{100} = 0.01$	$0.99 + 0.01 = 1.00$
	Total = 100	Total = 1.00	

Table 1.8 Frequency Table of Soccer Player Height

The data in this table have been **grouped** into the following intervals:

- 59.95 to 61.95 inches
- 61.95 to 63.95 inches
- 63.95 to 65.95 inches
- 65.95 to 67.95 inches
- 67.95 to 69.95 inches
- 69.95 to 71.95 inches
- 71.95 to 73.95 inches
- 73.95 to 75.95 inches

In this sample, there are **five** players whose heights fall within the interval 59.95–61.95 inches, **three** players whose heights fall within the interval 61.95–63.95 inches, **15** players whose heights fall within the interval 63.95–65.95 inches, **40** players whose heights fall within the interval 65.95–67.95 inches, **17** players whose heights fall within the interval 67.95–69.95 inches, **12** players whose heights fall within the interval 69.95–71.95, **seven** players whose heights fall within the interval 71.95–73.95, and **one** player whose heights fall within the interval 73.95–75.95. All heights fall between the endpoints of an interval and not at the endpoints.

Example 1.4.14

From Table 1.4.8, find the percentage of heights that are less than 65.95 inches.

Exercise 1.4.14

Table 1.4.9 shows the amount, in inches, of annual rainfall in a sample of towns.

Table 1.4.9

Rainfall (inches)	Frequency	Relative frequency	Cumulative relative frequency
2.95–4.97	6	$\frac{6}{50} = 0.12$	0.12
4.97–6.99	7	$\frac{7}{50} = 0.14$	$0.12 + 0.14 = 0.26$
6.99–9.01	15	$\frac{15}{50} = 0.30$	$0.26 + 0.30 = 0.56$

Rainfall (inches)	Frequency	Relative frequency	Cumulative relative frequency
9.01–11.03	8	$\frac{8}{50} = 0.16$	$0.56 + 0.16 = 0.72$
11.03–13.05	9	$\frac{9}{50} = 0.18$	$0.72 + 0.18 = 0.90$
13.05–15.07	5	$\frac{5}{50} = 0.10$	$0.90 + 0.10 = 1.00$
	Total = 50	Total = 1.00	

From Table 1.4.9, find the percentage of rainfall that is less than 9.01 inches.

Example 1.4.15

From Table 1.4.8, find the percentage of heights that fall between 61.95 and 65.95 inches.

Answer

Solution 1.15

Add the relative frequencies in the second and third rows: $0.03 + 0.15 = 0.18$ or 18%.

Exercise 1.4.15

From Table 1.4.9, find the percentage of rainfall that is between 6.99 and 13.05 inches.

Example 1.4.16

Use the heights of the 100 male semiprofessional soccer players in Table 1.4.8. Fill in the blanks and check your answers.

- The percentage of heights that are from 67.95 to 71.95 inches is: ____.
- The percentage of heights that are from 67.95 to 73.95 inches is: ____.
- The percentage of heights that are more than 65.95 inches is: ____.
- The number of players in the sample who are between 61.95 and 71.95 inches tall is: ____.
- What kind of data are the heights?
- Describe how you could gather this data (the heights) so that the data are characteristic of all male semiprofessional soccer players.

Remember, you **count frequencies**. To find the relative frequency, divide the frequency by the total number of data values. To find the cumulative relative frequency, add all of the previous relative frequencies to the relative frequency for the current row.

Answer

Solution 1.16

- 29%
- 36%
- 77%
- 87
- quantitative continuous
- get rosters from each team and choose a simple random sample from each

Example 1.4.17

Nineteen people were asked how many miles, to the nearest mile, they commute to work each day. The data are as follows: 2; 5; 7; 3; 2; 10; 18; 15; 20; 7; 10; 18; 5; 12; 13; 12; 4; 5; 10. Table 1.4.10 was produced:

Table 1.4.10 Frequency of Commuting Distances

Data	Frequency	Relative frequency	Cumulative relative frequency

Data	Frequency	Relative frequency	Cumulative relative frequency
3	3	$\frac{3}{19}$	0.1579
4	1	$\frac{1}{19}$	0.2105
5	3	$\frac{3}{19}$	0.1579
7	2	$\frac{2}{19}$	0.2632
10	3	$\frac{4}{19}$	0.4737
12	2	$\frac{2}{19}$	0.7895
13	1	$\frac{1}{19}$	0.8421
15	1	$\frac{1}{19}$	0.8948
18	1	$\frac{1}{19}$	0.9474
20	1	$\frac{1}{19}$	1.0000

1. Is the table correct? If it is not correct, what is wrong?
2. True or False: Three percent of the people surveyed commute three miles. If the statement is not correct, what should it be? If the table is incorrect, make the corrections.
3. What fraction of the people surveyed commute five or seven miles?
4. What fraction of the people surveyed commute 12 miles or more? Less than 12 miles? Between five and 13 miles (not including five and 13 miles)?

Answer

Solution 1.17

1. No. The frequency column sums to 18, not 19. Not all cumulative relative frequencies are correct.
2. False. The frequency for three miles should be one; for two miles (left out), two. The cumulative relative frequency column should read: 0.1052, 0.1579, 0.2105, 0.3684, 0.4737, 0.6316, 0.7368, 0.7895, 0.8421, 0.9474, 1.0000.
3. $\frac{5}{19}$
4. $(\frac{7}{19}, \frac{12}{19}, \frac{7}{19})$

Exercise 1.4.17

Table 1.4.9 represents the amount, in inches, of annual rainfall in a sample of towns. What fraction of towns surveyed get between 11.03 and 13.05 inches of rainfall each year?

Example 1.4.18

Table 1.4.11 contains the total number of deaths worldwide as a result of earthquakes for the period from 2000 to 2012.

Year	Total number of deaths
2000	231
2001	21,357
2002	11,685
2003	33,819
2004	228,802
2005	88,003
2006	6,605
2007	712

Year	Total number of deaths
2008	88,011
2009	1,790
2010	320,120
2011	21,953
2012	768
Total	823,856

Table1.11

Answer the following questions.

1. What is the frequency of deaths measured from 2006 through 2009?
2. What percentage of deaths occurred after 2009?
3. What is the relative frequency of deaths that occurred in 2003 or earlier?
4. What is the percentage of deaths that occurred in 2004?
5. What kind of data are the numbers of deaths?
6. The Richter scale is used to quantify the energy produced by an earthquake. Examples of Richter scale numbers are 2.3, 4.0, 6.1, and 7.0. What kind of data are these numbers?

Answer

Solution 1.18

1. 97,118 (11.8%)
2. 41.6%
3. $67,092/823,356$ or 0.081 or 8.1 %
4. 27.8%
5. Quantitative discrete
6. Quantitative continuous

Exercise 1.4.18

Table 1.4.12 contains the total number of fatal motor vehicle traffic crashes in the United States for the period from 1994 to 2011.

Year	Total number of crashes	Year	Total number of crashes
1994	36,254	2004	38,444
1995	37,241	2005	39,252
1996	37,494	2006	38,648
1997	37,324	2007	37,435
1998	37,107	2008	34,172
1999	37,140	2009	30,862
2000	37,526	2010	30,296
2001	37,862	2011	29,757
2002	38,491	Total	653,782
2003	38,477		

Table1.12

Answer the following questions.

- a. What is the frequency of deaths measured from 2000 through 2004?
- b. What percentage of deaths occurred after 2006?
- c. What is the relative frequency of deaths that occurred in 2000 or before?
- d. What is the percentage of deaths that occurred in 2011?
- e. What is the cumulative relative frequency for 2006? Explain what this number tells you about the data.

This page titled [1.4: Levels of Measurement](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [1.3: Levels of Measurement](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.5: Experimental Design and Ethics

Does aspirin reduce the risk of heart attacks? Is one brand of fertilizer more effective at growing roses than another? Is fatigue as dangerous to a driver as the influence of alcohol? Questions like these are answered using randomized experiments. In this module, you will learn important aspects of experimental design. Proper study design ensures the production of reliable, accurate data.

The purpose of an experiment is to investigate the relationship between two variables. When one variable causes change in another, we call the first variable the **independent variable** or **explanatory variable**. The affected variable is called the **dependent variable** or **response variable**: stimulus, response. In a randomized experiment, the researcher manipulates values of the explanatory variable and measures the resulting changes in the response variable. The different values of the explanatory variable are called **treatments**. An **experimental unit** is a single object or individual to be measured.

You want to investigate the effectiveness of vitamin E in preventing disease. You recruit a group of subjects and ask them if they regularly take vitamin E. You notice that the subjects who take vitamin E exhibit better health on average than those who do not. Does this prove that vitamin E is effective in disease prevention? It does not. There are many differences between the two groups compared in addition to vitamin E consumption. People who take vitamin E regularly often take other steps to improve their health: exercise, diet, other vitamin supplements, choosing not to smoke. Any one of these factors could be influencing health. As described, this study does not prove that vitamin E is the key to disease prevention.

Additional variables that can cloud a study are called **lurking variables**. In order to prove that the explanatory variable is causing a change in the response variable, it is necessary to isolate the explanatory variable. The researcher must design her experiment in such a way that there is only one difference between groups being compared: the planned treatments. This is accomplished by the **random assignment** of experimental units to treatment groups. When subjects are assigned treatments randomly, all of the potential lurking variables are spread equally among the groups. At this point the only difference between groups is the one imposed by the researcher. Different outcomes measured in the response variable, therefore, must be a direct result of the different treatments. In this way, an experiment can prove a cause-and-effect connection between the explanatory and response variables.

The power of suggestion can have an important influence on the outcome of an experiment. Studies have shown that the expectation of the study participant can be as important as the actual medication. In one study of performance-enhancing drugs, researchers noted:

Results showed that believing one had taken the substance resulted in [performance] times almost as fast as those associated with consuming the drug itself. In contrast, taking the drug without knowledge yielded no significant performance increment. (McClung, M. Collins, D. "Because I know it will!": placebo effects of an ergogenic aid on athletic performance. Journal of Sport & Exercise Psychology. 2007 Jun. 29(3):382-94. Web. April 30, 2013.)

When participation in a study prompts a physical response from a participant, it is difficult to isolate the effects of the explanatory variable. To counter the power of suggestion, researchers set aside one treatment group as a **control group**. This group is given a **placebo** treatment—a treatment that cannot influence the response variable. The control group helps researchers balance the effects of being in an experiment with the effects of the active treatments. Of course, if you are participating in a study and you know that you are receiving a pill which contains no actual medication, then the power of suggestion is no longer a factor. **Blinding** in a randomized experiment preserves the power of suggestion. When a person involved in a research study is blinded, he does not know who is receiving the active treatment(s) and who is receiving the placebo treatment. A **double-blind experiment** is one in which both the subjects and the researchers involved with the subjects are blinded.

Example 1.5.19

The Smell & Taste Treatment and Research Foundation conducted a study to investigate whether smell can affect learning. Subjects completed mazes multiple times while wearing masks. They completed the pencil and paper mazes three times wearing floral-scented masks, and three times with unscented masks. Participants were assigned at random to wear the floral mask during the first three trials or during the last three trials. For each trial, researchers recorded the time it took to complete the maze and the subject's impression of the mask's scent: positive, negative, or neutral.

- Describe the explanatory and response variables in this study.
- What are the treatments?
- Identify any lurking variables that could interfere with this study.
- Is it possible to use blinding in this study?

Answer

Solution 1.19

The explanatory variable is scent, and the response variable is the time it takes to complete the maze. There are two treatments: a floral-scented mask and an unscented mask. All subjects experienced both treatments. The order of treatments was randomly assigned so there were no differences between the treatment groups. Random assignment eliminates the problem of lurking variables. Subjects will clearly know whether they can smell flowers or not, so subjects cannot be blinded in this study. Researchers timing the mazes can be blinded, though. The researcher who is observing a subject will not know which mask is being worn.

This page titled [1.5: Experimental Design and Ethics](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [1.4: Experimental Design and Ethics](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.6: Chapter Key Terms

Average

also called mean or arithmetic mean; a number that describes the central tendency of the data

Blinding

not telling participants which treatment a subject is receiving

Categorical Variable

variables that take on values that are names or labels

Cluster Sampling

a method for selecting a random sample and dividing the population into groups (clusters); use simple random sampling to select a set of clusters. Every individual in the chosen clusters is included in the sample.

Continuous Random Variable

a random variable (RV) whose outcomes are measured; the height of trees in the forest is a continuous RV.

Control Group

a group in a randomized experiment that receives an inactive treatment but is otherwise managed exactly as the other groups

Convenience Sampling

a nonrandom method of selecting a sample; this method selects individuals that are easily accessible and may result in biased data.

Cumulative Relative Frequency

The term applies to an ordered set of observations from smallest to largest. The cumulative relative frequency is the sum of the relative frequencies for all values that are less than or equal to the given value.

Data

a set of observations (a set of possible outcomes); most data can be put into two groups: **qualitative** (an attribute whose value is indicated by a label) or **quantitative** (an attribute whose value is indicated by a number). Quantitative data can be separated into two subgroups: **discrete** and **continuous**. Data is discrete if it is the result of counting (such as the number of students of a given ethnic group in a class or the number of books on a shelf). Data is continuous if it is the result of measuring (such as distance traveled or weight of luggage)

Discrete Random Variable

a random variable (RV) whose outcomes are counted

Double-blinding

the act of blinding both the subjects of an experiment and the researchers who work with the subjects

Experimental Unit

any individual or object to be measured

Explanatory Variable

the **independent variable** in an experiment; the value controlled by researchers

Frequency

the number of times a value of the data occurs

Informed Consent

Any human subject in a research study must be cognizant of any risks or costs associated with the study. The subject has the right to know the nature of the treatments included in the study, their potential risks, and their potential benefits. Consent must

be given freely by an informed, fit participant.

Institutional Review Board

a committee tasked with oversight of research programs that involve human subjects

Lurking Variable

a variable that has an effect on a study even though it is neither an explanatory variable nor a response variable

Mathematical Models

a description of a phenomenon using mathematical concepts, such as equations, inequalities, distributions, etc.

Nonsampling Error

an issue that affects the reliability of sampling data other than natural variation; it includes a variety of human errors including poor study design, biased sampling methods, inaccurate information provided by study participants, data entry errors, and poor analysis.

Numerical Variable

variables that take on values that are indicated by numbers

Observational Study

a study in which the independent variable is not manipulated by the researcher

Parameter

a number that is used to represent a population characteristic and that generally cannot be determined easily

Placebo

an inactive treatment that has no real effect on the explanatory variable

Population

all individuals, objects, or measurements whose properties are being studied

Probability

a number between zero and one, inclusive, that gives the likelihood that a specific event will occur

Proportion

the number of successes divided by the total number in the sample

Qualitative Data

See [Data](#).

Quantitative Data

See [Data](#).

Random Assignment

the act of organizing experimental units into treatment groups using random methods

Random Sampling

a method of selecting a sample that gives every member of the population an equal chance of being selected.

Relative Frequency

the ratio of the number of times a value of the data occurs in the set of all outcomes to the number of all outcomes to the total number of outcomes

Representative Sample

a subset of the population that has the same characteristics as the population

Response Variable

the **dependent variable** in an experiment; the value that is measured for change at the end of an experiment

Sample

a subset of the population studied

Sampling Bias

not all members of the population are equally likely to be selected

Sampling Error

the natural variation that results from selecting a sample to represent a larger population; this variation decreases as the sample size increases, so selecting larger samples reduces sampling error.

Sampling with Replacement

Once a member of the population is selected for inclusion in a sample, that member is returned to the population for the selection of the next individual.

Sampling without Replacement

A member of the population may be chosen for inclusion in a sample only once. If chosen, the member is not returned to the population before the next selection.

Simple Random Sampling

a straightforward method for selecting a random sample; give each member of the population a number. Use a random number generator to select a set of labels. These randomly selected labels identify the members of your sample.

Statistic

a numerical characteristic of the sample; a statistic estimates the corresponding population parameter.

Statistical Models

a description of a phenomenon using probability distributions that describe the expected behavior of the phenomenon and the variability in the expected observations.

Stratified Sampling

a method for selecting a random sample used to ensure that subgroups of the population are represented adequately; divide the population into groups (strata). Use simple random sampling to identify a proportionate number of individuals from each stratum.

Survey

a study in which data is collected as reported by individuals.

Systematic Sampling

a method for selecting a random sample; list the members of the population. Use simple random sampling to select a starting point in the population. Let $k = (\text{number of individuals in the population}) / (\text{number of individuals needed in the sample})$. Choose every k th individual in the list starting with the one that was randomly selected. If necessary, return to the beginning of the population list to complete your sample.

Treatments

different values or components of the explanatory variable applied in an experiment

Variable

a characteristic of interest for each person or object in a population

This page titled [1.6: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

1.7: Chapter References

Definitions of Statistics, Probability, and Key Terms

- The Data and Story Library, lib.stat.cmu.edu/DASL/Stories...stDummies.html (accessed May 1, 2013).

Data, Sampling, and Variation in Data and Sampling

- Gallup-Healthways Well-Being Index. <http://www.well-beingindex.com/default.asp> (accessed May 1, 2013).
- Gallup-Healthways Well-Being Index. <http://www.well-beingindex.com/methodology.asp> (accessed May 1, 2013).
- Gallup-Healthways Well-Being Index. <http://www.gallup.com/poll/146822/ga...questions.aspx> (accessed May 1, 2013).
- Data from www.bookofodds.com/Relationsh...-the-President
- Dominic Lusinchi, “President’ Landon and the 1936 *Literary Digest* Poll: Were Automobile and Telephone Owners to Blame?” *Social Science History* 36, no. 1: 23-54 (2012), ssh.dukejournals.org/content/36/1/23.abstract (accessed May 1, 2013).
- “The Literary Digest Poll,” Virtual Laboratories in Probability and Statistics <http://www.math.uah.edu/stat/data/LiteraryDigest.html> (accessed May 1, 2013).
- “Gallup Presidential Election Trial-Heat Trends, 1936–2008,” Gallup Politics <http://www.gallup.com/poll/110548/ga...9362004.aspx#4> (accessed May 1, 2013).
- The Data and Story Library, lib.stat.cmu.edu/DASL/Datafiles/USCrime.html (accessed May 1, 2013).
- LBCC Distance Learning (DL) program data in 2010-2011, <http://de.lbcc.edu/reports/2010-11/f...hts.html#focus> (accessed May 1, 2013).
- Data from San Jose Mercury News

Levels of Measurement

- “State & County QuickFacts,” U.S. Census Bureau. quickfacts.census.gov/qfd/download_data.html (accessed May 1, 2013).
- “State & County QuickFacts: Quick, easy access to facts about people, business, and geography,” U.S. Census Bureau. quickfacts.census.gov/qfd/index.html (accessed May 1, 2013).
- “Table 5: Direct hits by mainland United States Hurricanes (1851-2004),” National Hurricane Center, <http://www.nhc.noaa.gov/gifs/table5.gif> (accessed May 1, 2013).
- “Levels of Measurement,” infinity.cos.edu/faculty/wood...ata_Levels.htm (accessed May 1, 2013).
- Courtney Taylor, “Levels of Measurement,” [about.com, http://statistics.about.com/od/Helpa...easurement.htm](http://statistics.about.com/od/Helpa...easurement.htm) (accessed May 1, 2013).
- David Lane. “Levels of Measurement,” [Connexions, http://cnx.org/content/m10809/latest/](http://cnx.org/content/m10809/latest/) (accessed May 1, 2013).

Experimental Design and Ethics

- “Vitamin E and Health,” Nutrition Source, Harvard School of Public Health, <http://www.hsph.harvard.edu/nutritio...rce/vitamin-e/> (accessed May 1, 2013).
- Stan Reents. “Don’t Underestimate the Power of Suggestion,” athletinme.com, www.athletinme.com/ArticleView.aspx?id=1053 (accessed May 1, 2013).
- Ankita Mehta. “Daily Dose of Aspiring Helps Reduce Heart Attacks: Study,” *International Business Times*, July 21, 2011. Also available online at <http://www.ibtimes.com/daily-dose-as...s-study-300443> (accessed May 1, 2013).
- The Data and Story Library, lib.stat.cmu.edu/DASL/Stories...dLearning.html (accessed May 1, 2013).
- M.L. Jacskon et al., “Cognitive Components of Simulated Driving Performance: Sleep Loss effect and Predictors,” *Accident Analysis and Prevention Journal*, Jan no. 50 (2013), <http://www.ncbi.nlm.nih.gov/pubmed/22721550> (accessed May 1, 2013).
- “Earthquake Information by Year,” U.S. Geological Survey. <http://earthquake.usgs.gov/earthquak...archives/year/> (accessed May 1, 2013).
- “Fatality Analysis Report Systems (FARS) Encyclopedia,” National Highway Traffic and Safety Administration. <http://www-fars.nhtsa.dot.gov/Main/index.aspx> (accessed May 1, 2013).
- Data from www.businessweek.com (accessed May 1, 2013).
- Data from www.forbes.com (accessed May 1, 2013).
- “America’s Best Small Companies,” <http://www.forbes.com/best-small-companies/list/> (accessed May 1, 2013).
- U.S. Department of Health and Human Services, Code of Federal Regulations Title 45 Public Welfare Department of Health and Human Services Part 46 Protection of Human Subjects revised January 15, 2009. Section 46.111:Criteria for IRB Approval of Research.

- “April 2013 Air Travel Consumer Report,” U.S. Department of Transportation, April 11 (2013), <http://www.dot.gov/airconsumer/april...onsumer-report> (accessed May 1, 2013).
- Lori Alden, “Statistics can be Misleading,” econoclass.com, <http://www.econoclass.com/misleadingstats.html> (accessed May 1, 2013).
- Maria de los A. Medina, “Ethics in Statistics,” Based on “Building an Ethics Module for Business, Science, and Engineering Students” by Jose A. Cruz-Cruz and William Frey, Connexions, <http://cnx.org/content/m15555/latest/> (accessed May 1, 2013).

This page titled [1.7: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **1.6: Chapter References** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.H: Sampling and Data (Homework)

1.1 Definitions of Statistics, Probability, and Key Terms

For each of the following eight exercises, identify: a. the population, b. the sample, c. the parameter, d. the statistic, e. the variable, and f. the data. Give examples where appropriate.

1.
A fitness center is interested in the mean amount of time a client exercises in the center each week.
2.
Ski resorts are interested in the mean age that children take their first ski and snowboard lessons. They need this information to plan their ski classes optimally.
3.
A cardiologist is interested in the mean recovery period of her patients who have had heart attacks.
4.
Insurance companies are interested in the mean health costs each year of their clients, so that they can determine the costs of health insurance.
5.
A politician is interested in the proportion of voters in his district who think he is doing a good job.
6.
A marriage counselor is interested in the proportion of clients she counsels who stay married.
7.
Political pollsters may be interested in the proportion of people who will vote for a particular cause.
8.
A marketing company is interested in the proportion of people who will buy a particular product.
Use the following information to answer the next three exercises: A Lake Tahoe Community College instructor is interested in the mean number of days Lake Tahoe Community College math students are absent from class during a quarter.
9.
What is the population she is interested in?
 - a. all Lake Tahoe Community College students
 - b. all Lake Tahoe Community College English students
 - c. all Lake Tahoe Community College students in her classes
 - d. all Lake Tahoe Community College math students
10.
Consider the following:
 X = number of days a Lake Tahoe Community College math student is absent
In this case, X is an example of a:
 - a. variable.
 - b. population.
 - c. statistic.
 - d. data.
11.
The instructor's sample produces a mean number of days absent of 3.5 days. This value is an example of a:
 - a. parameter.

- b. data.
- c. statistic.
- d. variable.

1.2 Data, Sampling, and Variation in Data and Sampling

For the following exercises, identify the type of data that would be used to describe a response (quantitative discrete, quantitative continuous, or qualitative), and give an example of the data.

12.

number of tickets sold to a concert

13.

percent of body fat

14.

favorite baseball team

15.

time in line to buy groceries

16.

number of students enrolled at Evergreen Valley College

17.

most-watched television show

18.

brand of toothpaste

19.

distance to the closest movie theatre

20.

age of executives in Fortune 500 companies

21.

number of competing computer spreadsheet software packages

Use the following information to answer the next two exercises: A study was done to determine the age, number of times per week, and the duration (amount of time) of resident use of a local park in San Jose. The first house in the neighborhood around the park was selected randomly and then every 8th house in the neighborhood around the park was interviewed.

22.

“Number of times per week” is what type of data?

- a. qualitative (categorical)
- b. quantitative discrete
- c. quantitative continuous

23.

“Duration (amount of time)” is what type of data?

- a. qualitative (categorical)
- b. quantitative discrete
- c. quantitative continuous

24.

Airline companies are interested in the consistency of the number of babies on each flight, so that they have adequate safety equipment. Suppose an airline conducts a survey. Over Thanksgiving weekend, it surveys six flights from Boston to Salt Lake City

to determine the number of babies on the flights. It determines the amount of safety equipment needed by the result of that study.

- a. Using complete sentences, list three things wrong with the way the survey was conducted.
- b. Using complete sentences, list three ways that you would improve the survey if it were to be repeated.

25.

Suppose you want to determine the mean number of students per statistics class in your state. Describe a possible sampling method in three to five complete sentences. Make the description detailed.

26.

Suppose you want to determine the mean number of cans of soda drunk each month by students in their twenties at your school. Describe a possible sampling method in three to five complete sentences. Make the description detailed.

27.

List some practical difficulties involved in getting accurate results from a telephone survey.

28.

List some practical difficulties involved in getting accurate results from a mailed survey.

29.

With your classmates, brainstorm some ways you could overcome these problems if you needed to conduct a phone or mail survey.

30.

The instructor takes her sample by gathering data on five randomly selected students from each Lake Tahoe Community College math class. The type of sampling she used is

- a. cluster sampling
- b. stratified sampling
- c. simple random sampling
- d. convenience sampling

31.

A study was done to determine the age, number of times per week, and the duration (amount of time) of residents using a local park in San Jose. The first house in the neighborhood around the park was selected randomly and then every eighth house in the neighborhood around the park was interviewed. The sampling method was:

- a. simple random
- b. systematic
- c. stratified
- d. cluster

32.

Name the sampling method used in each of the following situations:

- a. A woman in the airport is handing out questionnaires to travelers asking them to evaluate the airport's service. She does not ask travelers who are hurrying through the airport with their hands full of luggage, but instead asks all travelers who are sitting near gates and not taking naps while they wait.
- b. A teacher wants to know if her students are doing homework, so she randomly selects rows two and five and then calls on all students in row two and all students in row five to present the solutions to homework problems to the class.
- c. The marketing manager for an electronics chain store wants information about the ages of its customers. Over the next two weeks, at each store location, 100 randomly selected customers are given questionnaires to fill out asking for information about age, as well as about other variables of interest.
- d. The librarian at a public library wants to determine what proportion of the library users are children. The librarian has a tally sheet on which she marks whether books are checked out by an adult or a child. She records this data for every fourth patron who checks out books.
- e. A political party wants to know the reaction of voters to a debate between the candidates. The day after the debate, the party's polling staff calls 1,200 randomly selected phone numbers. If a registered voter answers the phone or is available to come to the

phone, that registered voter is asked whom he or she intends to vote for and whether the debate changed his or her opinion of the candidates.

33.

A “random survey” was conducted of 3,274 people of the “microprocessor generation” (people born since 1971, the year the microprocessor was invented). It was reported that 48% of those individuals surveyed stated that if they had \$2,000 to spend, they would use it for computer equipment. Also, 66% of those surveyed considered themselves relatively savvy computer users.

- Do you consider the sample size large enough for a study of this type? Why or why not?
- Based on your “gut feeling,” do you believe the percents accurately reflect the U.S. population for those individuals born since 1971? If not, do you think the percents of the population are actually higher or lower than the sample statistics? Why?
Additional information: The survey, reported by Intel Corporation, was filled out by individuals who visited the Los Angeles Convention Center to see the Smithsonian Institute's road show called “America’s Smithsonian.”
- With this additional information, do you feel that all demographic and ethnic groups were equally represented at the event? Why or why not?
- With the additional information, comment on how accurately you think the sample statistics reflect the population parameters.

34.

The Well-Being Index is a survey that follows trends of U.S. residents on a regular basis. There are six areas of health and wellness covered in the survey: Life Evaluation, Emotional Health, Physical Health, Healthy Behavior, Work Environment, and Basic Access. Some of the questions used to measure the Index are listed below.

Identify the type of data obtained from each question used in this survey: qualitative(categorical), quantitative discrete, or quantitative continuous.

- Do you have any health problems that prevent you from doing any of the things people your age can normally do?
- During the past 30 days, for about how many days did poor health keep you from doing your usual activities?
- In the last seven days, on how many days did you exercise for 30 minutes or more?
- Do you have health insurance coverage?

35.

In advance of the 1936 Presidential Election, a magazine titled Literary Digest released the results of an opinion poll predicting that the republican candidate Alf Landon would win by a large margin. The magazine sent post cards to approximately 10,000,000 prospective voters. These prospective voters were selected from the subscription list of the magazine, from automobile registration lists, from phone lists, and from club membership lists. Approximately 2,300,000 people returned the postcards.

- Think about the state of the United States in 1936. Explain why a sample chosen from magazine subscription lists, automobile registration lists, phone books, and club membership lists was not representative of the population of the United States at that time.
- What effect does the low response rate have on the reliability of the sample?
- Are these problems examples of sampling error or nonsampling error?
- During the same year, George Gallup conducted his own poll of 30,000 prospective voters. These researchers used a method they called “quota sampling” to obtain survey answers from specific subsets of the population. Quota sampling is an example of which sampling method described in this module?

36.

Crime-related and demographic statistics for 47 US states in 1960 were collected from government agencies, including the FBI's *Uniform Crime Report*. One analysis of this data found a strong connection between education and crime indicating that higher levels of education in a community correspond to higher crime rates.

Which of the potential problems with samples discussed in [Example 1.H.4](#) could explain this connection?

37.

YouPolls is a website that allows anyone to create and respond to polls. One question posted April 15 asks:

“Do you feel happy paying your taxes when members of the Obama administration are allowed to ignore their tax liabilities?” (lastbaldeagle. 2013. On Tax Day, House to Call for Firing Federal Workers Who Owe Back Taxes. Opinion poll posted online at: <http://www.youpolls.com/details.aspx?id=12328> (accessed May 1, 2013).)

As of April 25, 11 people responded to this question. Each participant answered “NO!”

Which of the potential problems with samples discussed in this module could explain this connection?

38.

A scholarly article about response rates begins with the following quote:

“Declining contact and cooperation rates in random digit dial (RDD) national telephone surveys raise serious concerns about the validity of estimates drawn from such research.” (Scott Keeter et al., “Gauging the Impact of Growing Nonresponse on Estimates from a National RDD Telephone Survey,” *Public Opinion Quarterly* 70 no. 5 (2006), <http://poq.oxfordjournals.org/content/70/5/759.full> (accessed May 1, 2013).)

The Pew Research Center for People and the Press admits:

“The percentage of people we interview – out of all we try to interview – has been declining over the past decade or more.” (Frequently Asked Questions, Pew Research Center for the People & the Press, <http://www.people-press.org/methodol...wer-your-polls> (accessed May 1, 2013).)

- What are some reasons for the decline in response rate over the past decade?
- Explain why researchers are concerned with the impact of the declining response rate on public opinion polls.

1.3 Levels of Measurement

39.

Fifty part-time students were asked how many courses they were taking this term. The (incomplete) results are shown below:

# of courses	Frequency	Relative frequency	Cumulative relative frequency
1	30	0.6	
2	15		
3			

Table1.13 Part-time Student Course Loads

- Fill in the blanks in Table 1.H. 13.
- What percent of students take exactly two courses?
- What percent of students take one or two courses?

40.

Sixty adults with gum disease were asked the number of times per week they used to floss before their diagnosis. The (incomplete) results are shown in **Table 1.H. 14**.

# flossing per week	Frequency	Relative frequency	Cumulative relative frequency
0	27	0.4500	
1	18		
3			0.9333
6	3	0.0500	
7	1	0.0167	

Table1.14 Flossing Frequency for Adults with Gum Disease

- Fill in the blanks in Table 1.H. 14.
- What percent of adults flossed six times per week?
- What percent flossed at most three times per week?

41.

Nineteen immigrants to the U.S were asked how many years, to the nearest year, they have lived in the U.S. The data are as follows: 2;5; 7; 2; 2; 10; 20; 15; 0; 7; 0; 20; 5; 12; 15; 12; 4; 5; 10 .

Table 1.H. 15 was produced.

Table 1.H. 15 Frequency of Immigrant Survey Responses

Data	Frequency	Relative frequency	Cumulative relative frequency
0	2	219219	0.1053
2	3	319319	0.2632
4	1	119119	0.3158
5	3	319319	0.4737
7	2	219219	0.5789
10	2	219219	0.6842
12	2	219219	0.7895
15	1	119119	0.8421
20	1	119119	1.0000

- Fix the errors in Table 1.H. 15. Also, explain how someone might have arrived at the incorrect number(s).
- Explain what is wrong with this statement: “47 percent of the people surveyed have lived in the U.S. for 5 years.”
- Fix the statement in **b** to make it correct.
- What fraction of the people surveyed have lived in the U.S. five or seven years?
- What fraction of the people surveyed have lived in the U.S. at most 12 years?
- What fraction of the people surveyed have lived in the U.S. fewer than 12 years?
- What fraction of the people surveyed have lived in the U.S. from five to 20 years, inclusive?

42.

How much time does it take to travel to work? **Table 1.H. 16** shows the mean commute time by state for workers at least 16 years old who are not working at home. Find the mean travel time, and round off the answer properly.

Table 1.H. 16

24.0	24.3	25.9	18.9	27.5	17.9	21.8	20.9	16.7	27.3
18.2	24.7	20.0	22.6	23.9	18.0	31.4	22.3	24.0	25.5
24.7	24.6	28.1	24.9	22.6	23.6	23.4	25.7	24.8	25.5
21.2	25.7	23.1	23.0	23.9	26.0	16.3	23.1	21.4	21.5
27.0	27.0	18.6	31.7	23.3	30.1	22.9	23.3	21.7	18.6

43.

Forbes magazine published data on the best small firms in 2012. These were firms which had been publicly traded for at least a year, have a stock price of at least \$5 per share, and have reported annual revenue between \$5 million and \$1 billion. **Table 1.H. 17** shows the ages of the chief executive officers for the first 60 ranked firms.

Table 1.H. 17

Age	Frequency	Relative frequency	Cumulative relative frequency
40–44	3		
45–49	11		
50–54	13		
55–59	16		
60–64	10		
65–69	6		
70–74	1		

- What is the frequency for CEO ages between 54 and 65?
- What percentage of CEOs are 65 years or older?
- What is the relative frequency of ages under 50?
- What is the cumulative relative frequency for CEOs younger than 55?
- Which graph shows the relative frequency and which shows the cumulative relative frequency?




Figure 1.H.11

Use the following information to answer the next two exercises: Table 1.H.18 contains data on hurricanes that have made direct hits on the U.S. Between 1851 and 2004. A hurricane is given a strength category rating based on the minimum wind speed generated by the storm.

Category	Number of direct hits	Relative frequency	Cumulative frequency
1	109	0.3993	0.3993
2	72	0.2637	0.6630
3	71	0.2601	
4	18		0.9890
5	3	0.0110	1.0000
	Total = 273		

Table 1.18 Frequency of Hurricane Direct Hits

44.

What is the relative frequency of direct hits that were category 4 hurricanes?

- 0.0768
- 0.0659
- 0.2601
- Not enough information to calculate

45.

What is the relative frequency of direct hits that were AT MOST a category 3 storm?

- 0.3480
- 0.9231
- 0.2601
- 0.3370

This page titled [1.H: Sampling and Data \(Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **1.H: Sampling and Data (Homework)** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.R: Sampling and Data (Review)

1.1 Definitions of Statistics, Probability, and Key Terms

The mathematical theory of statistics is easier to learn when you know the language. This module presents important terms that will be used throughout the text.

1.2 Data, Sampling, and Variation in Data and Sampling

Data are individual items of information that come from a population or sample. Data may be classified as qualitative (categorical), quantitative continuous, or quantitative discrete.

Because it is not practical to measure the entire population in a study, researchers use samples to represent the population. A random sample is a representative group from the population chosen by using a method that gives each individual in the population an equal chance of being included in the sample. Random sampling methods include simple random sampling, stratified sampling, cluster sampling, and systematic sampling. Convenience sampling is a nonrandom method of choosing a sample that often produces biased data.

Samples that contain different individuals result in different data. This is true even when the samples are well-chosen and representative of the population. When properly selected, larger samples model the population more closely than smaller samples. There are many different potential problems that can affect the reliability of a sample. Statistical data needs to be critically analyzed, not simply accepted.

1.3 Levels of Measurement

Some calculations generate numbers that are artificially precise. It is not necessary to report a value to eight decimal places when the measures that generated that value were only accurate to the nearest tenth. Round off your final answer to one more decimal place than was present in the original data. This means that if you have data measured to the nearest tenth of a unit, report the final statistic to the nearest hundredth.

In addition to rounding your answers, you can measure your data using the following four levels of measurement.

- **Nominal scale level:** data that cannot be ordered nor can it be used in calculations
- **Ordinal scale level:** data that can be ordered; the differences cannot be measured
- **Interval scale level:** data with a definite ordering but no starting point; the differences can be measured, but there is no such thing as a ratio.
- **Ratio scale level:** data with a starting point that can be ordered; the differences have meaning and ratios can be calculated.

When organizing data, it is important to know how many times a value appears. How many statistics students study five hours or more for an exam? What percent of families on our block own two pets? Frequency, relative frequency, and cumulative relative frequency are measures that answer questions like these.

1.4 Experimental Design and Ethics

A poorly designed study will not produce reliable data. There are certain key components that must be included in every experiment. To eliminate lurking variables, subjects must be assigned randomly to different treatment groups. One of the groups must act as a control group, demonstrating what happens when the active treatment is not applied. Participants in the control group receive a placebo treatment that looks exactly like the active treatments but cannot influence the response variable. To preserve the integrity of the placebo, both researchers and subjects may be blinded. When a study is designed properly, the only difference between treatment groups is the one imposed by the researcher. Therefore, when groups respond differently to different treatments, the difference must be due to the influence of the explanatory variable.

“An ethics problem arises when you are considering an action that benefits you or some cause you support, hurts or reduces benefits to others, and violates some rule.” (Andrew Gelman, “Open Data and Open Methods,” Ethics and Statistics, <http://www.stat.columbia.edu/~gelman...nceEthics1.pdf> (accessed May 1, 2013).) Ethical violations in statistics are not always easy to spot. Professional associations and federal agencies post guidelines for proper conduct. It is important that you learn basic statistical procedures so that you can recognize proper data analysis.

This page titled [1.R: Sampling and Data \(Review\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [1.R: Sampling and Data \(Review\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

1.S: Sampling and Data (Solutions)

2.

- a. all children who take ski or snowboard lessons
- b. a group of these children
- c. the population mean age of children who take their first snowboard lesson
- d. the sample mean age of children who take their first snowboard lesson
- e. X = the age of one child who takes his or her first ski or snowboard lesson
- f. values for X , such as 3, 7, and so on

4.

- a. the clients of the insurance companies
- b. a group of the clients
- c. the mean health costs of the clients
- d. the mean health costs of the sample
- e. X = the health costs of one client
- f. values for X , such as 34, 9, 82, and so on

6.

- a. all the clients of this counselor
- b. a group of clients of this marriage counselor
- c. the proportion of all her clients who stay married
- d. the proportion of the sample of the counselor's clients who stay married
- e. X = the number of couples who stay married
- f. yes, no

8.

- a. all people (maybe in a certain geographic area, such as the United States)
- b. a group of the people
- c. the proportion of all people who will buy the product
- d. the proportion of the sample who will buy the product
- e. X = the number of people who will buy it
- f. buy, not buy

10.

a

12.

quantitative discrete, 150

14.

qualitative, Oakland A's

16.

quantitative discrete, 11,234 students

18.

qualitative, Crest

20.

quantitative continuous, 47.3 years

22.

b

24.

- a. The survey was conducted using six similar flights.
The survey would not be a true representation of the entire population of air travelers.
Conducting the survey on a holiday weekend will not produce representative results.
- b. Conduct the survey during different times of the year.
Conduct the survey using flights to and from various locations.
Conduct the survey on different days of the week.

26.

Answers will vary. Sample Answer: You could use a systematic sampling method. Stop the tenth person as they leave one of the buildings on campus at 9:50 in the morning. Then stop the tenth person as they leave a different building on campus at 1:50 in the afternoon.

28.

Answers will vary. Sample Answer: Many people will not respond to mail surveys. If they do respond to the surveys, you can't be sure who is responding. In addition, mailing lists can be incomplete.

30.

b

32.

convenience cluster stratified systematic simple random

34.

- a. qualitative(categorical)
- b. quantitative discrete
- c. quantitative discrete
- d. qualitative(categorical)

36.

Causality: The fact that two variables are related does not guarantee that one variable is influencing the other. We cannot assume that crime rate impacts education level or that education level impacts crime rate.

Confounding: There are many factors that define a community other than education level and crime rate. Communities with high crime rates and high education levels may have other lurking variables that distinguish them from communities with lower crime rates and lower education levels. Because we cannot isolate these variables of interest, we cannot draw valid conclusions about the connection between education and crime. Possible lurking variables include police expenditures, unemployment levels, region, average age, and size.

38.

- a. Possible reasons: increased use of caller id, decreased use of landlines, increased use of private numbers, voice mail, privacy managers, hectic nature of personal schedules, decreased willingness to be interviewed
- b. When a large number of people refuse to participate, then the sample may not have the same characteristics of the population. Perhaps the majority of people willing to participate are doing so because they feel strongly about the subject of the survey.

40.

Table 1.19

# flossing per week	Frequency	Relative frequency	Cumulative relative frequency
0	27	0.4500	0.4500
1	18	0.3000	0.7500
3	11	0.1833	0.9333
6	3	0.0500	0.9833

# flossing per week	Frequency	Relative frequency	Cumulative relative frequency
7	1	0.0167	1

- b. 5.00%
- c. 93.33%

42.

The sum of the travel times is 1,173.1. Divide the sum by 50 to calculate the mean value: 23.462. Because each state's travel time was measured to the nearest tenth, round this calculation to the nearest hundredth: 23.46.

44.

b

This page titled [1.S: Sampling and Data \(Solutions\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [1.S: Sampling and Data \(Solutions\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

2: Descriptive Statistics

- [2.1: introduction to Descriptive Statistics](#)
- [2.2: Display Data](#)
- [2.3: Measures of the Location of the Data](#)
- [2.4: Measures of the Center of the Data](#)
- [2.5: Sigma Notation and Calculating the Arithmetic Mean](#)
- [2.6: Geometric Mean](#)
- [2.7: Skewness and the Mean, Median, and Mode](#)
- [2.8: Measures of the Spread of the Data](#)
- [2.9: Chapter Key Terms](#)
- [2.10: Chapter Formula Review](#)
- [2.11: Chapter Homework](#)
- [2.12: Chapter Practice](#)
- [2.13: Homework](#)
- [2.14: Chapter References](#)
- [2.15: Chapter Homework Solutions](#)
- [2.R: Descriptive Statistics \(Review\)](#)

This page titled [2: Descriptive Statistics](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

2.1: introduction to Descriptive Statistics



Figure 2.1.1 When you have large amounts of data, you will need to organize it in a way that makes sense. These ballots from an election are rolled together with similar ballots to keep them organized. (credit: William Greeson)

Once you have collected data, what will you do with it? Data can be described and presented in many different formats. For example, suppose you are interested in buying a house in a particular area. You may have no clue about the house prices, so you might ask your real estate agent to give you a sample data set of prices. Looking at all the prices in the sample often is overwhelming. A better way might be to look at the median price and the variation of prices. The median and variation are just two ways that you will learn to describe data. Your agent might also provide you with a graph of the data.

In this chapter, you will study numerical and graphical ways to describe and display your data. This area of statistics is called "**Descriptive Statistics.**" You will learn how to calculate, and even more importantly, how to interpret these measurements and graphs.

A statistical graph is a tool that helps you learn about the shape or distribution of a sample or a population. A graph can be a more effective way of presenting data than a mass of numbers because we can see where data clusters and where there are only a few data values. Newspapers and the Internet use graphs to show trends and to enable readers to compare facts and figures quickly. Statisticians often graph data first to get a picture of the data. Then, more formal tools may be applied.

Some of the types of graphs that are used to summarize and organize data are the dot plot, the bar graph, the histogram, the stem-and-leaf plot, the frequency polygon (a type of broken line graph), the pie chart, and the box plot. In this chapter, we will briefly look at stem-and-leaf plots, line graphs, and bar graphs, as well as frequency polygons, and time series graphs. Our emphasis will be on histograms and box plots.

This page titled [2.1: introduction to Descriptive Statistics](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.0: introduction to Descriptive Statistics](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.2: Display Data

Stem-and-Leaf Graphs (Stemplots), Line Graphs, and Bar Graphs

One simple graph, the **stem-and-leaf graph** or **stemplot**, comes from the field of exploratory data analysis. It is a good choice when the data sets are small. To create the plot, divide each observation of data into a stem and a leaf. The leaf consists of a **final significant digit**. For example, 23 has stem two and leaf three. The number 432 has stem 43 and leaf two. Likewise, the number 5,432 has stem 543 and leaf two. The decimal 9.3 has stem nine and leaf three. Write the stems in a vertical line from smallest to largest. Draw a vertical line to the right of the stems. Then write the leaves in increasing order next to their corresponding stem.

Example 2.2.2.1

For Susan Dean's spring pre-calculus class, scores for the first exam were as follows (smallest to largest):

33; 42; 49; 49; 53; 55; 55; 61; 63; 67; 68; 68; 69; 69; 72; 73; 74; 78; 80; 83; 88; 88; 88; 88; 90; 92; 94; 94; 94; 94; 96; 100

Table 2.2.2.1 Stem-and-Leaf Graph

Stem	Leaf
3	3
4	2 9 9
5	3 5 5
6	1 3 7 8 8 9 9
7	2 3 4 8
8	0 3 8 8 8
9	0 2 4 4 4 4 6
10	0

The stemplot shows that most scores fell in the 60s, 70s, 80s, and 90s. Eight out of the 31 scores or approximately 26% (8/31) were in the 90s or 100, a fairly high number of As.

Exercise 2.2.2.1

For the Park City basketball team, scores for the last 30 games were as follows (smallest to largest):

32; 32; 33; 34; 38; 40; 42; 42; 43; 44; 46; 47; 47; 48; 48; 48; 49; 50; 50; 51; 52; 52; 52; 53; 54; 56; 57; 57; 60; 61

Construct a stem plot for the data.

The stemplot is a quick way to graph data and gives an exact picture of the data. You want to look for an overall pattern and any outliers. An **outlier** is an observation of data that does not fit the rest of the data. It is sometimes called an **extreme value**. When you graph an outlier, it will appear not to fit the pattern of the graph. Some outliers are due to mistakes (for example, writing down 50 instead of 500) while others may indicate that something unusual is happening. It takes some background information to explain outliers, so we will cover them in more detail later.

Example 2.2.2.2

The data are the distances (in kilometers) from a home to local supermarkets. Create a stemplot using the data:

1.1; 1.5; 2.3; 2.5; 2.7; 3.2; 3.3; 3.3; 3.5; 3.8; 4.0; 4.2; 4.5; 4.5; 4.7; 4.8; 5.5; 5.6; 6.5; 6.7; 12.3

Do the data seem to have any concentration of values?

NOTE

The leaves are to the right of the decimal.

Answer

The value 12.3 may be an outlier. Values appear to concentrate at three and four kilometers.

Table 2.2.2.2

Stem	Leaf
1	1 5
2	3 5 7
3	2 3 3 5 8
4	0 2 5 5 7 8
5	5 6
6	5 7
7	
8	
9	
10	
11	
12	3

Exercise 2.2.2.2

The following data show the distances (in miles) from the homes of off-campus statistics students to the college. Create a stem plot using the data and identify any outliers:

0.5; 0.7; 1.1; 1.2; 1.2; 1.3; 1.3; 1.5; 1.5; 1.7; 1.7; 1.8; 1.9; 2.0; 2.2; 2.5; 2.6; 2.8; 2.8; 2.8; 3.5; 3.8; 4.4; 4.8; 4.9; 5.2; 5.5; 5.7; 5.8; 8.0

Example 2.2.2.3

A **side-by-side stem-and-leaf plot** allows a comparison of the two data sets in two columns. In a side-by-side stem-and-leaf plot, two sets of leaves share the same stem. The leaves are to the left and the right of the stems. Table 2.2.2.4 and Table 2.2.2.5 show the ages of presidents at their inauguration and at their death. Construct a side-by-side stem-and-leaf plot using this data.

Answer

Table 2.2.2.3

Ages at Inauguration		Ages at Death
9 9 8 7 7 7 6 3 2	4	6 9
8 7 7 7 7 6 6 6 5 5 5 5 4 4 4 4 4 2 2 1 1 1 1 1 1 0	5	3 6 6 7 7 8
9 8 5 4 4 2 1 1 1 0	6	0 0 3 3 4 4 5 6 7 7 7 8
	7	0 0 1 1 1 4 7 8 8 9
	8	0 1 3 5 8
	9	0 0 3 3

Table 2.2.2.4 Presidential Ages at Inauguration

President	Age	President	Age	President	Age
Washington	57	Lincoln	52	Hoover	54
J. Adams	61	A. Johnson	56	F. Roosevelt	51
Jefferson	57	Grant	46	Truman	60
Madison	57	Hayes	54	Eisenhower	62
Monroe	58	Garfield	49	Kennedy	43
J. Q. Adams	57	Arthur	51	L. Johnson	55
Jackson	61	Cleveland	47	Nixon	56
Van Buren	54	B. Harrison	55	Ford	61
W. H. Harrison	68	Cleveland	55	Carter	52
Tyler	51	McKinley	54	Reagan	69
Polk	49	T. Roosevelt	42	G.H.W. Bush	64
Taylor	64	Taft	51	Clinton	47
Fillmore	50	Wilson	56	G. W. Bush	54
Pierce	48	Harding	55	Obama	47
Buchanan	65	Coolidge	51	Trump	70

Table 2.2.2.5 Presidential Age at Death

President	Age	President	Age	President	Age
Washington	67	Lincoln	56	Hoover	90
J. Adams	90	A. Johnson	66	F. Roosevelt	63
Jefferson	83	Grant	63	Truman	88
Madison	85	Hayes	70	Eisenhower	78
Monroe	73	Garfield	49	Kennedy	46
J. Q. Adams	80	Arthur	56	L. Johnson	64
Jackson	78	Cleveland	71	Nixon	81
Van Buren	79	B. Harrison	67	Ford	93
W. H. Harrison	68	Cleveland	71	Reagan	93
Tyler	71	McKinley	58		
Polk	53	T. Roosevelt	60		
Taylor	65	Taft	72		
Fillmore	74	Wilson	67		
Pierce	64	Harding	57		
Buchanan	77	Coolidge	60		

Another type of graph that is useful for specific data values is a **line graph**. In the particular line graph shown in [Example 2.2.4](#), the **x-axis**(horizontal axis) consists of **data values** and the **y-axis** (vertical axis) consists of **frequency points**. The frequency points are connected using line segments.

Example 2.2.2.4

In a survey, 40 mothers were asked how many times per week a teenager must be reminded to do his or her chores. The results are shown in Table 2.2.2.6 and in Figure 2.2.2.2.

Number of times teenager is reminded	Frequency
0	2
1	5
2	8
3	14
4	7
5	4

Table 2.6




Figure 2.2

Exercise 2.2.4

In a survey, 40 people were asked how many times per year they had their car in the shop for repairs. The results are shown in Table 2.2.7. Construct a line graph.

Number of times in shop	Frequency
0	7
1	10
2	14
3	9

Table 2.2.7

Bar graphs consist of bars that are separated from each other. The bars can be rectangles or they can be rectangular boxes (used in three-dimensional plots), and they can be vertical or horizontal. The **bar graph** shown in Example 2.2.5 has age groups represented on the **x-axis** and proportions on the **y-axis**.

Exercise 2.2.1

Add exercises text here.

Answer

Solution 2.5




Figure 2.2.2.3

Example 2.2.5

By the end of 2011, Facebook had over 146 million users in the United States. Table 2.2.2.8 shows three age groups, the number of users in each age group, and the proportion (%) of users in each age group. Construct a bar graph using this data.

Age groups	Number of Facebook users	Proportion (%) of Facebook users
13-25	65,082,280	45%
26-44	53,300,200	36%
45-64	27,885,100	19%

Table 2.2.8

Solution

Exercise 2.2.5

Add exercises text here.

Answer

The population in Park City is made up of children, working-age adults, and retirees. Table 2.2.9 shows the three age groups, the number of people in the town from each age group, and the proportion (%) of people in each age group. Construct a bar graph showing the proportions.

Age groups	Number of people	Proportion of population
Children	67,059	19%
Working-age adults	152,198	43%
Retirees	131,662	38%

Table 2.2.9

Example 2.2.2.6

The columns in Table 2.2.2.10 contain: the race or ethnicity of students in U.S. Public Schools for the class of 2011, percentages for the Advanced Placement examinee population for that class, and percentages for the overall student population. Create a bar graph with the student race or ethnicity (qualitative data) on the x -axis, and the Advanced Placement examinee population percentages on the y -axis.

Race/ethnicity	AP examinee population	Overall student population

Race/ethnicity	AP examinee population	Overall student population
1 = Asian, Asian American or Pacific Islander	10.3%	5.7%
2 = Black or African American	9.0%	14.7%
3 = Hispanic or Latino	17.0%	17.6%
4 = American Indian or Alaska Native	0.6%	1.1%
5 = White	57.1%	59.2%
6 = Not reported/other	6.0%	1.7%

Table 2.2.10

Answer

Solution 2.6




Figure 2.2.2.4

Exercise 2.2.2.6

Add exercises text here.

Answer

Park city is broken down into six voting districts. The table shows the percent of the total registered voter population that lives in each district as well as the percent total of the entire population that lives in each district. Construct a bar graph that shows the registered voter population by district.

Table 2.2.2.11

District	Registered voter population	Overall city population
1	15.5%	19.4%
2	12.2%	15.6%
3	9.8%	9.0%
4	17.4%	18.5%
5	22.8%	20.7%
6	22.3%	16.8%

Example 2.2.2.7

Below is a two-way table showing the types of pets owned by men and women:

Table 2.2.2.12

--

	Dogs	Cats	Fish	Total
Men	4	2	2	8
Women	4	6	2	12
Total	8	8	4	20

Given these data, calculate the conditional distributions for the subpopulation of men who own each pet type.

Answer

- Men who own dogs = $4/8 = 0.5$
- Men who own cats = $2/8 = 0.25$
- Men who own fish = $2/8 = 0.25$

Note: The sum of all of the conditional distributions must equal one. In this case, $0.5 + 0.25 + 0.25 = 1$; therefore, the solution "checks".

Histograms, Frequency Polygons, and Time Series Graphs

For most of the work you do in this book, you will use a histogram to display the data. One advantage of a histogram is that it can readily display large data sets. A rule of thumb is to use a histogram when the data set consists of 100 values or more.

A **histogram** consists of contiguous (adjoining) boxes. It has both a horizontal axis and a vertical axis. The horizontal axis is labeled with what the data represents (for instance, distance from your home to school). The vertical axis is labeled either **frequency** or **relative frequency** (or percent frequency or probability). The graph will have the same shape with either label. The histogram (like the stemplot) can give you the shape of the data, the center, and the spread of the data.

The relative frequency is equal to the frequency for an observed value of the data divided by the total number of data values in the sample. (Remember, frequency is defined as the number of times an answer occurs.) If:

- f = frequency
- n = total number of data values (or the sum of the individual frequencies), and
- RF = relative frequency,

then:

$$RF = \frac{f}{n}$$

For example, if three students in Mr. Ahab's English class of 40 students received from 90% to 100%, then, $f = 3$, $n = 40$, and $RF = \frac{f}{n} = \frac{3}{40} = 0.075$. 7.5% of the students received 90–100%. 90–100% are quantitative measures.

To construct a histogram, first decide how many **bars** or **intervals**, also called classes, represent the data. Many histograms consist of five to 15 bars or classes for clarity. The number of bars needs to be chosen. Choose a starting point for the first interval to be less than the smallest data value. A **convenient starting point** is a lower value carried out to one more decimal place than the value with the most decimal places. For example, if the value with the most decimal places is 6.1 and this is the smallest value, a convenient starting point is 6.05 ($6.1 - 0.05 = 6.05$). We say that 6.05 has more precision. If the value with the most decimal places is 2.23 and the lowest value is 1.5, a convenient starting point is 1.495 ($1.5 - 0.005 = 1.495$). If the value with the most decimal places is 3.234 and the lowest value is 1.0, a convenient starting point is 0.9995 ($1.0 - 0.0005 = 0.9995$). If all the data happen to be integers and the smallest value is two, then a convenient starting point is 1.5 ($2 - 0.5 = 1.5$). Also, when the starting point and other boundaries are carried to one additional decimal place, no data value will fall on a boundary. The next two examples go into detail about how to construct a histogram using continuous data and how to create a histogram using discrete data.

Example 2.2.2.8

The following data are the heights (in inches to the nearest half inch) of 100 male semiprofessional soccer players. The heights are **continuous** data, since height is measured.

60; 60.5; 61; 61; 61.5
63.5; 63.5; 63.5

64; 64; 64; 64; 64; 64; 64; 64; 64.5; 64.5; 64.5; 64.5; 64.5; 64.5; 64.5; 64.5
 66; 66; 66; 66; 66; 66; 66; 66; 66; 66; 66.5; 66.5; 66.5; 66.5; 66.5; 66.5; 66.5; 66.5; 66.5; 66.5; 66.5; 67; 67; 67; 67; 67; 67; 67; 67; 67; 67; 67; 67.5; 67.5; 67.5; 67.5; 67.5; 67.5; 67.5
 68; 68; 69; 69; 69; 69; 69; 69; 69; 69; 69; 69.5; 69.5; 69.5; 69.5
 70; 70; 70; 70; 70; 70; 70.5; 70.5; 70.5; 71; 71; 71
 72; 72; 72; 72.5; 72.5; 73; 73.5
 74

The smallest data value is 60. Since the data with the most decimal places has one decimal (for instance, 61.5), we want our starting point to have two decimal places. Since the numbers 0.5, 0.05, 0.005, etc. are convenient numbers, use 0.05 and subtract it from 60, the smallest value, for the convenient starting point.

$60 - 0.05 = 59.95$ which is more precise than, say, 61.5 by one decimal place. The starting point is, then, 59.95.

The largest value is 74, so $74 + 0.05 = 74.05$ is the ending value.

Next, calculate the width of each bar or class interval. To calculate this width, subtract the starting point from the ending value and divide by the number of bars (you must choose the number of bars you desire). Suppose you choose eight bars.

$$\frac{74.05 - 59.95}{8} = 1.76 \text{ \nonumber}$$

NOTE

We will round up to two and make each bar or class interval two units wide. Rounding up to two is one way to prevent a value from falling on a boundary. Rounding to the next number is often necessary even if it goes against the standard rules of rounding. For this example, using 1.76 as the width would also work. A guideline that is followed by some for the width of a bar or class interval is to take the square root of the number of data values and then round to the nearest whole number, if necessary. For example, if there are 150 values of data, take the square root of 150 and round to 12 bars or intervals.

The boundaries are:

- 59.95
- $59.95 + 2 = 61.95$
- $61.95 + 2 = 63.95$
- $63.95 + 2 = 65.95$
- $65.95 + 2 = 67.95$
- $67.95 + 2 = 69.95$
- $69.95 + 2 = 71.95$
- $71.95 + 2 = 73.95$
- $73.95 + 2 = 75.95$

The heights 60 through 61.5 inches are in the interval 59.95–61.95. The heights that are 63.5 are in the interval 61.95–63.95. The heights that are 64 through 64.5 are in the interval 63.95–65.95. The heights 66 through 67.5 are in the interval 65.95–67.95. The heights 68 through 69.5 are in the interval 67.95–69.95. The heights 70 through 71 are in the interval 69.95–71.95. The heights 72 through 73.5 are in the interval 71.95–73.95. The height 74 is in the interval 73.95–75.95.

The following histogram displays the heights on the *x*-axis and relative frequency on the *y*-axis.




Figure 2.2.2.5

Exercise 2.2.2.8

The following data are the shoe sizes of 50 male students. The sizes are continuous data since shoe size is measured. Construct a histogram and calculate the width of each bar or class interval. Suppose you choose six bars.

9; 9; 9.5; 9.5; 10; 10; 10; 10; 10; 10; 10.5; 10.5; 10.5; 10.5; 10.5; 10.5; 10.5; 10.5
 11; 11; 11; 11; 11; 11; 11; 11; 11; 11; 11; 11; 11; 11.5; 11.5; 11.5; 11.5; 11.5; 11.5; 11.5
 12; 12; 12; 12; 12; 12; 12; 12.5; 12.5; 12.5; 12.5; 14

Example 2.2.2.9

Create a histogram for the following data: the number of books bought by 50 part-time college students at ABC College. The number of books is **discrete data**, since books are counted.

1; 1; 1; 1; 1; 1; 1; 1; 1; 1; 1
 2; 2; 2; 2; 2; 2; 2; 2; 2; 2
 3; 3; 3; 3; 3; 3; 3; 3; 3; 3; 3; 3; 3; 3; 3
 4; 4; 4; 4; 4
 5; 5; 5; 5
 6; 6

Eleven students buy one book. Ten students buy two books. Sixteen students buy three books. Six students buy four books. Five students buy five books. Two students buy six books.

Because the data are integers, subtract 0.5 from 1, the smallest data value and add 0.5 to 6, the largest data value. Then the starting point is 0.5 and the ending value is 6.5.

Next, calculate the width of each bar or class interval. If the data are discrete and there are not too many different values, a width that places the data values in the middle of the bar or class interval is the most convenient. Since the data consist of the numbers 1, 2, 3, 4, 5, 6, and the starting point is 0.5, a width of one places the 1 in the middle of the interval from 0.5 to 1.5, the 2 in the middle of the interval from 1.5 to 2.5, the 3 in the middle of the interval from 2.5 to 3.5, the 4 in the middle of the interval from _____ to _____, the 5 in the middle of the interval from _____ to _____, and the _____ in the middle of the interval from _____ to _____.

Solution

Calculate the number of bars as follows:

$$\frac{6.5 - 0.5}{\text{number of bars}} = 1$$

where 1 is the width of a bar. Therefore, bars = 6.

The following histogram displays the number of books on the x-axis and the frequency on the y-axis.




Figure 2.2.2.6

Example 2.2.2.10

Using this data set, construct a histogram.

Table 2.2.2.13

Number of hours my classmates spent playing video games on weekends				
9.95	10	2.25	16.75	0
19.5	22.5	7.5	15	12.75
5.5	11	10	20.75	17.5
23	21.9	24	23.75	18
20	15	22.9	18.8	20.5

Answer

Solution 2.10



Figure 2.2.2.7

Some values in this data set fall on boundaries for the class intervals. A value is counted in a class interval if it falls on the left boundary, but not if it falls on the right boundary. Different researchers may set up histograms for the same data in different ways. There is more than one correct way to set up a histogram.

Frequency Polygons

Frequency polygons are analogous to line graphs, and just as line graphs make continuous data visually easy to interpret, so too do frequency polygons.

To construct a frequency polygon, first examine the data and decide on the number of intervals, or class intervals, to use on the x-axis and y-axis. After choosing the appropriate ranges, begin plotting the data points. After all the points are plotted, draw line segments to connect them.

Example 2.2.2.11

A frequency polygon was constructed from the frequency table below.

Table 2.2.2.14: Frequency distribution for calculus final test scores

Lower bound	Upper bound	Frequency	Cumulative frequency
49.5	59.5	5	5
59.5	69.5	10	15
69.5	79.5	30	45
79.5	89.5	40	85
89.5	99.5	15	100




Figure 2.2.2.8

The first label on the x -axis is 44.5. This represents an interval extending from 39.5 to 49.5. Since the lowest test score is 54.5, this interval is used only to allow the graph to touch the x -axis. The point labeled 54.5 represents the next interval, or the first “real” interval from the table, and contains five scores. This reasoning is followed for each of the remaining intervals with the point 104.5 representing the interval from 99.5 to 109.5. Again, this interval contains no data and is only used so that the graph will touch the x -axis. Looking at the graph, we say that this distribution is skewed because one side of the graph does not mirror the other side.

Exercise 2.2.2.11

Construct a frequency polygon of U.S. Presidents’ ages at inauguration shown in Table 2.2.15

Age at inauguration	Frequency
41.5–46.5	4
46.5–51.5	11
51.5–56.5	14
56.5–61.5	9
61.5–66.5	4
66.5–71.5	2

Table 2.2.15

Frequency polygons are useful for comparing distributions. This is achieved by overlaying the frequency polygons drawn for different data sets.

Example 2.2.2.12

We will construct an overlay frequency polygon comparing the scores from [Example 2.2.11](#) with the students' final numeric grade.

Table 2.2.2.16: Frequency distribution for calculus final test scores

Lower bound	Upper bound	Frequency	Cumulative frequency
49.5	59.5	5	5
59.5	69.5	10	15
69.5	79.5	30	45
79.5	89.5	40	85
89.5	99.5	15	100

Table 2.2.2.17: Frequency distribution for calculus final grades

Lower bound	Upper bound	Frequency	Cumulative frequency
49.5	59.5	10	10
59.5	69.5	10	20
69.5	79.5	30	50
79.5	89.5	45	95
89.5	99.5	5	100




Figure 2.2.2.9

Constructing a Time Series Graph

Suppose that we want to study the temperature range of a region for an entire month. Every day at noon we note the temperature and write this down in a log. A variety of statistical studies could be done with these data. We could find the mean or the median temperature for the month. We could construct a histogram displaying the number of days that temperatures reach a certain range of values. However, all of these methods ignore a portion of the data that we have collected.

One feature of the data that we may want to consider is that of time. Since each date is paired with the temperature reading for the day, we don't have to think of the data as being random. We can instead use the times given to impose a chronological order on the data. A graph that recognizes this ordering and displays the changing temperature as the month progresses is called a time series graph.

To construct a time series graph, we must look at both pieces of our **paired data set**. We start with a standard Cartesian coordinate system. The horizontal axis is used to plot the date or time increments, and the vertical axis is used to plot the values of the variable that we are measuring. By doing this, we make each point on the graph correspond to a date and a measured quantity. The points on the graph are typically connected by straight lines in the order in which they occur.

Example 2.2.2.13

The following data shows the Annual Consumer Price Index, each month, for ten years. Construct a time series graph for the Annual Consumer Price Index data only.

Table 2.2.2.18

Year	Jan	Feb	Mar	Apr	May	Jun	Jul
2003	181.7	183.1	184.2	183.8	183.5	183.7	183.9
2004	185.2	186.2	187.4	188.0	189.1	189.7	189.4
2005	190.7	191.8	193.3	194.6	194.4	194.5	195.4
2006	198.3	198.7	199.8	201.5	202.5	202.9	203.5
2007	202.416	203.499	205.352	206.686	207.949	208.352	208.299
2008	211.080	211.693	213.528	214.823	216.632	218.815	219.964
2009	211.143	212.193	212.709	213.240	213.856	215.693	215.351
2010	216.687	216.741	217.631	218.009	218.178	217.965	218.011
2011	220.223	221.309	223.467	224.906	225.964	225.722	225.922
2012	226.665	227.663	229.392	230.085	229.815	229.478	229.104

Table 2.2.2.19

Year	Aug	Sep	Oct	Nov	Dec	Annual
2003	184.6	185.2	185.0	184.5	184.3	184.0
2004	189.5	189.9	190.9	191.0	190.3	188.9
2005	196.4	198.8	199.2	197.6	196.8	195.3
2006	203.9	202.9	201.8	201.5	201.8	201.6
2007	207.917	208.490	208.936	210.177	210.036	207.342
2008	219.086	218.783	216.573	212.425	210.228	215.303
2009	215.834	215.969	216.177	216.330	215.949	214.537
2010	218.312	218.439	218.711	218.803	219.179	218.056
2011	226.545	226.889	226.421	226.230	225.672	224.939
2012	230.379	231.407	231.317	230.221	229.601	229.594

Answer

Solution 2.13




Figure 2.2.2.10

Exercise 2.2.2.13

The following table is a portion of a data set from www.worldbank.org. Use the table to construct a time series graph for CO₂ emissions for the United States.

Table 2.2.20: CO₂ emissions

Year	Ukraine	United Kingdom	United States
2003	352,259	540,640	5,681,664
2004	343,121	540,409	5,790,761
2005	339,029	541,990	5,826,394
2006	327,797	542,045	5,737,615
2007	328,357	528,631	5,828,697
2008	323,657	522,247	5,656,839
2009	272,176	474,579	5,299,563

Uses of a Time Series Graph

Time series graphs are important tools in various applications of statistics. When recording values of the same variable over an extended period of time, sometimes it is difficult to discern any trend or pattern. However, once the same data points are displayed graphically, some features jump out. Time series graphs make trends easy to spot.

How NOT to Lie with Statistics

It is important to remember that the very reason we develop a variety of methods to present data is to develop insights into the subject of what the observations represent. We want to get a "sense" of the data. Are the observations all very much alike or are they spread across a wide range of values, are they bunched at one end of the spectrum or are they distributed evenly and so on. We are trying to get a visual picture of the numerical data. Shortly we will develop formal mathematical measures of the data, but our visual graphical presentation can say much. It can, unfortunately, also say much that is distracting, confusing and simply wrong in terms of the impression the visual leaves. Many years ago Darrell Huff wrote the book *How to Lie with Statistics*. It has been through 25 plus printings and sold more than one and one-half million copies. His perspective was a harsh one and used many actual examples that were designed to mislead. He wanted to make people aware of such deception, but perhaps more importantly to educate so that others do not make the same errors inadvertently.

Again, the goal is to enlighten with visuals that tell the story of the data. Pie charts have a number of common problems when used to convey the message of the data. Too many pieces of the pie overwhelm the reader. More than perhaps five or six categories ought to give an idea of the relative importance of each piece. This is after all the goal of a pie chart, what subset matters most relative to the others. If there are more components than this then perhaps an alternative approach would be better or perhaps some can be consolidated into an "other" category. Pie charts cannot show changes over time, although we see this attempted all too often. In federal, state, and city finance documents pie charts are often presented to show the components of revenue available to the governing body for appropriation: income tax, sales tax motor vehicle taxes and so on. In and of itself this is interesting information and can be nicely done with a pie chart. The error occurs when two years are set side-by-side. Because the total revenues change year to year, but the size of the pie is fixed, no real information is provided and the relative size of each piece of the pie cannot be meaningfully compared.

Histograms can be very helpful in understanding the data. Properly presented, they can be a quick visual way to present probabilities of different categories by the simple visual of comparing relative areas in each category. Here the error, purposeful or not, is to vary the width of the categories. This of course makes comparison to the other categories impossible. It does embellish the importance of the category with the expanded width because it has a greater area, inappropriately, and thus visually "says" that that category has a higher probability of occurrence.

Time series graphs perhaps are the most abused. A plot of some variable across time should never be presented on axes that change part way across the page either in the vertical or horizontal dimension. Perhaps the time frame is changed from years to months. Perhaps this is to save space or because monthly data was not available for early years. In either case this confounds the presentation and destroys any value of the graph. If this is not done to purposefully confuse the reader, then it certainly is either lazy or sloppy work.

Changing the units of measurement of the axis can smooth out a drop or accentuate one. If you want to show large changes, then measure the variable in small units, penny rather than thousands of dollars. And of course to continue the fraud, be sure that the axis does not begin at zero, zero. If it begins at zero, zero, then it becomes apparent that the axis has been manipulated.

Perhaps you have a client that is concerned with the volatility of the portfolio you manage. An easy way to present the data is to use long time periods on the time series graph. Use months or better, quarters rather than daily or weekly data. If that doesn't get the volatility down then spread the time axis relative to the rate of return or portfolio valuation axis. If you want to show "quick" dramatic growth, then shrink the time axis. Any positive growth will show visually "high" growth rates. Do note that if the growth is negative then this trick will show the portfolio is collapsing at a dramatic rate.

Again, the goal of descriptive statistics is to convey meaningful visuals that tell the story of the data. Purposeful manipulation is fraud and unethical at the worst, but even at its best, making these type of errors will lead to confusion on the part of the analysis.

This page titled [2.2: Display Data](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.1: Display Data](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.3: Measures of the Location of the Data

The common measures of location are **quartiles** and **percentiles**

Quartiles are special percentiles. The first quartile, Q_1 , is the same as the 25th percentile, and the third quartile, Q_3 , is the same as the 75th percentile. The median, M , is called both the second quartile and the 50th percentile.

To calculate quartiles and percentiles, the data must be ordered from smallest to largest. Quartiles divide ordered data into quarters. Percentiles divide ordered data into hundredths. To score in the 90th percentile of an exam does not mean, necessarily, that you received 90% on a test. It means that 90% of test scores are the same or less than your score and 10% of the test scores are the same or greater than your test score.

Percentiles are useful for comparing values. For this reason, universities and colleges use percentiles extensively. One instance in which colleges and universities use percentiles is when SAT results are used to determine a minimum testing score that will be used as an acceptance factor. For example, suppose Duke accepts SAT scores at or above the 75th percentile. That translates into a score of at least 1220.

Percentiles are mostly used with very large populations. Therefore, if you were to say that 90% of the test scores are less (and not the same or less) than your score, it would be acceptable because removing one particular data value is not significant.

The **median** is a number that measures the "center" of the data. You can think of the median as the "middle value," but it does not actually have to be one of the observed values. It is a number that separates ordered data into halves. Half the values are the same number or smaller than the median, and half the values are the same number or larger. For example, consider the following data.

1; 11.5; 6; 7.2; 4; 8; 9; 10; 6.8; 8.3; 2; 2; 10; 1

Ordered from smallest to largest:

1; 1; 2; 2; 4; 6; 6.8; 7.2; 8; 8.3; 9; 10; 10; 11.5

Since there are 14 observations, the median is between the seventh value, 6.8, and the eighth value, 7.2. To find the median, add the two values together and divide by two.

$$\frac{6.8 + 7.2}{2} = 7$$

The median is seven. Half of the values are smaller than seven and half of the values are larger than seven.

Quartiles are numbers that separate the data into quarters. Quartiles may or may not be part of the data. To find the quartiles, first find the median or second quartile. The first quartile, Q_1 , is the middle value of the lower half of the data, and the third quartile, Q_3 , is the middle value, or median, of the upper half of the data. To get the idea, consider the same data set:

1; 1; 2; 2; 4; 6; 6.8; 7.2; 8; 8.3; 9; 10; 10; 11.5

The median or **second quartile** is seven. The lower half of the data are 1, 1, 2, 2, 4, 6, 6.8. The middle value of the lower half is two.

1; 1; 2; 2; 4; 6; 6.8

The number two, which is part of the data, is the **first quartile**. One-fourth of the entire sets of values are the same as or less than two and three-fourths of the values are more than two.

The upper half of the data is 7.2, 8, 8.3, 9, 10, 10, 11.5. The middle value of the upper half is nine.

The **third quartile**, Q_3 , is nine. Three-fourths (75%) of the ordered data set are less than nine. One-fourth (25%) of the ordered data set are greater than nine. The third quartile is part of the data set in this example.

The **interquartile range** is a number that indicates the spread of the middle half or the middle 50% of the data. It is the difference between the third quartile (Q_3) and the first quartile (Q_1).

$$IQR = Q_3 - Q_1$$

The IQR can help to determine potential **outliers**. A value is suspected to be a potential outlier if it is less than $\frac{1}{2}(1.5)(IQR)$ below the first quartile or more than $\frac{1}{2}(1.5)(IQR)$ above the third quartile. Potential outliers always require further investigation.

potential outlier

A potential outlier is a data point that is significantly different from the other data points. These special data points may be errors or some kind of abnormality or they may be a key to understanding the data.

Example 2.3.14

For the following 13 real estate prices, calculate the IQR and determine if any prices are potential outliers. Prices are in dollars.
389,950; 230,500; 158,000; 479,000; 639,000; 114,950; 5,500,000; 387,000; 659,000; 529,000; 575,000; 488,800;
1,095,000

Answer

Solution 2.14

Order the data from smallest to largest.

114, 950; 158, 000; 230, 500; 387, 000; 389, 950; 479, 000; 488, 800; 529, 000; 575, 000; 639, 000; 659, 000; 1, 095, 000; 5, 500, 000

$$M = 488, 800$$

$$Q_1 = \frac{230,500+387,000}{2} = 308, 750$$

$$Q_3 = \frac{639,000+659,000}{2} = 649, 000$$

$$IQR = 649, 000 - 308, 750 = 340, 250$$

$$(1.5)(IQR) = (1.5)(340, 250) = 510, 375$$

$$Q_1 - (1.5)(IQR) = 308, 750 - 510, 375 = -201, 625$$

$$Q_3 + (1.5)(IQR) = 649, 000 + 510, 375 = 1, 159, 375$$

No house price is less than $-201, 625$. However, $5, 500, 000$ is more than $1, 159, 375$. Therefore, $5, 500, 000$ is a potential **outlier**.

Example 2.3.15

For the two data sets in the [test scores example](#), find the following:

- The interquartile range. Compare the two interquartile ranges.
- Any outliers in either set.

Answer

Solution 2.15

The five number summary for the day and night classes is

Table 2.3.21

	Minimum	Q_1	Median	Q_3	Maximum
Day	32	56	74.5	82.5	99
Night	25.5	78	81	89	98

- a. The IQR for the day group is $Q_3 - Q_1 = 82.5 - 56 = 26.5$

The IQR for the night group is $Q_3 - Q_1 = 89 - 78 = 11$

The interquartile range (the spread or variability) for the day class is larger than the night class IQR . This suggests more variation will be found in the day class's class test scores.

- b. Day class outliers are found using the IQR times 1.5 rule. So,

- $Q_1 - IQR(1.5) = 56 - 26.5(1.5) = 16.25$
- $Q_3 + IQR(1.5) = 82.5 + 26.5(1.5) = 122.25$

Since the minimum and maximum values for the day class are greater than 16.25 and less than 122.25 there are no outliers.

Night class outliers are calculated as:

- $Q_1 - IQR(1.5) = 78 - 11(1.5) = 61.5$
- $Q_3 + IQR(1.5) = 89 + 11(1.5) = 105.5$

For this class, any test score less than 61.5 is an outlier. Therefore, the scores of 45 and 25.5 are outliers. Since no test score is greater than 105.5, there is no upper end outlier.

Example 2.3.16

Fifty statistics students were asked how much sleep they get per school night (rounded to the nearest hour). The results were:

Table 2.3.22

Amount of sleep per school night (hours)	Frequency	Relative frequency	Cumulative relative frequency
4	2	0.04	0.04

Amount of sleep per school night (hours)	Frequency	Relative frequency	Cumulative relative frequency
5	5	0.10	0.14
6	7	0.14	0.28
7	12	0.24	0.52
8	14	0.28	0.80
9	7	0.14	0.94
10	3	0.06	1.00

Find the 28th percentile. Notice the 0.28 in the "cumulative relative frequency" column. Twenty-eight percent of 50 data values is 14 values. There are 14 values less than the 28th percentile. They include the two 4s, the five 5s, and the seven 6s. The 28th percentile is between the last six and the first seven. **The 28th percentile is 6.5.**

Find the median. Look again at the "cumulative relative frequency" column and find 0.52. The median is the 50th percentile or the second quartile. 50% of 50 is 25. There are 25 values less than the median. They include the two 4s, the five 5s, the seven 6s, and eleven of the 7s. The median or 50th percentile is between the 25th, or seven, and 26th, or seven, values. **The median is seven.**

Find the third quartile. The third quartile is the same as the 75th percentile. You can "eyeball" this answer. If you look at the "cumulative relative frequency" column, you find 0.52 and 0.80. When you have all the fours, fives, sixes and sevens, you have 52% of the data. When you include all the 8s, you have 80% of the data. **The 75th percentile, then, must be an eight.** Another way to look at the problem is to find 75% of 50, which is 37.5, and round up to 38. The third quartile, Q_3 , is the 38th value, which is an eight. You can check this answer by counting the values. (There are 37 values below the third quartile and 12 values above.)

Exercise 2.3.16

Forty bus drivers were asked how many hours they spend each day running their routes (rounded to the nearest hour). Find the 65th percentile.

Table 2.3.23

Amount of time spent on route (hours)	Frequency	Relative frequency	Cumulative relative frequency
2	12	0.30	0.30
3	14	0.35	0.65
4	10	0.25	0.90
5	4	0.10	1.00

Example 2.3.17

Using Table 2.3.22

- Find the 80th percentile.
- Find the 90th percentile.
- Find the first quartile. What is another name for the first quartile?

Answer

Solution 2.17

Using the data from the frequency table, we have:

- The 80th percentile is between the last eight and the first nine in the table (between the 40th and 41st values). Therefore, we need to take the mean of the 40th and 41st values. The 80th percentile = $\frac{8+9}{2} = 8.5$
- The 90th percentile will be the 45th data value (location is $0.90(50) = 45$) and the 45th data value is nine.
- Q_1 is also the 25th percentile. The 25th percentile location calculation: $P_{25} = 0.25(50) = 12.5 \approx 13$ the 13th data value. Thus, the 25th percentile is six.

A Formula for Finding the k th Percentile

If you were to do a little research, you would find several formulas for calculating the k^{th} percentile. Here is one of them.

k = the k^{th} percentile. It may or may not be part of the data.

i = the index (ranking or position of a data value)

n = the total number of data points, or observations

- Order the data from smallest to largest.
- Calculate $i = \frac{k}{100}(n + 1)$
- If i is an integer, then the k^{th} percentile is the data value in the i^{th} position in the ordered set of data.
- If i is not an integer, then round i up and round i down to the nearest integers. Average the two data values in these two positions in the ordered data set. This is easier to understand in an example.

Example 2.3.18

Listed are 29 ages for Academy Award winning best actors *in order from smallest to largest*.

18; 21; 22; 25; 26; 27; 29; 30; 31; 33; 36; 37; 41; 42; 47; 52; 55; 57; 58; 62; 64; 67; 69; 71; 72; 73; 74; 76; 77

1. Find the 70th percentile.
2. Find the 83rd percentile.

Answer

Solution 2.18

1.

- $k = 70$
- $i =$ the index
- $n = 29$

$i = \frac{k}{100}(n + 1) = \left(\frac{70}{100}\right)(29 + 1) = 21$. Twenty-one is an integer, and the data value in the 21st position in the ordered data set is 64.

The 70th percentile is 64 years.

2.

- $k = 83^{\text{rd}}$ percentile
- $i =$ the index
- $n = 29$

$i = \frac{k}{100}(n + 1) = \left(\frac{83}{100}\right)(29 + 1) = 24.9$, which is NOT an integer. Round it down to 24 and up to 25. The age in the 24th position is 71 and the age in the 25th position is 72. Average 71 and 72. The 83rd percentile is 71.5 years.

Exercise 2.3.18

Listed are 29 ages for Academy Award winning best actors *in order from smallest to largest*.

18; 21; 22; 25; 26; 27; 29; 30; 31; 33; 36; 37; 41; 42; 47; 52; 55; 57; 58; 62; 64; 67; 69; 71; 72; 73; 74; 76; 77

Calculate the 20th percentile and the 55th percentile.

A Formula for Finding the Percentile of a Value in a Data Set

- Order the data from smallest to largest.
- x = the number of data values counting from the bottom of the data list up to but not including the data value for which you want to find the percentile.
- y = the number of data values equal to the data value for which you want to find the percentile.
- n = the total number of data.
- Calculate $\frac{x+0.5y}{n}(100)$. Then round to the nearest integer.

Example 2.3.19

Listed are 29 ages for Academy Award winning best actors *in order from smallest to largest*.

18; 21; 22; 25; 26; 27; 29; 30; 31; 33; 36; 37; 41; 42; 47; 52; 55; 57; 58; 62; 64; 67; 69; 71; 72; 73; 74; 76; 77

1. Find the percentile for 58.

2. Find the percentile for 25.

Answer

Solution 2.19

1. Counting from the bottom of the list, there are 18 data values less than 58. There is one value of 58.

$$x = 18 \text{ and } y = 1. \frac{x+0.5y}{n}(100) = \frac{18+0.5(1)}{29}(100) = 63.80 \text{ 58 is the } 64^{\text{th}} \text{ percentile.}$$

2. Counting from the bottom of the list, there are three data values less than 25. There is one value of 25.

$$x = 3 \text{ and } y = 1. \frac{x+0.5y}{n}(100) = \frac{3+0.5(1)}{29}(100) = 12.07. \text{ Twenty-five is the } 12^{\text{th}} \text{ percentile.}$$

Interpreting Percentiles, Quartiles, and Median

A percentile indicates the relative standing of a data value when data are sorted into numerical order from smallest to largest. Percentages of data values are less than or equal to the p th percentile. For example, 15% of data values are less than or equal to the 15th percentile.

- Low percentiles always correspond to lower data values.
- High percentiles always correspond to higher data values.

A percentile may or may not correspond to a value judgment about whether it is "good" or "bad." The interpretation of whether a certain percentile is "good" or "bad" depends on the context of the situation to which the data applies. In some situations, a low percentile would be considered "good;" in other contexts a high percentile might be considered "good". In many situations, there is no value judgment that applies.

Understanding how to interpret percentiles properly is important not only when describing data, but also when calculating probabilities in later chapters of this text.

NOTE

When writing the interpretation of a percentile in the context of the given data, the sentence should contain the following information.

- information about the context of the situation being considered
- the data value (value of the variable) that represents the percentile
- the percent of individuals or items with data values below the percentile
- the percent of individuals or items with data values above the percentile.

Example 2.3.20

On a timed math test, the first quartile for time it took to finish the exam was 35 minutes. Interpret the first quartile in the context of this situation.

Answer

Solution 2.20

Twenty-five percent of students finished the exam in 35 minutes or less. Seventy-five percent of students finished the exam in 35 minutes or more. A low percentile could be considered good, as finishing more quickly on a timed exam is desirable. (If you take too long, you might not be able to finish.)

Example 2.3.21

On a 20 question math test, the 70th percentile for number of correct answers was 16. Interpret the 70th percentile in the context of this situation.

Answer

Solution 2.21

Seventy percent of students answered 16 or fewer questions correctly. Thirty percent of students answered 16 or more questions correctly. A higher percentile could be considered good, as answering more questions correctly is desirable.

Exercise 2.3.21

On a 60 point written assignment, the 80th percentile for the number of points earned was 49. Interpret the 80th percentile in the context of this situation.

Example 2.3.22

At a community college, it was found that the 30th percentile of credit units that students are enrolled for is seven units. Interpret the 30th percentile in the context of this situation.

Answer

Solution 2.22

- Thirty percent of students are enrolled in seven or fewer credit units.
- Seventy percent of students are enrolled in seven or more credit units.
- In this example, there is no "good" or "bad" value judgment associated with a higher or lower percentile. Students attend community college for varied reasons and needs, and their course load varies according to their needs.

Example 2.3.23

Sharpe Middle School is applying for a grant that will be used to add fitness equipment to the gym. The principal surveyed 15 anonymous students to determine how many minutes a day the students spend exercising. The results from the 15 anonymous students are shown.

0 minutes; 40 minutes; 60 minutes; 30 minutes; 60 minutes

10 minutes; 45 minutes; 30 minutes; 300 minutes; 90 minutes;

30 minutes; 120 minutes; 60 minutes; 0 minutes; 20 minutes

Determine the following five values.

- Min = 0
- $Q_1 = 20$
- Med = 40
- $Q_3 = 60$
- Max = 300

If you were the principal, would you be justified in purchasing new fitness equipment? Since 75% of the students exercise for 60 minutes or less daily, and since the IQR is 40 minutes ($60 - 20 = 40$), we know that half of the students surveyed exercise between 20 minutes and 60 minutes daily. This seems a reasonable amount of time spent exercising, so the principal would be justified in purchasing the new equipment.

However, the principal needs to be careful. The value 300 appears to be a potential outlier.

$$Q_3 + 1.5(IQR) = 60 + (1.5)(40) = 120.$$

The value 300 is greater than 120 so it is a potential outlier. If we delete it and calculate the five values, we get the following values:

- Min = 0
- $Q_1 = 20$
- $Q_3 = 60$
- Max = 120

We still have 75% of the students exercising for 60 minutes or less daily and half of the students exercising between 20 and 60

This page titled [2.3: Measures of the Location of the Data](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.2: Measures of the Location of the Data](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.4: Measures of the Center of the Data

The "center" of a data set is also a way of describing location. The two most widely used measures of the "center" of the data are the **mean**(average) and the **median**. To calculate the **mean weight** of 50 people, add the 50 weights together and divide by 50. Technically this is the arithmetic mean. We will discuss the geometric mean later. To find the **median weight** of the 50 people, order the data and find the number that splits the data into two equal parts meaning an equal number of observations on each side. The weight of 25 people are below this weight and 25 people are heavier than this weight. The median is generally a better measure of the center when there are extreme values or outliers because it is not affected by the precise numerical values of the outliers. The mean is the most common measure of the center.

NOTE

The words "mean" and "average" are often used interchangeably. The substitution of one word for the other is common practice. The technical term is "arithmetic mean" and "average" is technically a center location. Formally, the arithmetic mean is called the first moment of the distribution by mathematicians. However, in practice among non-statisticians, "average" is commonly accepted for "arithmetic mean."

When each value in the data set is not unique, the mean can be calculated by multiplying each distinct value by its frequency and then dividing the sum by the total number of data values. The letter used to represent the **sample mean** is an x with a bar over it (pronounced " x bar"): \bar{x} .

The Greek letter μ (pronounced "mew") represents the **population mean**. One of the requirements for the **sample mean** to be a good estimate of the **population mean** is for the sample taken to be truly random.

To see that both ways of calculating the mean are the same, consider the sample:

1; 1; 1; 2; 2; 3; 4; 4; 4; 4; 4

$$\bar{x} = \frac{1 + 1 + 1 + 2 + 2 + 3 + 4 + 4 + 4 + 4 + 4}{11} = 2.7$$

$$\bar{x} = \frac{3(1) + 2(2) + 1(3) + 5(4)}{11} = 2.7$$

In the second calculation, the frequencies are 3, 2, 1, and 5.

You can quickly find the location of the median by using the expression $\frac{n+1}{2}$.

The letter n is the total number of data values in the sample. If n is an odd number, the median is the middle value of the ordered data (ordered smallest to largest). If n is an even number, the median is equal to the two middle values added together and divided by two after the data has been ordered. For example, if the total number of data values is 97, then $\frac{n+1}{2} = \frac{97+1}{2} = 49$. The median is the 49th value in the ordered data. If the total number of data values is 100, then $\frac{n+1}{2} = \frac{100+1}{2} = 50.5$. The median occurs midway between the 50th and 51st values. The location of the median and the value of the median are **not** the same. The upper case letter M is often used to represent the median. The next example illustrates the location of the median and the value of the median.

Example 2.24

AIDS data indicating the number of months a patient with AIDS lives after taking a new antibody drug are as follows (smallest to largest):
3; 4; 8; 8; 10; 11; 12; 13; 14; 15; 15; 16; 16; 17; 17; 18; 21; 22; 22; 24; 24; 25; 26; 26; 27; 27; 29; 29; 31; 32; 33; 33; 34; 34; 35; 37; 40; 44; 44; 47;
Calculate the mean and the median.

Answer

Solution 2.24

The calculation for the mean is:

$$\bar{x} = \frac{[3+4+(8)(2)+10+11+12+13+14+(15)(2)+\dots+35+37+40+(44)(2)+47]}{40} = 23.6$$

To find the median, M , first use the formula for the location. The location is:

$$\frac{n+1}{2} = \frac{40+1}{2} = 20.5$$

Starting at the smallest value, the median is located between the 20th and 21st values (the two 24s):

3; 4; 8; 8; 10; 11; 12; 13; 14; 15; 15; 16; 16; 17; 17; 18; 21; 22; 22; 24; 24; 25; 26; 26; 27; 27; 29; 29; 31; 32; 33; 33; 34; 34; 35; 37; 40; 44; 44; 47;

$$M = \frac{24+24}{2} = 24$$

Example 2.25

Suppose that in a small town of 50 people, one person earns \$5,000,000 per year and the other 49 each earn \$30,000. Which is the better measure of the "center": the mean or the median?

Answer

Solution 2.25

$$\bar{x} = \frac{5,000,000+49(30,000)}{50} = 129,400$$

$M = 30,000$

(There are 49 people who earn \$30,000 and one person who earns \$5,000,000.)

The median is a better measure of the "center" than the mean because 49 of the values are 30,000 and one is 5,000,000. The 5,000,000 is an outlier. The 30,000 gives us a better sense of the middle of the data.

Another measure of the center is the mode. The **mode** is the most frequent value. There can be more than one mode in a data set as long as those values have the same frequency and that frequency is the highest. A data set with two modes is called bimodal.

Example 2.26

Statistics exam scores for 20 students are as follows:

50; 53; 59; 59; 63; 63; 72; 72; 72; 72; 72; 72; 76; 78; 81; 83; 84; 84; 84; 90; 93

Find the mode.

Answer

Solution 2.26

The most frequent score is 72, which occurs five times. Mode = 72.

Example 2.27

Five real estate exam scores are 430, 430, 480, 480, 495. The data set is bimodal because the scores 430 and 480 each occur twice.

When is the mode the best measure of the "center"? Consider a weight loss program that advertises a mean weight loss of six pounds the first week of the program. The mode might indicate that most people lose two pounds the first week, making the program less appealing.

NOTE

The mode can be calculated for qualitative data as well as for quantitative data. For example, if the data set is: red, red, red, green, green, yellow, purple, black, blue, the mode is red.

Calculating the Arithmetic Mean of Grouped Frequency Tables

When only grouped data is available, you do not know the individual data values (we only know intervals and interval frequencies); therefore, you cannot compute an exact mean for the data set. What we must do is estimate the actual mean by calculating the mean of a frequency table. A frequency table is a data representation in which grouped data is displayed along with the corresponding frequencies. To calculate the mean from a grouped frequency table we can apply the basic definition of mean: $mean = \frac{\text{data sum}}{\text{number of data values}}$. We simply need to modify the definition to fit within the restrictions of a frequency table.

Since we do not know the individual data values we can instead find the midpoint of each interval. The midpoint is $\frac{\text{lower boundary} + \text{upper boundary}}{2}$. We can now modify the mean definition to be **Mean of Frequency Table** = $\frac{\sum fm}{\sum f}$ where f = the frequency of the interval and m = the midpoint of the interval.

Example 2.28

A frequency table displaying professor Blount's last statistic test is shown. Find the best estimate of the class mean.

Table 2.24

Grade interval	Number of students
50–56.5	1
56.5–62.5	0
62.5–68.5	4
68.5–74.5	4
74.5–80.5	2
80.5–86.5	3
86.5–92.5	4
92.5–98.5	1

Answer

Solution 2.28

Find the midpoints for all intervals

Table 2.25

Grade interval	Midpoint
50–56.5	53.25
56.5–62.5	59.5
62.5–68.5	65.5
68.5–74.5	71.5
74.5–80.5	77.5
80.5–86.5	83.5
86.5–92.5	89.5
92.5–98.5	95.5

- Calculate the sum of the product of each interval frequency and midpoint. $\sum fm$
 $53.25(1) + 59.5(0) + 65.5(4) + 71.5(4) + 77.5(2) + 83.5(3) + 89.5(4) + 95.5(1) = 1460.25$
- $\mu = \frac{\sum fm}{\sum f} = \frac{1460.25}{19} = 76.86$

Exercise 2.28

Maris conducted a study on the effect that playing video games has on memory recall. As part of her study, she compiled the following data:

Table 2.26

Hours teenagers spend on video games	Number of teenagers
0–3.5	3
3.5–7.5	7
7.5–11.5	12
11.5–15.5	7
15.5–19.5	9

What is the best estimate for the mean number of hours spent playing video games?

This page titled [2.4: Measures of the Center of the Data](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.3: Measures of the Center of the Data](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.5: Sigma Notation and Calculating the Arithmetic Mean

Formula for Population Mean

$$\mu = \frac{1}{N} \sum_{i=1}^N x_i$$

Formula for Sample Mean

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

This unit is here to remind you of material that you once studied and said at the time “I am sure that I will never need this!”

Here are the formulas for a population mean and the sample mean. The Greek letter μ is the symbol for the population mean and \bar{x} is the symbol for the sample mean. Both formulas have a mathematical symbol that tells us how to make the calculations. It is called Sigma notation because the symbol is the Greek capital letter sigma: Σ . Like all mathematical symbols it tells us what to do: just as the plus sign tells us to add and the x tells us to multiply. These are called mathematical operators. The Σ symbol tells us to add a specific list of numbers.

Let's say we have a sample of animals from the local animal shelter and we are interested in their average age. If we list each value, or observation, in a column, you can give each one an index number. The first number will be number 1 and the second number 2 and so on.

Table 2.5.27

Animal	Age
1	9
2	1
3	8.5
4	10.5
5	10
6	8.5
7	12
8	8
9	1
10	9.5

Each observation represents a particular animal in the sample. Purr is animal number one and is a 9 year old cat, Toto is animal number 2 and is a 1 year old puppy and so on.

To calculate the mean we are told by the formula to add up all these numbers, ages in this case, and then divide the sum by 10, the total number of animals in the sample.

Animal number one, the cat Purr, is designated as X_1 , animal number 2, Toto, is designated as X_2 and so on through Dundee who is animal number 10 and is designated as X_{10} .

The i in the formula tells us which of the observations to add together. In this case it is X_1 through X_{10} which is all of them. We know which ones to add by the indexing notation, the $i = 1$ and the n or capital N for the population. For this example the indexing notation would be $i = 1$ and because it is a sample we use a small n on the top of the Σ which would be 10.

The standard deviation requires the same mathematical operator and so it would be helpful to recall this knowledge from your past.

The sum of the ages is found to be 78 and dividing by 10 gives us the sample mean age as 7.8 years.

This page titled [2.5: Sigma Notation and Calculating the Arithmetic Mean](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.4: Sigma Notation and Calculating the Arithmetic Mean](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.6: Geometric Mean

The mean (Arithmetic), median and mode are all measures of the “center” of the data, the “average”. They are all in their own way trying to measure the “common” point within the data, that which is “normal”. In the case of the arithmetic mean this is solved by finding the value from which all points are equal linear distances. We can imagine that all the data values are combined through addition and then distributed back to each data point in equal amounts. The sum of all the values is what is redistributed in equal amounts such that the total sum remains the same.

The geometric mean redistributes not the sum of the values but the product of multiplying all the individual values and then redistributing them in equal portions such that the total product remains the same. This can be seen from the formula for the geometric mean, \tilde{x} : (*Pronounced x-tilde*)

$$\tilde{x} = \left(\prod_{i=1}^n x_i \right)^{\frac{1}{n}} = \sqrt[n]{x_1 \cdot x_2 \cdots x_n} = (x_1 \cdot x_2 \cdots x_n)^{\frac{1}{n}}$$

where π is another mathematical operator, that tells us to multiply all the x_i numbers in the same way capital Greek sigma tells us to add all the x_i numbers. Remember that a fractional exponent is calling for the n th root of the number thus an exponent of $1/3$ is the cube root of the number.

The geometric mean answers the question, “if all the quantities had the same value, what would that value have to be in order to achieve the same product?” The geometric mean gets its name from the fact that when redistributed in this way the sides form a geometric shape for which all sides have the same length. To see this, take the example of the numbers 10, 51.2 and 8. The geometric mean is the product of multiplying these three numbers together (4,096) and taking the cube root because there are three numbers among which this product is to be distributed. Thus the geometric mean of these three numbers is 16. This describes a cube $16 \times 16 \times 16$ and has a volume of 4,096 units.

The geometric mean is relevant in Economics and Finance for dealing with growth: growth of markets, in investment, population and other variables the growth in which there is an interest. Imagine that our box of 4,096 units (perhaps dollars) is the value of an investment after three years and that the investment returns in percents were the three numbers in our example. The geometric mean will provide us with the answer to the question, what is the average rate of return: 16 percent. The arithmetic mean of these three numbers is 23.6 percent. The reason for this difference, 16 versus 23.6, is that the arithmetic mean is additive and thus does not account for the interest on the interest, compound interest, embedded in the investment growth process. The same issue arises when asking for the average rate of growth of a population or sales or market penetration, etc., knowing the annual rates of growth. The formula for the geometric mean rate of return, or any other growth rate, is:

$$r_s = (x_1 \cdot x_2 \cdots x_n)^{\frac{1}{n}} - 1$$

Manipulating the formula for the geometric mean can also provide a calculation of the average rate of growth between two periods knowing only the initial value a_0 and the ending value a_n and the number of periods, n . The following formula provides this information:

$$\left(\frac{a_n}{a_0} \right)^{\frac{1}{n}} = \tilde{x}$$

Finally, we note that the formula for the geometric mean requires that all numbers be positive, greater than zero. The reason of course is that the root of a negative number is undefined for use outside of mathematical theory. There are ways to avoid this problem however. In the case of rates of return and other simple growth problems we can convert the negative values to meaningful positive equivalent values. Imagine that the annual returns for the past three years are +12%, -8%, and +2%. Using the decimal multiplier equivalents of 1.12, 0.92, and 1.02, allows us to compute a geometric mean of 1.0167. Subtracting 1 from this value gives the geometric mean of +1.67% as a net rate of population growth (or financial return). From this example we can see that the geometric mean provides us with this formula for calculating the geometric (mean) rate of return for a series of annual rates of return:

$$r_s = \tilde{x} - 1$$

where r_s is average rate of return and \tilde{x} is the geometric mean of the returns during some number of time periods. Note that the length of each time period must be the same.

As a general rule one should convert the percent values to its decimal equivalent multiplier. It is important to recognize that when dealing with percents, the geometric mean of percent values does not equal the geometric mean of the decimal multiplier equivalents and it is the decimal multiplier equivalent geometric mean that is relevant.

This page titled [2.6: Geometric Mean](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.5: Geometric Mean](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.7: Skewness and the Mean, Median, and Mode

Consider the following data set.

4; 5; 6; 6; 6; 7; 7; 7; 7; 7; 7; 8; 8; 8; 9; 10

This data set can be represented by following histogram. Each interval has width one, and each value is located in the middle of an interval.




Figure 2.11

The histogram displays a **symmetrical** distribution of data. A distribution is symmetrical if a vertical line can be drawn at some point in the histogram such that the shape to the left and the right of the vertical line are mirror images of each other. The mean, the median, and the mode are each seven for these data. **In a perfectly symmetrical distribution, the mean and the median are the same.** This example has one mode (unimodal), and the mode is the same as the mean and median. In a symmetrical distribution that has two modes (bimodal), the two modes would be different from the mean and median.

The histogram for the data: 4; 5; 6; 6; 6; 7; 7; 7; 8 is not symmetrical. The right-hand side seems "chopped off" compared to the left side. A distribution of this type is called **skewed to the left** because it is pulled out to the left. We can formally measure the skewness of a distribution just as we can mathematically measure the center weight of the data or its "spreadness". The mathematical formula for skewness is:

$$a_3 = \frac{\sum (x_t - \bar{x})^3}{n s^3}.$$

The greater the deviation from zero indicates a greater degree of skewness. If the skewness is negative then the distribution is skewed left as in Figure 2.7.13




Figure 2.12

The mean is 6.3, the median is 6.5, and the mode is seven. **Notice that the mean is less than the median, and they are both less than the mode.** The mean and the median both reflect the skewing, but the mean reflects it more so.

The histogram for the data: 6; 7; 7; 7; 7; 8; 8; 8; 9; 10, is also not symmetrical. It is **skewed to the right**.




Figure 2.13

The mean is 7.7, the median is 7.5, and the mode is seven. Of the three statistics, **the mean is the largest, while the mode is the smallest**. Again, the mean reflects the skewing the most.

To summarize, generally if the distribution of data is skewed to the left, the mean is less than the median, which is often less than the mode. If the distribution of data is skewed to the right, the mode is often less than the median, which is less than the mean.

As with the mean, median and mode, and as we will see shortly, the variance, there are mathematical formulas that give us precise measures of these characteristics of the distribution of the data. Again looking at the formula for skewness we see that this is a relationship between the mean of the data and the individual observations cubed.

$$a_3 = \sum \frac{(x_i - \bar{x})^3}{ns^3}$$

where ss is the sample standard deviation of the data, X_i , and \bar{x} is the arithmetic mean and n is the sample size.

Formally the arithmetic mean is known as the first moment of the distribution. The second moment we will see is the variance, and skewness is the third moment. The variance measures the squared differences of the data from the mean and skewness measures the cubed differences of the data from the mean. While a variance can never be a negative number, the measure of skewness can and this is how we determine if the data are skewed right of left. The skewness for a normal distribution is zero, and any symmetric data should have skewness near zero. Negative values for the skewness indicate data that are skewed left and positive values for the skewness indicate data that are skewed right. By skewed left, we mean that the left tail is long relative to the right tail. Similarly, skewed right means that the right tail is long relative to the left tail. The skewness characterizes the degree of asymmetry of a distribution around its mean. While the mean and standard deviation are *dimensional* quantities (this is why we will take the square root of the variance) that is, have the same units as the measured quantities X_i , the skewness is conventionally defined in such a way as to make it *nondimensional*. It is a pure number that characterizes only the shape of the distribution. A positive value of skewness signifies a distribution with an asymmetric tail extending out towards more positive X and a negative value signifies a distribution whose tail extends out towards more negative X . A zero measure of skewness will indicate a symmetrical distribution.

Skewness and symmetry become important when we discuss probability distributions in later chapters.

This page titled [2.7: Skewness and the Mean, Median, and Mode](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.6: Skewness and the Mean, Median, and Mode](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.8: Measures of the Spread of the Data

An important characteristic of any set of data is the variation in the data. In some data sets, the data values are concentrated closely near the mean; in other data sets, the data values are more widely spread out from the mean. The most common measure of variation, or spread, is the standard deviation. The **standard deviation** is a number that measures how far data values are from their mean.

The standard deviation

- provides a numerical measure of the overall amount of variation in a data set, and
- can be used to determine whether a particular data value is close to or far from the mean.

The standard deviation provides a measure of the overall variation in a data set

The standard deviation is always positive or zero. The standard deviation is small when the data are all concentrated close to the mean, exhibiting little variation or spread. The standard deviation is larger when the data values are more spread out from the mean, exhibiting more variation.

Suppose that we are studying the amount of time customers wait in line at the checkout at supermarket *A* and supermarket *B*. The average wait time at both supermarkets is five minutes. At supermarket *A*, the standard deviation for the wait time is two minutes; at supermarket *B*. The standard deviation for the wait time is four minutes.

Because supermarket *B* has a higher standard deviation, we know that there is more variation in the wait times at supermarket *B*. Overall, wait times at supermarket *B* are more spread out from the average; wait times at supermarket *A* are more concentrated near the average.

Calculating the Standard Deviation

If x is a number, then the difference " x minus the mean" is called its **deviation**. In a data set, there are as many deviations as there are items in the data set. The deviations are used to calculate the standard deviation. If the numbers belong to a population, in symbols a deviation is $x - \mu$. For sample data, in symbols a deviation is $x - \bar{x}$.

The procedure to calculate the standard deviation depends on whether the numbers are the entire population or are data from a sample. The calculations are similar, but not identical. Therefore the symbol used to represent the standard deviation depends on whether it is calculated from a population or a sample. The lower case letter s represents the sample standard deviation and the Greek letter σ (sigma, lower case) represents the population standard deviation. If the sample has the same characteristics as the population, then s should be a good estimate of σ .

To calculate the standard deviation, we need to calculate the variance first. The **variance** is the **average of the squares of the deviations** (the $x - \bar{x}$ values for a sample, or the $x - \mu$ values for a population). The symbol σ^2 represents the population variance; the population standard deviation σ is the square root of the population variance. The symbol s^2 represents the sample variance; the sample standard deviation s is the square root of the sample variance. You can think of the standard deviation as a special average of the deviations. Formally, the variance is the second moment of the distribution or the first moment around the mean. Remember that the mean is the first moment of the distribution.

If the numbers come from a census of the entire **population** and not a sample, when we calculate the average of the squared deviations to find the variance, we divide by N , the number of items in the population. If the data are from a **sample** rather than a population, when we calculate the average of the squared deviations, we divide by $n - 1$, one less than the number of items in the sample.

Formulas for the Sample Standard Deviation

- $s = \sqrt{\frac{\sum(x - \bar{x})^2}{n - 1}}$ or $s = \sqrt{\frac{\sum f(x - \bar{x})^2}{n - 1}}$ or $s = \sqrt{\frac{(\sum_{i=1}^n x_i^2) - nx^2}{n - 1}}$
- For the sample standard deviation, the denominator is $n - 1$, that is the sample size minus 1.

Formulas for the Population Standard Deviation

- $\sigma = \sqrt{\frac{\sum(x - \mu)^2}{N}}$ or $\sigma = \sqrt{\frac{\sum f(x - \mu)^2}{N}}$ or $\sigma = \sqrt{\frac{\sum_{i=1}^N x_i^2}{N} - \mu^2}$
- For the population standard deviation, the denominator is N , the number of items in the population.

In these formulas, f represents the frequency with which a value appears. For example, if a value appears once, f is one. If a value appears three times in the data set or population, f is three. Two important observations concerning the variance and standard deviation: the deviations are measured from the mean and the deviations are squared. In principle, the deviations could be measured from any point, however, our interest is measurement from the center weight of the data, what is the "normal" or most usual value of the observation. Later we will be trying to measure the "unusualness" of an observation or a sample mean and thus we need a measure from the mean. The second observation is that the deviations are squared. This does two things, first it makes the deviations all positive and second it changes the units of measurement from that of the mean and the original observations. If the data are weights then the mean is measured in pounds, but the variance is measured in pounds-squared. One reason to use the standard deviation is to return to the original units of measurement by taking the square root of the variance. Further, when the deviations are squared it explodes their value. For example, a deviation of 10 from the mean when squared is 100, but a deviation of 100 from the mean is 10,000. What this does is place great weight on outliers when calculating the variance.

Types of Variability in Samples

When trying to study a population, a sample is often used, either for convenience or because it is not possible to access the entire population. Variability is the term used to describe the differences that may occur in these outcomes. Common types of variability include the following:

- Observational or measurement variability
- Natural variability
- Induced variability
- Sample variability

Here are some examples to describe each type of variability.

Example 1: Measurement variability

Measurement variability occurs when there are differences in the instruments used to measure or in the people using those instruments. If we are gathering data on how long it takes for a ball to drop from a height by having students measure the time of the drop with a stopwatch, we may experience measurement variability if the two stopwatches used were made by different manufacturers: For example, one stopwatch measures to the nearest second, whereas the other one measures to the nearest tenth of a second. We also may experience measurement variability because two different people are gathering the data. Their reaction times in pressing the button on the stopwatch may differ; thus, the outcomes will vary accordingly. The differences in outcomes may be affected by measurement variability.

Example 2: Natural variability

Natural variability arises from the differences that naturally occur because members of a population differ from each other. For example, if we have two identical corn plants and we expose both plants to the same amount of water and sunlight, they may still grow at different rates simply because they are two different corn plants. The difference in outcomes may be explained by natural variability.

Example 3: Induced variability

Induced variability is the counterpart to natural variability; this occurs because we have artificially induced an element of variation (that, by definition, was not present naturally): For example, we assign people to two different groups to study memory, and we induce a variable in one group by limiting the amount of sleep they get. The difference in outcomes may be affected by induced variability.

Example 4: Sample variability

Sample variability occurs when multiple random samples are taken from the same population. For example, if I conduct four surveys of 50 people randomly selected from a given population, the differences in outcomes may be affected by sample variability.

Example 2.8.29

In a fifth grade class, the teacher was interested in the average age and the sample standard deviation of the ages of her students. The following data are the ages for a SAMPLE of $n = 20$ fifth grade students. The ages are rounded to the nearest half year:

9; 9.5; 9.5; 10; 10; 10; 10; 10.5; 10.5; 10.5; 10.5; 11; 11; 11; 11; 11; 11; 11.5; 11.5; 11.5;

$$\bar{x} = \frac{9 + 9.5(2) + 10(4) + 10.5(4) + 11(6) + 11.5(3)}{20} = 10.525$$

The average age is 10.53 years, rounded to two places.

The variance may be calculated by using a table. Then the standard deviation is calculated by taking the square root of the variance. We will explain the parts of the table after calculating s .

Table 2.8.28

Data	Freq.	Deviations	Deviations ²	(Freq.)(Deviations ²)
x	f	$(x - \bar{x})$	$(x - \bar{x})^2$	$(f)(x - \bar{x})^2$
9	1	$9 - 10.525 = -1.525$	$(-1.525)^2 = 2.325625$	$1 \times 2.325625 = 2.325625$
9.5	2	$9.5 - 10.525 = -1.025$	$(-1.025)^2 = 1.050625$	$2 \times 1.050625 = 2.101250$
10	4	$10 - 10.525 = -0.525$	$(-0.525)^2 = 0.275625$	$4 \times 0.275625 = 1.1025$
10.5	4	$10.5 - 10.525 = -0.025$	$(-0.025)^2 = 0.000625$	$4 \times 0.000625 = 0.0025$
11	6	$11 - 10.525 = 0.475$	$(0.475)^2 = 0.225625$	$6 \times 0.225625 = 1.35375$
11.5	3	$11.5 - 10.525 = 0.975$	$(0.975)^2 = 0.950625$	$3 \times 0.950625 = 2.851875$
				The total is 9.7375

The sample variance, s^2 , is equal to the sum of the last column (9.7375) divided by the total number of data values minus one ($20 - 1$):

$$s^2 = \frac{9.7375}{20-1} = 0.5125$$

The **sample standard deviation** s is equal to the square root of the sample variance:

$$s = \sqrt{0.5125} = 0.715891, \text{ which is rounded to two decimal places, } s = 0.72.$$

Explanation of the standard deviation calculation shown in the table

The deviations show how spread out the data are about the mean. The data value 11.5 is farther from the mean than is the data value 11 which is indicated by the deviations 0.97 and 0.47. A positive deviation occurs when the data value is greater than the mean, whereas a negative deviation occurs when the data value is less than the mean. The deviation is -1.525 for the data value nine. **If you add the deviations, the sum is always zero.** (For [Example 2.8.29](#), there are $n = 20$ deviations.) So you cannot simply add the deviations to get the spread of the data. By squaring the deviations, you make them positive numbers, and the sum will also be positive. The variance, then, is the average squared deviation. By squaring the deviations we are placing an extreme penalty on observations that are far from the mean; these observations get greater weight in the calculations of the variance. We will see later on that the variance (standard deviation) plays the critical role in determining our conclusions in inferential statistics. We can begin now by using the standard deviation as a measure of "unusualness." "How did you do on the test?" "Terrific! Two standard deviations above the mean." This, we will see, is an unusually good exam grade.

The variance is a squared measure and does not have the same units as the data. Taking the square root solves the problem. The standard deviation measures the spread in the same units as the data.

Notice that instead of dividing by $n = 20$, the calculation divided by $n - 1 = 20 - 1 = 19$ because the data is a sample. For the **sample** variance, we divide by the sample size minus one ($n - 1$). Why not divide by n ? The answer has to do with the population variance. **The sample variance is an estimate of the population variance.** This estimate requires us to use an estimate of the population mean rather than the actual population mean. Based on the theoretical mathematics that lies behind these calculations, dividing by $(n - 1)$ gives a better estimate of the population variance.

The standard deviation, s or σ , is either zero or larger than zero. Describing the data with reference to the spread is called "variability". The variability in data depends upon the method by which the outcomes are obtained; for example, by measuring or by random sampling. When the standard deviation is zero, there is no spread; that is, all the data values are equal to each other. The standard deviation is small when the data are all concentrated close to the mean, and is larger when the data values show more variation from the mean. When the standard deviation is a lot larger than zero, the data values are very spread out about the mean; outliers can make s or σ very large.

Example 2.8.30

Use the following data (first exam scores) from Susan Dean's spring pre-calculus class:

33; 42; 49; 49; 53; 55; 55; 61; 63; 67; 68; 68; 69; 69; 72; 73; 74; 78; 80; 83; 88; 88; 88; 90; 92; 94; 94; 94; 94; 96; 100

- a. Create a chart containing the data, frequencies, relative frequencies, and cumulative relative frequencies to three decimal places.
- b. Calculate the following to one decimal place:
 - i. The sample mean
 - ii. The sample standard deviation
 - iii. The median
 - iv. The first quartile
 - v. The third quartile
 - vi. IQR

Answer

Solution 2.30

a. See Table 2.8.29

b.

- i. The sample mean = 73.5
- ii. The sample standard deviation = 17.9
- iii. The median = 73
- iv. The first quartile = 61
- v. The third quartile = 90
- vi. $IQR = 90 - 61 = 29$

Table 2.8.29

Data	Frequency	Relative frequency	Cumulative relative frequency
33	1	0.032	0.032
42	1	0.032	0.064
49	2	0.065	0.129
53	1	0.032	0.161
55	2	0.065	0.226
61	1	0.032	0.258
63	1	0.032	0.29
67	1	0.032	0.322
68	2	0.065	0.387
69	2	0.065	0.452
72	1	0.032	0.484

Data	Frequency	Relative frequency	Cumulative relative frequency
73	1	0.032	0.516
74	1	0.032	0.548
78	1	0.032	0.580
80	1	0.032	0.612
83	1	0.032	0.644
88	3	0.097	0.741
90	1	0.032	0.773
92	1	0.032	0.805
94	4	0.129	0.934
96	1	0.032	0.966
100	1	0.032	0.998 (Why isn't this value 1? Answer: Rounding)

Standard deviation of Grouped Frequency Tables

Recall that for grouped data we do not know individual data values, so we cannot describe the typical value of the data with precision. In other words, we cannot find the exact mean, median, or mode. We can, however, determine the best estimate of the measures of center by finding the mean of the grouped data with the formula: Mean of Frequency Table = $\frac{\sum \backslash(f \backslash)m}{\sum f}$ where f = interval frequencies and m = interval midpoints.

Just as we could not find the exact mean, neither can we find the exact standard deviation. Remember that standard deviation describes numerically the expected deviation a data value has from the mean. In simple English, the standard deviation allows us to compare how “unusual” individual data is compared to the mean.

Example 2.8.31

Find the standard deviation for the data in [Table 2.8.30](#)

Table 2.8.30

Class	Frequency, f	Midpoint, m	$f \cdot m$	$f(m - \bar{x})^2$
0-2	1	1	$1 \cdot 1 = 1$	$1(1 - 6.88)^2 = 34.57$
3-5	6	4	$6 \cdot 4 = 24$	$6(4 - 6.88)^2 = 49.77$
6-8	10	7	$10 \cdot 7 = 70$	$10(7 - 6.88)^2 = 0.14$
9-11	7	10	$7 \cdot 10 = 70$	$7(10 - 6.88)^2 = 68.14$
12-14	0	13	$0 \cdot 13 = 0$	$0(13 - 6.88)^2 = 0$
	$n = 24$		$\bar{x} = 16524 = 6.88$	$s^2 = 152.6224 - 1 = 6.64$

For this data set, we have the mean, $\bar{x} = 6.88$ and the standard deviation, $s_x = 2.58$. This means that a randomly selected data value would be expected to be 2.58 units from the mean. If we look at the first class, we see that the class midpoint is equal to

one. This is almost three standard deviations from the mean. While the formula for calculating the standard deviation is not complicated,

$$s_x = \sqrt{\frac{\sum(m - \bar{x})^2 f}{n - 1}}$$

where s_x = sample standard deviation, \bar{x} = sample mean, the calculations are tedious. It is usually best to use technology when performing the calculations.

Comparing Values from Different Data Sets

The standard deviation is useful when comparing data values that come from different data sets. If the data sets have different means and standard deviations, then comparing the data values directly can be misleading.

- For each data value x , calculate how many standard deviations away from its mean the value is.
- Use the formula: $x = \text{mean} + (\text{\#of STDEVs})(\text{standard deviation})$; solve for #of STDEVs.
- # of *STDEVs* = $\frac{x - \text{mean}}{\text{standard deviation}}$
- Compare the results of this calculation.

#of STDEVs is often called a "z-score"; we can use the symbol z . In symbols, the formulas become:

Table 2.8.31

Sample	$x = \bar{x} + zs$	$z = \frac{x - \bar{x}}{s}$
Population	$x = \mu + z\sigma$	$z = \frac{x - \mu}{\sigma}$

Example 2.8.32

Two students, John and Ali, from different high schools, wanted to find out who had the highest GPA when compared to his school. Which student had the highest GPA when compared to his school?

Table 2.8.32

Student	GPA	School mean GPA	School standard deviation
John	2.85	3.0	0.7
Ali	77	80	10

Answer

Solution 2.32

For each student, determine how many standard deviations (#of STDEVs) his GPA is away from the average, for his school. Pay careful attention to signs when comparing and interpreting the answer.

$$z = \# \text{ of STDEVs} = \frac{\text{value} - \text{mean}}{\text{standard deviation}} = \frac{x - \mu}{\sigma}$$

$$\text{For John, } z = \# \text{ of STDEVs} = \frac{2.85 - 3.0}{0.7} = -0.21$$

$$\text{For Ali, } z = \# \text{ of STDEVs} = \frac{77 - 80}{10} = -0.3$$

John has the better GPA when compared to his school because his GPA is 0.21 standard deviations **below** his school's mean while Ali's GPA is 0.3 standard deviations **below** his school's mean.

John's z-score of -0.21 is higher than Ali's z-score of -0.3 . For GPA, higher values are better, so we conclude that John has the better GPA when compared to his school.

Exercise 2.8.32

Add exercises text here.

Answer

Two swimmers, Angie and Beth, from different teams, wanted to find out who had the fastest time for the 50 meter freestyle when compared to her team. Which swimmer had the fastest time when compared to her team?

Table 2.8.33

Swimmer	Time (seconds)	Team mean time	Team standard deviation
Angie	26.2	27.2	0.8
Beth	27.3	30.1	1.4

The following lists give a few facts that provide a little more insight into what the standard deviation tells us about the distribution of the data.

For ANY data set, no matter what the distribution of the data is:

- At least 75% of the data is within two standard deviations of the mean.
- At least 89% of the data is within three standard deviations of the mean.
- At least 95% of the data is within 4.5 standard deviations of the mean.
- This is known as Chebyshev's Rule.

For data having a Normal Distribution, which we will examine in great detail later:

- Approximately 68% of the data is within one standard deviation of the mean.
- Approximately 95% of the data is within two standard deviations of the mean.
- More than 99% of the data is within three standard deviations of the mean.
- This is known as the Empirical Rule.
- It is important to note that this rule only applies when the shape of the distribution of the data is bell-shaped and symmetric. We will learn more about this when studying the "Normal" or "Gaussian" probability distribution in later chapters.

Coefficient of Variation

Another useful way to compare distributions besides simple comparisons of means or standard deviations is to adjust for differences in the scale of the data being measured. Quite simply, a large variation in data with a large mean is different than the same variation in data with a small mean. To adjust for the scale of the underlying data the Coefficient of Variation (CV) has been developed. Mathematically:

$$CV = \frac{s}{\bar{x}} * 100 \text{ conditioned upon } \bar{x} \neq 0, \text{ where } s \text{ is the standard deviation of the data and } \bar{x}$$

We can see that this measures the variability of the underlying data as a percentage of the mean value; the center weight of the data set. This measure is useful in comparing risk where an adjustment is warranted because of differences in scale of two data sets. In effect, the scale is changed to common scale, percentage differences, and allows direct comparison of the two or more magnitudes of variation of different data sets.

This page titled [2.8: Measures of the Spread of the Data](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.7: Measures of the Spread of the Data](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.9: Chapter Key Terms

Frequency

the number of times a value of the data occurs

Frequency Table

a data representation in which grouped data is displayed along with the corresponding frequencies

Histogram

a graphical representation in x-y form of the distribution of data in a data set; x represents the data and y represents the frequency, or relative frequency. The graph consists of contiguous rectangles.

Interquartile Range

or *IQR*, is the range of the middle 50 percent of the data values; the *IQR* is found by subtracting the first quartile from the third quartile.

Mean (arithmetic)

a number that measures the central tendency of the data; a common name for mean is 'average.' The term 'mean' is a shortened form of 'arithmetic mean.' By definition, the mean for a sample (denoted by \bar{x}) is $\bar{x} = \frac{\text{Sum of all values in the sample}}{\text{Number of values in the sample}}$, and the mean for a population (denoted by μ) is $\mu = \frac{\text{Sum of all values in the population}}{\text{Number of values in the population}}$

Mean (geometric)

a measure of central tendency that provides a measure of average geometric growth over multiple time periods.

Median

a number that separates ordered data into halves; half the values are the same number or smaller than the median and half the values are the same number or larger than the median. The median may or may not be part of the data.

Midpoint

the mean of an interval in a frequency table

Mode

the value that appears most frequently in a set of data

Outlier

an observation that does not fit the rest of the data

Percentile

a number that divides ordered data into hundredths; percentiles may or may not be part of the data. The median of the data is the second quartile and the 50th percentile. The first and third quartiles are the 25th and the 75th percentiles, respectively.

Quartiles

the numbers that separate the data into quarters; quartiles may or may not be part of the data. The second quartile is the median of the data.

Relative Frequency

the ratio of the number of times a value of the data occurs in the set of all outcomes to the number of all outcomes

Standard Deviation

a number that is equal to the square root of the variance and measures how far data values are from their mean; notation: s for sample standard deviation and σ for population standard deviation.

Variance

mean of the squared deviations from the mean, or the square of the standard deviation; for a set of data, a deviation can be represented as $x - \bar{x}$ where x is a value of the data and \bar{x} is the sample mean. The sample variance is equal to the sum of the squares of the deviations divided by the difference of the sample size and one.

This page titled [2.9: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

2.10: Chapter Formula Review

2.2 Measures of the Location of the Data

$$i = \left(\frac{k}{100}\right)(n + 1)$$

where i = the ranking or position of a data value,

k = the k th percentile,

n = total number of data.

$$\text{Expression for finding the percentile of a data value: } \left(\frac{x+0.5y}{n}\right)(100)$$

where x = the number of values counting from the bottom of the data list up to but not including the data value for which you want to find the percentile,

y = the number of data values equal to the data value for which you want to find the percentile,

n = total number of data

2.3 Measures of the Center of the Data

$$\mu = \frac{\sum fm}{\sum f} \text{ Where } f = \text{interval frequencies and } m = \text{interval midpoints.}$$

The arithmetic mean for a sample (denoted by \bar{x}) is $\bar{x} = \frac{\text{Sum of all values in the sample}}{\text{Number of values in the sample}}$

The arithmetic mean for a population (denoted by μ) is $\mu = \frac{\text{Sum of all values in the population}}{\text{Number of values in the population}}$

2.5 Geometric Mean

The Geometric Mean: $\bar{x} = \left(\prod_{i=1}^n x_i\right)^{\frac{1}{n}} = \sqrt[n]{x_1 \cdot x_2 \cdots x_n} = (x_1 \cdot x_2 \cdots x_n)^{\frac{1}{n}}$

2.6 Skewness and the Mean, Median, and Mode

$$\text{Formula for skewness: } a_3 = \frac{\sum (x_i - \bar{x})^3}{ns^2}$$

Formula for Coefficient of Variation: $CV = \frac{s}{\bar{x}} \cdot 100$ conditioned upon $\bar{x} \neq 0$

2.7 Measures of the Spread of the Data

$$s_x = \sqrt{\frac{\sum fm^2}{n} - \bar{x}^2} \text{ where } s_x = \text{sample standard deviation}$$

\bar{x} = sample mean

Formulas for Sample Standard Deviation $s = \sqrt{\frac{\sum(x-\bar{x})^2}{n-1}}$ or $s = \sqrt{\frac{\sum f(x-\bar{x})^2}{n-1}}$ or $s = \sqrt{\frac{(\sum_{i=1}^n x_i^2) - nx^2}{n-1}}$ For the sample standard deviation, the denominator is $n - 1$, that is the sample size - 1.

Formulas for Population Standard Deviation $\sigma = \sqrt{\frac{\sum(x-\mu)^2}{N}}$ or $\sigma = \sqrt{\frac{\sum f(x-\mu)^2}{N}}$ or $\sigma = \sqrt{\frac{\sum_{i=1}^N x_i^2}{N} - \mu^2}$ For the population standard deviation, the denominator is N , the number of items in the population.

This page titled [2.10: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.9: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.11: Chapter Homework

2.1 Display Data

84.

Table 2.11.63 contains the 2010 obesity rates in U.S. states and Washington, DC.

Table 2.11.63

State	Percent (%)	State	Percent (%)	State	Percent (%)
Alabama	32.2	Kentucky	31.3	North Dakota	27.2
Alaska	24.5	Louisiana	31.0	Ohio	29.2
Arizona	24.3	Maine	26.8	Oklahoma	30.4
Arkansas	30.1	Maryland	27.1	Oregon	26.8
California	24.0	Massachusetts	23.0	Pennsylvania	28.6
Colorado	21.0	Michigan	30.9	Rhode Island	25.5
Connecticut	22.5	Minnesota	24.8	South Carolina	31.5
Delaware	28.0	Mississippi	34.0	South Dakota	27.3
Washington, DC	22.2	Missouri	30.5	Tennessee	30.8
Florida	26.6	Montana	23.0	Texas	31.0
Georgia	29.6	Nebraska	26.9	Utah	22.5
Hawaii	22.7	Nevada	22.4	Vermont	23.2
Idaho	26.5	New Hampshire	25.0	Virginia	26.0
Illinois	28.2	New Jersey	23.8	Washington	25.5
Indiana	29.6	New Mexico	25.1	West Virginia	32.5
Iowa	28.4	New York	23.9	Wisconsin	26.3
Kansas	29.4	North Carolina	27.8	Wyoming	25.1

- Use a random number generator to randomly pick eight states. Construct a bar graph of the obesity rates of those eight states.
- Construct a bar graph for all the states beginning with the letter "A."
- Construct a bar graph for all the states beginning with the letter "M."

85.

Suppose that three book publishers were interested in the number of fiction paperbacks adult consumers purchase per month. Each publisher conducted a survey. In the survey, adult consumers were asked the number of fiction paperbacks they had purchased the previous month. The results are as follows:

Table 2.11.64 Publisher A

# of books	Freq.	Rel. freq.
0	10	
1	12	
2	16	
3	12	
4	8	
5	6	
6	2	

# of books	Freq.	Rel. freq.
8	2	

Table 2.11.65 Publisher B

# of books	Freq.	Rel. freq.
0	18	
1	24	
2	24	
3	22	
4	15	
5	10	
7	5	
9	1	

Table 2.11.66 Publisher C

# of books	Freq.	Rel. freq.
0–1	20	
2–3	35	
4–5	12	
6–7	2	
8–9	1	

- Find the relative frequencies for each survey. Write them in the charts.
- Use the frequency column to construct a histogram for each publisher's survey. For Publishers A and B, make bar widths of one. For Publisher C, make bar widths of two.
- In complete sentences, give two reasons why the graphs for Publishers A and B are not identical.
- Would you have expected the graph for Publisher C to look like the other two graphs? Why or why not?
- Make new histograms for Publisher A and Publisher B. This time, make bar widths of two.
- Now, compare the graph for Publisher C to the new graphs for Publishers A and B. Are the graphs more similar or more different? Explain your answer.

86.

Often, cruise ships conduct all on-board transactions, with the exception of gambling, on a cashless basis. At the end of the cruise, guests pay one bill that covers all onboard transactions. Suppose that 60 single travelers and 70 couples were surveyed as to their on-board bills for a seven-day cruise from Los Angeles to the Mexican Riviera. Following is a summary of the bills for each group.

Table 2.11.67 Singles

Amount(\$)	Frequency	Rel. frequency
51–100	5	
101–150	10	
151–200	15	
201–250	15	
251–300	10	
301–350	5	

Table 2.11.68 Couples

--

Amount(\$)	Frequency	Rel. frequency
100–150	5	
201–250	5	
251–300	5	
301–350	5	
351–400	10	
401–450	10	
451–500	10	
501–550	10	
551–600	5	
601–650	5	

- Fill in the relative frequency for each group.
- Construct a histogram for the singles group. Scale the x -axis by \$50 widths. Use relative frequency on the y -axis.
- Construct a histogram for the couples group. Scale the x -axis by \$50 widths. Use relative frequency on the y -axis.
- Compare the two graphs:
 - List two similarities between the graphs.
 - List two differences between the graphs.
 - Overall, are the graphs more similar or different?
- Construct a new graph for the couples by hand. Since each couple is paying for two individuals, instead of scaling the x -axis by \$50, scale it by \$100. Use relative frequency on the y -axis.
- Compare the graph for the singles with the new graph for the couples:
 - List two similarities between the graphs.
 - Overall, are the graphs more similar or different?
- How did scaling the couples graph differently change the way you compared it to the singles graph?
- Based on the graphs, do you think that individuals spend the same amount, more or less, as singles as they do person by person as a couple? Explain why in one or two complete sentences.

87.


Twenty-five randomly selected students were asked the number of movies they watched the previous week. The results are as follows.

Table 2.11.69

# of movies	Frequency	Relative frequency	Cumulative relative frequency
0	5		
1	9		
2	6		
3	4		
4	1		

- Construct a histogram of the data.
- Complete the columns of the chart.

Use the following information to answer the next two exercises: Suppose one hundred eleven people who shopped in a special t-shirt store were asked the number of t-shirts they own costing more than \$19 each.



88.

The percentage of people who own at most three t-shirts costing more than \$19 each is approximately:

- a. 21
- b. 59
- c. 41
- d. Cannot be determined

89.

If the data were collected by asking the first 111 people who entered the store, then the type of sampling is:

- a. cluster
- b. simple random
- c. stratified
- d. convenience

90.

Following are the 2010 obesity rates by U.S. states and Washington, DC.

Table 2.11.70

State	Percent (%)	State	Percent (%)	State	Percent (%)
Alabama	32.2	Kentucky	31.3	North Dakota	27.2
Alaska	24.5	Louisiana	31.0	Ohio	29.2
Arizona	24.3	Maine	26.8	Oklahoma	30.4
Arkansas	30.1	Maryland	27.1	Oregon	26.8
California	24.0	Massachusetts	23.0	Pennsylvania	28.6
Colorado	21.0	Michigan	30.9	Rhode Island	25.5
Connecticut	22.5	Minnesota	24.8	South Carolina	31.5
Delaware	28.0	Mississippi	34.0	South Dakota	27.3
Washington, DC	22.2	Missouri	30.5	Tennessee	30.8
Florida	26.6	Montana	23.0	Texas	31.0
Georgia	29.6	Nebraska	26.9	Utah	22.5
Hawaii	22.7	Nevada	22.4	Vermont	23.2
Idaho	26.5	New Hampshire	25.0	Virginia	26.0
Illinois	28.2	New Jersey	23.8	Washington	25.5
Indiana	29.6	New Mexico	25.1	West Virginia	32.5
Iowa	28.4	New York	23.9	Wisconsin	26.3
Kansas	29.4	North Carolina	27.8	Wyoming	25.1

Construct a bar graph of obesity rates of your state and the four states closest to your state. Hint: Label the x-axis with the states.

2.2 Measures of the Location of the Data

91.

The median age for U.S. blacks currently is 30.9 years; for U.S. whites it is 42.3 years.

- Based upon this information, give two reasons why the black median age could be lower than the white median age.
- Does the lower median age for blacks necessarily mean that blacks die younger than whites? Why or why not?
- How might it be possible for blacks and whites to die at approximately the same age, but for the median age for whites to be higher?

92.

Six hundred adult Americans were asked by telephone poll, "What do you think constitutes a middle-class income?" The results are in [Table 2.71](#). Also, include left endpoint, but not the right endpoint.

Table 2.11.71

Salary (\$)	Relative frequency
< 20,000	0.02
20,000–25,000	0.09
25,000–30,000	0.19
30,000–40,000	0.26
40,000–50,000	0.18
50,000–75,000	0.17
75,000–99,999	0.02
100,000+	0.01

- What percentage of the survey answered "not sure"?
- What percentage think that middle-class is from \$25,000 to \$50,000?
- Construct a histogram of the data.
 - Should all bars have the same width, based on the data? Why or why not?
 - How should the <20,000 and the 100,000+ intervals be handled? Why?
- Find the 40th and 80th percentiles
- Construct a bar graph of the data

2.3 Measures of the Center of the Data

93.

The most obese countries in the world have obesity rates that range from 11.4% to 74.6%. This data is summarized in the following table.

Table 2.11.72

Percent of population obese	Number of countries
11.4–20.45	29
20.45–29.45	13
29.45–38.45	4
38.45–47.45	0
47.45–56.45	2
56.45–65.45	1
65.45–74.45	0
74.45–83.45	1

- a. What is the best estimate of the average obesity percentage for these countries?
- b. The United States has an average obesity rate of 33.9%. Is this rate above average or below?
- c. How does the United States compare to other countries?

94.

Table 2.11.73 gives the percent of children under five considered to be underweight. What is the best estimate for the mean percentage of underweight children?

Table 2:73

Percent of underweight children	Number of countries
16–21.45	23
21.45–26.9	4
26.9–32.35	9
32.35–37.8	7
37.8–43.25	6
43.25–48.7	1

2.4 Sigma Notation and Calculating the Arithmetic Mean

95.

A sample of 10 prices is chosen from a population of 100 similar items. The values obtained from the sample, and the values for the population, are given in Table 2.11.74 and Table 2.11.75 respectively.

- a. Is the mean of the sample within \$1 of the population mean?
- b. What is the difference in the sample and population means?

Table 2.11.74

Prices of the sample
\$21
\$23
\$21
\$24
\$22
\$22
\$25
\$21
\$20
\$24

Table 2.11.75

Prices of the population	Frequency
\$20	20
\$21	35
\$22	15
\$23	10
\$24	18
\$25	2

96.

A standardized test is given to ten people at the beginning of the school year with the results given in Table 2.11.76 below. At the end of the year the same people were again tested.

- What is the average improvement?
- Does it matter if the means are subtracted, or if the individual values are subtracted?

Table 2.11.76

Student	Beginning score	Ending score
1	1100	1120
2	980	1030
3	1200	1208
4	998	1000
5	893	948
6	1015	1030
7	1217	1224
8	1232	1245
9	967	988
10	988	997

97.

A small class of 7 students has a mean grade of 82 on a test. If six of the grades are 80, 82, 86, 90, 90, and 95, what is the other grade?

98.

A class of 20 students has a mean grade of 80 on a test. Nineteen of the students has a mean grade between 79 and 82, inclusive.

- What is the lowest possible grade of the other student?
- What is the highest possible grade of the other student?

99.

If the mean of 20 prices is \$10.39, and 5 of the items with a mean of \$10.99 are sampled, what is the mean of the other 15 prices?

2.5 Geometric Mean

100.

An investment grows from \$10,000 to \$22,000 in five years. What is the average rate of return?

101.

An initial investment of \$20,000 grows at a rate of 9% for five years. What is its final value?

102.

A culture contains 1,300 bacteria. The bacteria grow to 2,000 in 10 hours. What is the rate at which the bacteria grow per hour to the nearest tenth of a percent?

103.

An investment of \$3,000 grows at a rate of 5% for one year, then at a rate of 8% for three years. What is the average rate of return to the nearest hundredth of a percent?

104.

An investment of \$10,000 goes down to \$9,500 in four years. What is the average return per year to the nearest hundredth of a percent?

2.6 Skewness and the Mean, Median, and Mode

105.

The median age of the U.S. population in 1980 was 30.0 years. In 1991, the median age was 33.1 years.

- What does it mean for the median age to rise?
- Give two reasons why the median age could rise.
- For the median age to rise, is the actual number of children less in 1991 than it was in 1980? Why or why not?

2.7 Measures of the Spread of the Data

Use the following information to answer the next nine exercises: The population parameters below describe the full-time equivalent number of students (FTES) each year at Lake Tahoe Community College from 1976–1977 through 2004–2005.

- $\mu = 1000$ FTES
- median = 1,014 FTES
- $\sigma = 474$ FTES
- first quartile = 528.5 FTES
- third quartile = 1,447.5 FTES
- $n = 29$ years

106.

A sample of 11 years is taken. About how many are expected to have a FTES of 1014 or above? Explain how you determined your answer.

107.

75% of all years have an FTES:

- at or below: _____
- at or above: _____

108.

The population standard deviation = _____

109.

What percent of the FTES were from 528.5 to 1447.5? How do you know?

110.

What is the *IQR*? What does the *IQR* represent?

111.

How many standard deviations away from the mean is the median?

Additional Information: The population FTES for 2005–2006 through 2010–2011 was given in an updated report. The data are reported here.

Table 2.11.77

Year	2005–06	2006–07	2007–08	2008–09	2009–10	2010–11
Total FTES	1,585	1,690	1,735	1,935	2,021	1,890

112.

Calculate the mean, median, standard deviation, the first quartile, the third quartile and the *IQR*. Round to one decimal place.

113.

Compare the *IQR* for the FTES for 1976–77 through 2004–2005 with the *IQR* for the FTES for 2005–2006 through 2010–2011. Why do you suppose the *IQRs* are so different?

114.

Three students were applying to the same graduate school. They came from schools with different grading systems. Which student had the best GPA when compared to other students at his school? Explain how you determined your answer.

Table 2.11.78

Student	GPA	School Average GPA	School Standard Deviation
Thuy	2.7	3.2	0.8
Vichet	87	75	20
Kamala	8.6	8	0.4

115.

A music school has budgeted to purchase three musical instruments. They plan to purchase a piano costing \$3,000, a guitar costing \$550, and a drum set costing \$600. The mean cost for a piano is \$4,000 with a standard deviation of \$2,500. The mean cost for a guitar is \$500 with a standard deviation of \$200. The mean cost for drums is \$700 with a standard deviation of \$100. Which cost is the lowest, when compared to other instruments of the same type? Which cost is the highest when compared to other instruments of the same type. Justify your answer.

116.

An elementary school class ran one mile with a mean of 11 minutes and a standard deviation of three minutes. Rachel, a student in the class, ran one mile in eight minutes. A junior high school class ran one mile with a mean of nine minutes and a standard deviation of two minutes. Kenji, a student in the class, ran 1 mile in 8.5 minutes. A high school class ran one mile with a mean of seven minutes and a standard deviation of four minutes. Nedda, a student in the class, ran one mile in eight minutes.

- Why is Kenji considered a better runner than Nedda, even though Nedda ran faster than he?
- Who is the fastest runner with respect to his or her class? Explain why.

117.

The most obese countries in the world have obesity rates that range from 11.4% to 74.6%. This data is summarized in Table 2.11.79

Table 2.11.79

Percent of population obese	Number of countries
11.4–20.45	29
20.45–29.45	13
29.45–38.45	4
38.45–47.45	0
47.45–56.45	2
56.45–65.45	1
65.45–74.45	0
74.45–83.45	1

What is the best estimate of the average obesity percentage for these countries? What is the standard deviation for the listed obesity rates? The United States has an average obesity rate of 33.9%. Is this rate above average or below? How “unusual” is the United States’ obesity rate compared to the average rate? Explain.

118.

Table 2.11.80 gives the percent of children under five considered to be underweight.

Table 2.11.80

Percent of underweight children	Number of countries
16–21.45	23
21.45–26.9	4

Percent of underweight children	Number of countries
26.9–32.35	9
32.35–37.8	7
37.8–43.25	6
43.25–48.7	1

What is the best estimate for the mean percentage of underweight children? What is the standard deviation? Which interval(s) could be considered unusual? Explain.

This page titled [2.11: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **2.10: Chapter Homework** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.12: Chapter Practice

2.1 Display Data



Figure 2.12.14

14.

Construct a frequency polygon for the following:

- a. Describe the relationship between the mode and the median of this distribution.




Figure 2.12.16

67.

Describe the relationship between the mean and the median of this distribution.




Figure 2.12.17

68.




Figure 2.12.18

69.

Describe the relationship between the mode and the median of this distribution.




Figure 2.12.19

70.

Are the mean and the median the exact same in this distribution? Why or why not?




Figure 2.12.20

71.

Describe the shape of this distribution.




Figure 2.12.21

72.

Describe the relationship between the mode and the median of this distribution.




Figure 2.12.22

73.

Describe the relationship between the mean and the median of this distribution.




Figure 2.12.23

74.

The mean and median for the data are the same.

3; 4; 5; 5; 6; 6; 6; 6; 7; 7; 7; 7; 7; 7; 7

Is the data perfectly symmetrical? Why or why not?

75.

Which is the greatest, the mean, the mode, or the median of the data set?

11; 11; 12; 12; 12; 12; 13; 15; 17; 22; 22; 22

76.

Which is the least, the mean, the mode, and the median of the data set?

56; 56; 56; 58; 59; 60; 62; 64; 64; 65; 67

77.

Of the three measures, which tends to reflect skewing the most, the mean, the mode, or the median? Why?

78.

In a perfectly symmetrical distribution, when would the mode be different from the mean and median?

2.7 Measures of the Spread of the Data

Use the following information to answer the next two exercises: The following data are the distances between 20 retail stores and a large distribution center. The distances are in miles.

29; 37; 38; 40; 58; 67; 68; 69; 76; 86; 87; 95; 96; 96; 99; 106; 112; 127; 145; 150

79.

Use a graphing calculator or computer to find the standard deviation and round to the nearest tenth.

80.

Find the value that is one standard deviation below the mean.

81.

Two baseball players, Fredo and Karl, on different teams wanted to find out who had the higher batting average when compared to his team. Which baseball player had the higher batting average when compared to his team?

Table 2:59 to find the value that is three standard deviations:

Baseball player	Batting average	Team batting average	Team standard deviation
Fredo	0.158	0.166	0.012
Karl	0.177	0.189	0.015

- o Find the standard deviation for the following frequency tables using the formula. Check the calculations with the TI 83/84. **83.**

Find the standard deviation for the following frequency tables using the formula. Check the calculations with the TI 83/84.

Table 2.12.60

a.

Grade	Frequency
49.5–59.5	2
59.5–69.5	3
69.5–79.5	8
79.5–89.5	12
89.5–99.5	5

Table 2.12.61

b.

Daily low temperature	Frequency
49.5–59.5	53
59.5–69.5	32
69.5–79.5	15
79.5–89.5	1
89.5–99.5	0

Table 2.12.62

c.

Points per game	Frequency
49.5–59.5	14
59.5–69.5	32
69.5–79.5	15
79.5–89.5	23
89.5–99.5	2

This page titled [2.12: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **2.14: Chapter Practice** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.13: Homework

119.

Javier and Ercilia are supervisors at a shopping mall. Each was given the task of estimating the mean distance that shoppers live from the mall. They each randomly surveyed 100 shoppers. The samples yielded the following information.

Table 2.13.81

	Javier	Ercilia
\bar{x}	6.0 miles	6.0 miles
s	4.0 miles	7.0 miles

- How can you determine which survey was correct ?
- Explain what the difference in the results of the surveys implies about the data.
- If the two histograms depict the distribution of values for each supervisor, which one depicts Ercilia's sample? How do you know?




Figure 2.24

Use the following information to answer the next three exercises: We are interested in the number of years students in a particular elementary statistics class have lived in California. The information in the following table is from the entire section.

Table 2.13.82

Number of years	Frequency	Number of years	Frequency
7	1	22	1
14	3	23	1
15	1	26	1
18	1	40	2
19	4	42	2
20	3		
			Total = 20

120.

What is the *IQR*?

- 8
- 11
- 15
- 35

121.

What is the mode?

- 19
- 19.5
- 14 and 20
- 22.65

122.

Is this a sample or the entire population?

- a. sample
- b. entire population
- c. neither

123.

Twenty-five randomly selected students were asked the number of movies they watched the previous week. The results are as follows:

Table 2.13.83

# of movies	Frequency
0	5
1	9
2	6
3	4
4	1

- a. Find the sample mean \bar{x} .
- b. Find the approximate sample standard deviation, s .

124.

Forty randomly selected students were asked the number of pairs of sneakers they owned. Let X = the number of pairs of sneakers owned. The results are as follows:

Table 2.13.84

X	Frequency
1	2
2	5
3	8
4	12
5	12
6	0
7	1

- a. Find the sample mean \bar{x}
- b. Find the sample standard deviation, s
- c. Construct a histogram of the data.
- d. Complete the columns of the chart.
- e. Find the first quartile.
- f. Find the median.
- g. Find the third quartile.
- h. What percent of the students owned at least five pairs?
- i. Find the 40th percentile.
- j. Find the 90th percentile.
- k. Construct a line graph of the data
- l. Construct a stemplot of the data

125.

Following are the published weights (in pounds) of all of the team members of the San Francisco 49ers from a previous year.

177; 205; 210; 210; 232; 205; 185; 185; 178; 210; 206; 212; 184; 174; 185; 242; 188; 212; 215; 247; 241; 223; 220; 260; 245; 259; 278; 270; 280; 295; 275; 285; 290; 272; 273; 280; 285; 286; 200; 215; 185; 230; 250; 241; 190; 260; 250; 302; 265; 290; 276; 228; 265

- Organize the data from smallest to largest value.
- Find the median.
- Find the first quartile.
- Find the third quartile.
- The middle 50% of the weights are from _____ to _____.
- If our population were all professional football players, would the above data be a sample of weights or the population of weights? Why?
- If our population included every team member who ever played for the San Francisco 49ers, would the above data be a sample of weights or the population of weights? Why?
- Assume the population was the San Francisco 49ers. Find:
 - the population mean, μ .
 - the population standard deviation, σ .
 - the weight that is two standard deviations below the mean.
 - When Steve Young, quarterback, played football, he weighed 205 pounds. How many standard deviations above or below the mean was he?
- That same year, the mean weight for the Dallas Cowboys was 240.08 pounds with a standard deviation of 44.38 pounds. Emmitt Smith weighed in at 209 pounds. With respect to his team, who was lighter, Smith or Young? How did you determine your answer?

126.

One hundred teachers attended a seminar on mathematical problem solving. The attitudes of a representative sample of 12 of the teachers were measured before and after the seminar. A positive number for change in attitude indicates that a teacher's attitude toward math became more positive. The 12 change scores are as follows:

3; 8; -1; 2; 0; 5; -3; 1; -1; 6; 5; -2

- What is the mean change score?
- What is the standard deviation for this population?
- What is the median change score?
- Find the change score that is 2.2 standard deviations below the mean.

127.

Refer to **Figure 2.13.25** determine which of the following are true and which are false. Explain your solution to each part in complete sentences.




Figure 2.25

- The medians for both graphs are the same.
- We cannot determine if any of the means for both graphs is different.
- The standard deviation for graph b is larger than the standard deviation for graph a.
- We cannot determine if any of the third quartiles for both graphs is different.

128.

In a recent issue of the *IEEE Spectrum*, 84 engineering conferences were announced. Four conferences lasted two days. Thirty-six lasted three days. Eighteen lasted four days. Nineteen lasted five days. Four lasted six days. One lasted seven days. One lasted eight

days. One lasted nine days. Let X = the length (in days) of an engineering conference.

- Organize the data in a chart.
- Find the median, the first quartile, and the third quartile.
- Find the 65th percentile.
- Find the 10th percentile.
- The middle 50% of the conferences last from _____ days to _____ days.
- Calculate the sample mean of days of engineering conferences.
- Calculate the sample standard deviation of days of engineering conferences.
- Find the mode.
- If you were planning an engineering conference, which would you choose as the length of the conference: mean; median; or mode? Explain why you made that choice.
- Give two reasons why you think that three to five days seem to be popular lengths of engineering conferences.

129.

A survey of enrollment at 35 community colleges across the United States yielded the following figures:

6414; 1550; 2109; 9350; 21828; 4300; 5944; 5722; 2825; 2044; 5481; 5200; 5853; 2750; 10012; 6357; 27000; 9414; 7681; 3200; 17500; 9200; 7380; 18314; 6557; 13713; 17768; 7493; 2771; 2861; 1263; 7285; 28165; 5080; 11622

- Organize the data into a chart with five intervals of equal width. Label the two columns "Enrollment" and "Frequency."
- Construct a histogram of the data.
- If you were to build a new community college, which piece of information would be more valuable: the mode or the mean?
- Calculate the sample mean.
- Calculate the sample standard deviation.
- A school with an enrollment of 8000 would be how many standard deviations away from the mean?

Use the following information to answer the next two exercises. X = the number of days per week that 100 clients use a particular exercise facility.

Table 2.13.85

x	Frequency
0	3
1	12
2	33
3	28
4	11
5	9
6	4

130.

The 80th percentile is _____

- 5
- 80
- 3
- 4

131.

The number that is 1.5 standard deviations BELOW the mean is approximately _____

- 0.7
- 4.8
- 2.8
- Cannot be determined

132.

Suppose that a publisher conducted a survey asking adult consumers the number of fiction paperback books they had purchased in the previous month. The results are summarized in the Table 2.13.86

# of books	Freq.	Rel. Freq.
0	18	
1	24	
2	24	
3	22	
4	15	
5	10	
7	5	
9	1	

Table 2.86

- Are there any outliers in the data? Use an appropriate numerical test involving the IQR to identify outliers, if any, and clearly state your conclusion.
- If a data value is identified as an outlier, what should be done about it?
- Are any data values further than two standard deviations away from the mean? In some situations, statisticians may use this criteria to identify data values that are unusual, compared to the other data values. (Note that this criteria is most appropriate to use for data that is mound-shaped and symmetric, rather than for skewed data.)
- Do parts a and c of this problem give the same answer?
- Examine the shape of the data. Which part, a or c, of this question gives a more appropriate result for this data?
- Based on the shape of the data which is the most appropriate measure of center for this data: mean, median or mode?

This page titled 2.13: Homework is shared under a [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/) license and was authored, remixed, and/or curated by [OpenStax](https://openstax.org/) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- 2.8: Homework by [OpenStax](https://openstax.org/) is licensed [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.14: Chapter References

This page titled [2.14: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

2.15: Chapter Homework Solutions

1.



Figure 2.15.26

3.




Figure 2.15.27

5.




Figure 2.15.28

7.




Figure 2.15.29

9.

65

11.

The relative frequency shows the *proportion* of data points that have each value. The frequency tells the *number* of data points that have each value.

13.

Answers will vary. One possible histogram is shown:




Figure 2.15.30

15.

Find the midpoint for each class. These will be graphed on the x -axis. The frequency values will be graphed on the y -axis values.




Figure 2.15.31

17.




Figure 2.15.32

19.

- The 40th percentile is 37 years.
- The 78th percentile is 70 years.

21.

Jesse graduated 37th out of a class of 180 students. There are $180 - 37 = 143$ students ranked below Jesse. There is one rank of 37.

$x = 143$ and $y = 1$. $\frac{x+0.5y}{n}(100) = \frac{143+0.5(1)}{180}(100) = 79.72$. Jesse's rank of 37 puts him at the 80th percentile.

23.

- For runners in a race it is more desirable to have a high percentile for speed. A high percentile means a higher speed which is faster.
- 40% of runners ran at speeds of 7.5 miles per hour or less (slower). 60% of runners ran at speeds of 7.5 miles per hour or more (faster).

25.

When waiting in line at the DMV, the 85th percentile would be a long wait time compared to the other people waiting. 85% of people had shorter wait times than Mina. In this context, Mina would prefer a wait time corresponding to a lower percentile. 85% of people at the DMV waited 32 minutes or less. 15% of people at the DMV waited 32 minutes or longer.

27.

The manufacturer and the consumer would be upset. This is a large repair cost for the damages, compared to the other cars in the sample. INTERPRETATION: 90% of the crash tested cars had damage repair costs of \$1700 or less; only 10% had damage repair costs of \$1700 or more.

29.

You can afford 34% of houses. 66% of the houses are too expensive for your budget. INTERPRETATION: 34% of houses cost \$240,000 or less. 66% of houses cost \$240,000 or more.

31.

4

33.

$$6 - 4 = 2$$

35.

6

37.

$$\text{Mean: } 16 + 17 + 19 + 20 + 20 + 21 + 23 + 24 + 25 + 25 + 25 + 26 + 26 + 27 + 27 + 27 + 28 + 29 + 30 + 32 + 33 + 33 + 34 + 35 + 37 + 39 + 40 = 738$$

$$\frac{738}{27} = 27.33$$

39.

The most frequent lengths are 25 and 27, which occur three times. Mode = 25, 27

41.

4

44.

39.48 in.

45.

\$21,574

46.

15.98 ounces

47.

81.56

48.

4 hours

49.

2.01 inches

50.

18.25

51.

10

52.

14.15

53.

14

54.

14.78

55.

44%

56.

100%

57.

6%

58.

33%

59.

The data are symmetrical. The median is 3 and the mean is 2.85. They are close, and the mode lies close to the middle of the data, so the data are symmetrical.

61.

The data are skewed right. The median is 87.5 and the mean is 88.2. Even though they are close, the mode lies to the left of the middle of the data, and there are many more instances of 87 than any other number, so the data are skewed right.

63.

When the data are symmetrical, the mean and median are close or the same.

65.

The distribution is skewed right because it looks pulled out to the right.

67.

The mean is 4.1 and is slightly greater than the median, which is four.

69.

The mode and the median are the same. In this case, they are both five.

71.

The distribution is skewed left because it looks pulled out to the left.

73.

The mean and the median are both six.

75.

The mode is 12, the median is 12.5, and the mean is 15.1. The mean is the largest.

77.

The mean tends to reflect skewing the most because it is affected the most by outliers.

79.

$s = 34.5$

81.

For Fredo: $z = \frac{0.158 - 0.166}{0.012} = -0.67$

For Karl: $z = \frac{0.177 - 0.189}{0.015} = -0.8$

Fredo's z-score of -0.67 is higher than Karl's z-score of -0.8 . For batting average, higher values are better, so Fredo has a better batting average compared to his team.

83.

a. $s_x = \sqrt{\frac{\sum fm^2}{n} - \bar{x}^2} = \sqrt{\frac{193157.45}{30} - 79.5^2} = 10.88$

b. $s_x = \sqrt{\frac{\sum fm^2}{n} - \bar{x}^2} = \sqrt{\frac{38045.3}{101} - 60.94^2} = 7.62$

c. $s_x = \sqrt{\frac{\sum fm^2}{n} - \bar{x}^2} = \sqrt{\frac{440051.5}{86} - 70.66^2} = 11.14$

84.

a. Example solution for using the random number generator for the TI-84+ to generate a simple random sample of 8 states. Instructions are as follows.

- o Number the entries in the table 1–51 (Includes Washington, DC; Numbered vertically)
- o Press MATH
- o Arrow over to PRB
- o Press 5:randInt(
- o Enter 51,1,8)

Eight numbers are generated (use the right arrow key to scroll through the numbers). The numbers correspond to the numbered states (for this example: {47 21 9 23 51 13 25 4}). If any numbers are repeated, generate a different number by using 5:randInt(51,1). Here, the states (and Washington DC) are {Arkansas, Washington DC, Idaho, Maryland, Michigan, Mississippi, Virginia, Wyoming}.

Corresponding percents are {30.1, 22.2, 26.5, 27.1, 30.9, 34.0, 26.0, 25.1}




Figure 2.15.33





Figure 2.15.34

b.



c.

Figure 2.15.35

86.

Amount(\$)	Frequency	Relative frequency
51–100	5	0.08
101–150	10	0.17
151–200	15	0.25
201–250	15	0.25
251–300	10	0.17
301–350	5	0.08

Table 2.87 Singles

Amount(\$)	Frequency	Relative frequency
100–150	5	0.07
201–250	5	0.07
251–300	5	0.07
301–350	5	0.07
351–400	10	0.14
401–450	10	0.14
451–500	10	0.14
501–550	10	0.14
551–600	5	0.07
601–650	5	0.07

Table 2.88 Couples

- See Table 2.15.87 and Table 2.15.88
- In the following histogram data values that fall on the right boundary are counted in the class interval, while values that fall on the left boundary are not counted (with the exception of the first interval where both boundary values are included).




Figure 2.15.36

c. In the following histogram, the data values that fall on the right boundary are counted in the class interval, while values that fall on the left boundary are not counted (with the exception of the first interval where values on both boundaries are included).




Figure 2.15.37

d. Compare the two graphs:

i. Answers may vary. Possible answers include:

- Both graphs have a single peak.
- Both graphs use class intervals with width equal to \$50.

ii. Answers may vary. Possible answers include:

- The couples graph has a class interval with no values.
- It takes almost twice as many class intervals to display the data for couples.

iii. Answers may vary. Possible answers include: The graphs are more similar than different because the overall patterns for the graphs are the same.

e. Check student's solution.

f. Compare the graph for the Singles with the new graph for the Couples:

- i.
 - Both graphs have a single peak.
 - Both graphs display 6 class intervals.
 - Both graphs show the same general pattern.

ii. Answers may vary. Possible answers include: Although the width of the class intervals for couples is double that of the class intervals for singles, the graphs are more similar than they are different.

g. Answers may vary. Possible answers include: You are able to compare the graphs interval by interval. It is easier to compare the overall patterns with the new scale on the Couples graph. Because a couple represents two individuals, the new scale leads to a more accurate comparison.

h. Answers may vary. Possible answers include: Based on the histograms, it seems that spending does not vary much from singles to individuals who are part of a couple. The overall patterns are the same. The range of spending for couples is approximately double the range for individuals.

88.

c

90.

Answers will vary.

92.

a. $1 - (0.02 + 0.09 + 0.19 + 0.26 + 0.18 + 0.17 + 0.02 + 0.01) = 0.06$

b. $0.19 + 0.26 + 0.18 = 0.63$

c. Check student's solution.

d. 40th percentile will fall between 30,000 and 40,000

80th percentile will fall between 50,000 and 75,000

e. Check student's solution.

94.

The mean percentage, $\bar{x} = \frac{1328.65}{50} = 26.75$

95.

a. Yes

b. The sample is 0.5 higher.

96.

a. 20

b. No

97.

51

98.

a. 42

b. 99

99.

\$10.19

100.

17%

101.

\$30,772.48

102.

4.4%

103.

7.24%

104.

-1.27%

106.

The median value is the middle value in the ordered list of data values. The median value of a set of 11 will be the 6th number in order. Six years will have totals at or below the median.

108.

474 FTES

110.

919

112.

- mean = 1,809.3
- median = 1,812.5

- standard deviation = 151.2
- first quartile = 1,690
- third quartile = 1,935
- $IQR = 245$

113.

Hint: Think about the number of years covered by each time period and what happened to higher education during those periods.

115.

For pianos, the cost of the piano is 0.4 standard deviations BELOW the mean. For guitars, the cost of the guitar is 0.25 standard deviations ABOVE the mean. For drums, the cost of the drum set is 1.0 standard deviations BELOW the mean. Of the three, the drums cost the lowest in comparison to the cost of other instruments of the same type. The guitar costs the most in comparison to the cost of other instruments of the same type.

117.

- $\bar{x} = 23.32$
- Using the TI 83/84, we obtain a standard deviation of: $s_x = 12.95$.
- The obesity rate of the United States is 10.58% higher than the average obesity rate.
- Since the standard deviation is 12.95, we see that $23.32 + 12.95 = 36.27$ is the obesity percentage that is one standard deviation from the mean. The United States obesity rate is slightly less than one standard deviation from the mean. Therefore, we can assume that the United States, while 34% obese, does not have an unusually high percentage of obese people.

120.

a

122.

b

123.

- a. 1.48
- b. 1.12

125.

- a. 174; 177; 178; 184; 185; 185; 185; 185; 188; 190; 200; 205; 205; 206; 210; 210; 210; 212; 212; 215; 215; 220; 223; 228; 230; 232; 241; 241; 242; 245; 247; 250; 250; 259; 260; 260; 265; 265; 270; 272; 273; 275; 276; 278; 280; 280; 285; 285; 286; 290; 290; 295; 302
- b. 241
- c. 205.5
- d. 272.5
- e. 205.5, 272.5
- f. sample
- g. population
- h. i. 236.34
ii. 37.50
iii. 161.34
iv. 0.84 std. dev. below the mean

i. Young

127.

- a. True
- b. True
- c. True
- d. False

129.

Table 2.15.89

a.

Enrollment	Frequency
1000-5000	10
5000-10000	16
10000-15000	3
15000-20000	3
20000-25000	1
25000-30000	2

b. Check student's solution.

c. mode

d. 8628.74

e. 6943.88

f. -0.09

131.

a

This page titled [2.15: Chapter Homework Solutions](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.13: Chapter Homework Solutions](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

2.R: Descriptive Statistics (Review)

2.1 Display Data

A **stem-and-leaf plot** is a way to plot data and look at the distribution. In a stem-and-leaf plot, all data values within a class are visible. The advantage in a stem-and-leaf plot is that all values are listed, unlike a histogram, which gives classes of data values. A **line graph** is often used to represent a set of data values in which a quantity varies with time. These graphs are useful for finding trends. That is, finding a general pattern in data sets including temperature, sales, employment, company profit or cost over a period of time. A **bar graph** is a chart that uses either horizontal or vertical bars to show comparisons among categories. One axis of the chart shows the specific categories being compared, and the other axis represents a discrete value. Some bar graphs present bars clustered in groups of more than one (grouped bar graphs), and others show the bars divided into subparts to show cumulative effect (stacked bar graphs). Bar graphs are especially useful when categorical data is being used.

A **histogram** is a graphic version of a frequency distribution. The graph consists of bars of equal width drawn adjacent to each other. The horizontal scale represents classes of quantitative data values and the vertical scale represents frequencies. The heights of the bars correspond to frequency values. Histograms are typically used for large, continuous, quantitative data sets. A frequency polygon can also be used when graphing large data sets with data points that repeat. The data usually goes on y -axis with the frequency being graphed on the x -axis. Time series graphs can be helpful when looking at large amounts of data for one variable over a period of time.

2.2 Measures of the Location of the Data

The values that divide a rank-ordered set of data into 100 equal parts are called percentiles. Percentiles are used to compare and interpret data. For example, an observation at the 50th percentile would be greater than 50 percent of the other observations in the set. Quartiles divide data into quarters. The first quartile (Q_1) is the 25th percentile, the second quartile (Q_2 or median) is 50th percentile, and the third quartile (Q_3) is the 75th percentile. The interquartile range, or IQR , is the range of the middle 50 percent of the data values. The IQR is found by subtracting Q_1 from Q_3 , and can help determine outliers by using the following two expressions.

- $Q_3 + IQR(1.5)$
- $Q_1 - IQR(1.5)$

2.3 Measures of the Center of the Data

The mean and the median can be calculated to help you find the "center" of a data set. The mean is the best estimate for the actual data set, but the median is the best measurement when a data set contains several outliers or extreme values. The mode will tell you the most frequently occurring datum (or data) in your data set. The mean, median, and mode are extremely helpful when you need to analyze your data, but if your data set consists of ranges which lack specific values, the mean may seem impossible to calculate. However, the mean can be approximated if you add the lower boundary with the upper boundary and divide by two to find the midpoint of each interval. Multiply each midpoint by the number of values found in the corresponding range. Divide the sum of these values by the total number of data values in the set.

2.6 Skewness and the Mean, Median, and Mode

Looking at the distribution of data can reveal a lot about the relationship between the mean, the median, and the mode. There are three types of distributions. A **right (or positive) skewed** distribution has a shape like Figure 2.R. 11.

2.7 Measures of the Spread of the Data

The standard deviation can help you calculate the spread of data. There are different equations to use if are calculating the standard deviation of a sample or of a population.

- The Standard Deviation allows us to compare individual data or classes to the data set mean numerically.
- $s = \sqrt{\frac{\sum(x-\bar{x})^2}{n-1}}$ or $s = \sqrt{\frac{\sum f(x-\bar{x})^2}{n-1}}$ is the formula for calculating the standard deviation of a sample. To calculate the standard deviation of a population, we would use the population mean, μ , and the formula $\sigma = \sqrt{\frac{\sum(x-\mu)^2}{N}}$ or $\sigma = \sqrt{\frac{\sum f(x-\mu)^2}{N}}$.

This page titled [2.R: Descriptive Statistics \(Review\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [2.R: Descriptive Statistics \(Review\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

3: Probability Topics

You have, more than likely, used probability. In fact, you probably have an intuitive sense of probability. Probability deals with the chance of an event occurring. Whenever you weigh the odds of whether or not to do your homework or to study for an exam, you are using probability. In this chapter, you will learn how to solve probability problems using a systematic approach.

- [3.1: Introduction to Probability](#)
- [3.2: Probability Terminology](#)
- [3.3: Independent and Mutually Exclusive Events](#)
- [3.4: Two Basic Rules of Probability](#)
- [3.5: Contingency Tables and Probability Trees](#)
- [3.6: Venn Diagrams](#)
- [3.7: Chapter Key Terms](#)
- [3.8: Chapter Review](#)
- [3.9: Chapter Formula Review](#)
- [3.10: Chapter Practice](#)
- [3.11: Chapter More Practice](#)
- [3.12: Chapter Homework](#)
- [3.13: Chapter Reference](#)
- [3.14: Chapter Solution \(Practice + Homework\)](#)

This page titled [3: Probability Topics](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

3.1: Introduction to Probability

It is often necessary to "guess" about the outcome of an event in order to make a decision. Politicians study polls to guess their likelihood of winning an election. Teachers choose a particular course of study based on what they think students can comprehend. Doctors choose the treatments needed for various diseases based on their assessment of likely results. You may have visited a casino where people play games chosen because of the belief that the likelihood of winning is good. You may have chosen your course of study based on the probable availability of jobs.




Figure 3.1 Meteor showers are rare, but the probability of them occurring can be calculated. (credit: Navicore/flickr)

You have, more than likely, used probability. In fact, you probably have an intuitive sense of probability. Probability deals with the chance of an event occurring. Whenever you weigh the odds of whether or not to do your homework or to study for an exam, you are using probability. In this chapter, you will learn how to solve probability problems using a systematic approach.

This page titled [3.1: Introduction to Probability](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **3.0: Introduction to Probability** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.2: Probability Terminology

Probability is a measure that is associated with how certain we are of outcomes of a particular experiment or activity. An **experiment** is a planned operation carried out under controlled conditions. If the result is not predetermined, then the experiment is said to be a **chance** experiment. Flipping one fair coin twice is an example of an experiment.

A result of an experiment is called an **outcome**. The **sample space** of an experiment is the set of all possible outcomes. Three ways to represent a sample space are: to list the possible outcomes, to create a tree diagram, or to create a Venn diagram. The uppercase letter S is used to denote the sample space. For example, if you flip one fair coin, $S = \{H, T\}$ where H = heads and T = tails are the outcomes.

An **event** is any combination of outcomes. Upper case letters like A and B represent events. For example, if the experiment is to flip one fair coin, event A might be getting at most one head. The probability of an event A is written $P(A)$.

The **probability** of any outcome is the **long-term relative frequency** of that outcome. **Probabilities are between zero and one, inclusive** (that is, zero and one and all numbers between these values). $P(A) = 0$ means the event A can never happen. $P(A) = 1$ means the event A always happens. $P(A) = 0.5$ means the event A is equally likely to occur or not to occur. For example, if you flip one fair coin repeatedly (from 20 to 2,000 to 20,000 times) the relative frequency of heads approaches 0.5 (the probability of heads).

Equally likely means that each outcome of an experiment occurs with equal probability. For example, if you toss a **fair**, six-sided die, each face (1, 2, 3, 4, 5, or 6) is as likely to occur as any other face. If you toss a fair coin, a Head (H) and a Tail (T) are equally likely to occur. If you randomly guess the answer to a true/false question on an exam, you are equally likely to select a correct answer or an incorrect answer.

To calculate the probability of an event A when all outcomes in the sample space are equally likely, count the number of outcomes for event A and divide by the total number of outcomes in the sample space. For example, if you toss a fair dime and a fair nickel, the sample space is $\{HH, TH, HT, TT\}$ where T = tails and H = heads. The sample space has four outcomes. A = getting one head. There are two outcomes that meet this condition $\{HT, TH\}$, so $P(A) = \frac{2}{4} = 0.5$.

Suppose you roll one fair six-sided die, with the numbers $\{1, 2, 3, 4, 5, 6\}$ on its faces. Let event E = rolling a number that is at least five. There are two outcomes $\{5, 6\}$. $P(E) = \frac{2}{6}$. If you were to roll the die only a few times, you would not be surprised if your observed results did not match the probability. If you were to roll the die a very large number of times, you would expect that, overall, $\frac{2}{6}$ of the rolls would result in an outcome of "at least five". You would not expect exactly $\frac{2}{6}$. The long-term relative frequency of obtaining this result would approach the theoretical probability of $\frac{2}{6}$ as the number of repetitions grows larger and larger.

This important characteristic of probability experiments is known as the **law of large numbers** which states that as the number of repetitions of an experiment is increased, the relative frequency obtained in the experiment tends to become closer and closer to the theoretical probability. Even though the outcomes do not happen according to any set pattern or order, overall, the long-term observed relative frequency will approach the theoretical probability. (The word **empirical** is often used instead of the word observed.)

It is important to realize that in many situations, the outcomes are not equally likely. A coin or die may be **unfair**, or **biased**. Two math professors in Europe had their statistics students test the Belgian one Euro coin and discovered that in 250 trials, a head was obtained 56% of the time and a tail was obtained 44% of the time. The data seem to show that the coin is not a fair coin; more repetitions would be helpful to draw a more accurate conclusion about such bias. Some dice may be biased. Look at the dice in a game you have at home; the spots on each face are usually small holes carved out and then painted to make the spots visible. Your dice may or may not be biased; it is possible that the outcomes may be affected by the slight weight differences due to the different numbers of holes in the faces. Gambling casinos make a lot of money depending on outcomes from rolling dice, so casino dice are made differently to eliminate bias. Casino dice have flat faces; the holes are completely filled with paint having the same density as the material that the dice are made out of so that each face is equally likely to occur. Later we will learn techniques to use to work with probabilities for events that are not equally likely.

" \cup " Event: The Union

An outcome is in the event $A \cup B$ if the outcome is in A or is in B or is in both A and B . For example, let $A = \{1, 2, 3, 4, 5\}$ and $B = \{4, 5, 6, 7, 8\}$. $A \cup B = \{1, 2, 3, 4, 5, 6, 7, 8\}$. Notice that 4 and 5 are NOT listed twice.

" \cap " Event: The Intersection

An outcome is in the event $A \cap B$ if the outcome is in both A and B at the same time. For example, let A and B be $\{1, 2, 3, 4, 5\}$ and $\{4, 5, 6, 7, 8\}$ respectively. Then $A \cap B = \{4, 5\}$.

The **complement** of event A is denoted A' (read "A prime"). A' consists of all outcomes that are **NOT** in A. Notice that $P(A) + P(A') = 1$. For example, let $S = \{1, 2, 3, 4, 5, 6\}$ and let $A = \{1, 2, 3, 4\}$. Then, $A' = \{5, 6\}$. $P(A) = \frac{4}{6}$, $P(A') = \frac{2}{6}$, and $P(A) + P(A') = \frac{4}{6} + \frac{2}{6} = 1$

The **conditional probability** of A given B is written $P(A|B)$. $P(A|B)$ is the probability that event A will occur given that the event B has already occurred. **A conditional reduces the sample space.** We calculate the probability of A from the reduced sample space B. The formula to calculate $P(A|B)$ is $P(A|B) = \frac{P(A \cap B)}{P(B)}$ where $P(B)$ is greater than zero.

For example, suppose we toss one fair, six-sided die. The sample space $S = \{1, 2, 3, 4, 5, 6\}$. Let $A =$ face is 2 or 3 and $B =$ face is even (2, 4, 6). To calculate $P(A|B)$, we count the number of outcomes 2 or 3 in the sample space $B = \{2, 4, 6\}$. Then we divide that by the number of outcomes B (rather than S).

We get the same result by using the formula. Remember that S has six outcomes.

$$P(A|B) = \frac{\begin{matrix} \text{(the number of outcomes that are 2 or 3 and even in S)} \\ 6 \end{matrix}}{\begin{matrix} \text{(the number of outcomes that are even in S)} \\ 6 \end{matrix}} = \frac{\frac{1}{6}}{\frac{3}{6}} = \frac{1}{3}$$

Odds

The odds of an event presents the probability as a ratio of success to failure. This is common in various gambling formats. Mathematically, the odds of an event can be defined as:

$$\frac{P(A)}{1 - P(A)}$$

where $P(A)$ is the probability of success and of course $1 - P(A)$ is the probability of failure. Odds are always quoted as "numerator to denominator," e.g. 2 to 1. Here the probability of winning is twice that of losing; thus, the probability of winning is 0.66. A probability of winning of 0.60 would generate odds in favor of winning of 3 to 2. While the calculation of odds can be useful in gambling venues in determining payoff amounts, it is not helpful for understanding probability or statistical theory.

Understanding Terminology and Symbols

It is important to read each problem carefully to think about and understand what the events are. Understanding the wording is the first very important step in solving probability problems. Reread the problem several times if necessary. Clearly identify the event of interest. Determine whether there is a condition stated in the wording that would indicate that the probability is conditional; carefully identify the condition, if any.

Solution 3.3

1. $P(M) = 0.52$
2. $P(F) = 0.48$
3. $P(R) = 0.87$
4. $P(L) = 0.13$
5. $P(M \cap R) = 0.43$
6. $P(F \cap L) = 0.04$
7. $P(M \cup F) = 1$
8. $P(M \cup R) = 0.96$
9. $P(F \cup L) = 0.57$
10. $P(M') = 0.48$
11. $P(R|M) = 0.8269$ (rounded to four decimal places)
12. $P(F|L) = 0.3077$ (rounded to four decimal places)
13. $P(L|F) = 0.0833$

This page titled [3.2: Probability Terminology](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **3.1: Probability Terminology** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.3: Independent and Mutually Exclusive Events

Independent and mutually exclusive do **not** mean the same thing.

Independent Events

Two events are independent if one of the following are true:

- Two events A and B are **independent** if the knowledge that one occurred does not affect the chance the other occurs. For example, the outcomes of two rolls of a fair die are independent events. The outcome of the first roll does not change the probability for the outcome of the second roll. To show two events are independent, you must show **only one** of the above conditions. If two events are NOT independent, then we say that they are **dependent**.

Sampling may be done **with** replacement or **without replacement**.

- If it is not known whether A and B are independent or dependent, **assume they are dependent until you can show otherwise**.

1. Compute $P(T)$.
2. Compute $P(T|F)$.
3. Are T and F independent?
4. Are F and S mutually exclusive?
5. Are F and S independent?

This page titled [3.3: Independent and Mutually Exclusive Events](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.2: Independent and Mutually Exclusive Events](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.4: Two Basic Rules of Probability

When calculating probability, there are two rules to consider when determining if two events are independent or dependent and if they are mutually exclusive or not.

The Multiplication Rule

If A and B are two events defined on a **sample space**, then: $P(A \cap B) = P(B)P(A|B)$. We can think of the intersection symbol as substituting for the word "and".

This rule may also be written as: $P(A|B) = \frac{P(A \cap B)}{P(B)}$

This equation is read as the probability of A given B equals the probability of A and B divided by the probability of B.

If A and B are **independent**, then $P(A|B) = P(A)$. Then $P(A \cap B) = P(A|B)P(B)$ becomes $P(A \cap B) = P(A)P(B)$ because the $P(A|B) = P(A)$ if A and B are independent.

One easy way to remember the multiplication rule is that the word "and" means that the event has to satisfy two conditions. For example the name drawn from the class roster is to be both a female and a sophomore. It is harder to satisfy two conditions than only one and of course when we multiply fractions the result is always smaller. This reflects the increasing difficulty of satisfying two conditions.

The Addition Rule

If A and B are defined on a sample space, then: $P(A \cup B) = P(A) + P(B) - P(A \cap B)$. We can think of the union symbol substituting for the word "or". The reason we subtract the intersection of A and B is to keep from double counting elements that are in both A and B.

If A and B are **mutually exclusive**, then $P(A \cap B) = 0$. Then $P(A \cup B) = P(A) + P(B) - P(A \cap B)$ becomes $P(A \cup B) = P(A) + P(B)$.

A student goes to the library. Let events B = the student checks out a book and D = the student checks out a DVD. Suppose

$P(B) = 0.40$, $P(D) = 0.30$, and $P(D|B) = 0.5$ that , and .

1. Find $P(B')$.
2. Find $P(D \cap B)$.
3. Find $P(B|D)$.
4. Find $P(D \cap B')$.
5. Find $P(D|B')$.

This page titled [3.4: Two Basic Rules of Probability](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.3: Two Basic Rules of Probability](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.5: Contingency Tables and Probability Trees

Contingency Tables

A **contingency table** provides a way of portraying data that can facilitate calculating probabilities. The table helps in determining conditional probabilities quite easily. The table displays sample values in relation to two different variables that may be dependent or contingent on one another. Later on, we will use contingency tables again, but in another manner.

Example 3.5.20

Suppose a study of speeding violations and drivers who use cell phones produced the following fictional data:

Table 3.5.2

	Speeding violation in the last year	No speeding violation in the last year	Total
Uses cell phone while driving	25	280	305
Does not use cell phone while driving	45	405	450
Total	70	685	755

The total number of people in the sample is 755. The row totals are 305 and 450. The column totals are 70 and 685. Notice that $305 + 450 = 755$ and $70 + 685 = 755$.

Calculate the following probabilities using the table.

- a. Find $P(\text{Driver is a cell phone user})$.

Answer

Solution 3.20

$$\text{a. } \frac{\text{number of cell phone users}}{\text{total number in study}} = \frac{305}{755}$$

- b. Find $P(\text{Driver had no violation in the last year})$.

Answer

Solution 3.20

$$\text{b. } \frac{\text{number that had no violation}}{\text{total number in study}} = \frac{685}{755}$$

- c. Find $P(\text{Driver had no violation in the last year} \cap \text{was a cell phone user})$.

Answer

Solution 3.20

$$\text{c. } \frac{280}{755}$$

- d. Find $P(\text{Driver is a cell phone user} \cup \text{driver had no violation in the last year})$.

Answer

Solution 3.20

$$\text{d. } \left(\frac{305}{755} + \frac{685}{755} \right) - \frac{280}{755} = \frac{710}{755}$$

- e. Find $P(\text{Driver is a cell phone user} \mid \text{driver had a violation in the last year})$.

Answer

Solution 3.20

e. $\frac{25}{70}$ (The sample space is reduced to the number of drivers who had a violation.)

f. Find $P(\text{Driver had no violation last year} \mid \text{driver was not a cell phone user})$

Answer

Solution 3.20

f. $\frac{405}{450}$ (The sample space is reduced to the number of drivers who were not cell phone users.)

Exercise 3.5.20

Table 3.5.3 shows the number of athletes who stretch before exercising and how many had injuries within the past year.

	Injury in last year	No injury in last year	Total
Stretches	55	295	350
Does not stretch	231	219	450
Total	286	514	800

Table3.3

1. What is $P(\text{athlete stretches before exercising})$?
2. What is $P(\text{athlete stretches before exercising} \mid \text{no injury in the last year})$?

Example 3.5.21

Table 3.5.4 shows a random sample of 100 hikers and the areas of hiking they prefer.

Sex	The coastline	Near lakes and streams	On mountain peaks	Total
Female	18	16	—	45
Male	—	—	14	55
Total	—	41	—	—

Table3.4 Hiking Area Preference

a. Complete the table.

Answer

Solution 3.21

a.

Table 3.5.5 Hiking Area Preference

Sex	The coastline	Near lakes and streams	On mountain peaks	Total
Female	18	16	11	45
Male	16	25	14	55
Total	34	41	25	100

b. Are the events "being female" and "preferring the coastline" independent events?

Let F = being female and let C = preferring the coastline.

1. Find $P(F \cap C)$.
2. Find $P(F)P(C)$

Are these two numbers the same? If they are, then F and C are independent. If they are not, then F and C are not independent.

Answer

Solution 3.21

b.

$$1. P(F \cap C) = \frac{18}{100} = 0.18$$

$$2. P(F)P(C) = \left(\frac{45}{100}\right) \left(\frac{34}{100}\right) = (0.45)(0.34) = 0.153$$

$P(F \cap C) \neq P(F)P(C)$, so the events F and C are not independent.

c. Find the probability that a person is male given that the person prefers hiking near lakes and streams. Let M = being male, and let L = prefers hiking near lakes and streams.

1. What word tells you this is a conditional?
2. Fill in the blanks and calculate the probability: $P(___|___) = ___.$
3. Is the sample space for this problem all 100 hikers? If not, what is it?

Answer

Solution 3.21

c.

1.The word 'given' tells you that this is a conditional.

$$2.P(M|L) = \frac{25}{41}$$

3.No, the sample space for this problem is the 41 hikers who prefer lakes and streams.

d. Find the probability that a person is female or prefers hiking on mountain peaks. Let F = being female, and let P= prefers mountain peaks.

1. Find $P(F)$.
2. Find $P(P)$.
3. Find $P(F \cap P)$.
4. Find $P(F \cup P)$.

Answer

Solution 3.21

d.

$$1. P(F) = \frac{45}{100}$$

$$2. P(P) = \frac{25}{100}$$

$$3. P(F \cap P) = \frac{11}{100}$$

$$4. P(F \cup P) = \frac{45}{100} + \frac{25}{100} - \frac{11}{100} = \frac{59}{100}$$

Exercise 3.5.21

Table 3.5.6 shows a random sample of 200 cyclists and the routes they prefer. Let M = males and H = hilly path.

Table 3.5.6

Gender	Lake path	Hilly path	Wooded path	Total

Gender	Lake path	Hilly path	Wooded path	Total
Female	45	38	27	110
Male	26	52	12	90
Total	71	90	39	200

- Out of the males, what is the probability that the cyclist prefers a hilly path?
- Are the events “being male” and “preferring the hilly path” independent events?

Example 3.5.22

Muddy Mouse lives in a cage with three doors. If Muddy goes out the first door, the probability that he gets caught by Alissa the cat is $\frac{1}{15}$ and the probability he is not caught is $\frac{4}{15}$. If he goes out the second door, the probability he gets caught by Alissa is $\frac{1}{12}$ and the probability he is not caught is $\frac{3}{12}$. The probability that Alissa catches Muddy coming out of the third door is $\frac{1}{6}$ and the probability she does not catch Muddy is $\frac{5}{6}$. It is equally likely that Muddy will choose any of the three doors so the probability of choosing each door is $\frac{1}{3}$.

Table 3.5.7 Door Choice

Caught or not	Door one	Door two	Door three	Total
Caught	$\frac{1}{15}$	$\frac{1}{12}$	$\frac{1}{6}$	—
Not caught	$\frac{4}{15}$	$\frac{3}{12}$	$\frac{5}{6}$	—
Total	—	—	—	1

- The first entry $\frac{1}{15} = \left(\frac{1}{5}\right) \left(\frac{1}{3}\right)$ is $P(\text{Door One} \cap \text{Caught})$
- The entry $\frac{4}{15} = \left(\frac{4}{5}\right) \left(\frac{1}{3}\right)$ is $P(\text{Door One} \cap \text{Not Caught})$

Verify the remaining entries.

- Complete the probability contingency table. Calculate the entries for the totals. Verify that the lower-right corner entry is 1.

Answer

Solution 3.22

a.

Table 3.5.8 Door Choice

Caught or not	Door one	Door two	Door three	Total
Caught	$\frac{1}{15}$	$\frac{1}{12}$	$\frac{1}{6}$	$\frac{19}{60}$
Not caught	$\frac{4}{15}$	$\frac{3}{12}$	$\frac{5}{6}$	$\frac{41}{60}$
Total	$\frac{5}{15}$	$\frac{4}{12}$	$\frac{2}{6}$	1

- What is the probability that Alissa does not catch Muddy?

Answer

Solution 3.22

b. $\frac{41}{60}$

- What is the probability that Muddy chooses Door One \cap Door Two given that Muddy is caught by Alissa?

Answer

Solution 3.22

c. $\frac{9}{19}$

Example 3.5.23

Table 3.5.9 contains the number of crimes per 100,000 inhabitants from 2008 to 2011 in the U.S.

Table 3.5.9 United States Crime Index Rates Per 100,000 Inhabitants 2008–2011

Year	Robbery	Burglary	Rape	Vehicle	Total
2008	145.7	732.1	29.7	314.7	
2009	133.1	717.7	29.1	259.2	
2010	119.3	701	27.7	239.1	
2011	113.7	702.2	26.8	229.6	
Total					

TOTAL each column and each row. Total data = 4,520.7

1. Find $P(2009 \cap \text{Robbery})$.
2. Find $P(2010 \cap \text{Burglary})$.
3. Find $P(2010 \cup \text{Burglary})$.
4. Find $P(2011|\text{Rape})$.
5. Find $P(\text{Vehicle}|2008)$.

Answer

Solution 3.23

1. 0.0294
2. 0.1551
3. 0.7165
4. 0.2365
5. 0.2575

Exercise 3.5.23

Table 3.5.10 relates the weights and heights of a group of individuals participating in an observational study.

Table 3.5.10

Weight/height	Tall	Medium	Short	Totals
Obese	18	28	14	
Normal	20	51	28	
Underweight	12	25	9	
Totals				

1. Find the total for each row and column
2. Find the probability that a randomly chosen individual from this group is Tall.
3. Find the probability that a randomly chosen individual from this group is Obese and Tall.
4. Find the probability that a randomly chosen individual from this group is Tall given that the individual is Obese.
5. Find the probability that a randomly chosen individual from this group is Obese given that the individual is Tall.
6. Find the probability a randomly chosen individual from this group is Tall and Underweight.
7. Are the events Obese and Tall independent?

Tree Diagrams

Sometimes, when the probability problems are complex, it can be helpful to graph the situation. Tree diagrams can be used to visualize and solve conditional probabilities.

Tree Diagrams

A **tree diagram** is a special type of graph used to determine the outcomes of an experiment. It consists of "branches" that are labeled with either frequencies or probabilities. Tree diagrams can make some probability problems easier to visualize and solve. The following example illustrates how to use a tree diagram.

Example 3.5.24

In an urn, there are 11 balls. Three balls are red (R) and eight balls are blue (B). Draw two balls, one at a time, **with replacement**. "With replacement" means that you put the first ball back in the urn before you select the second ball. The tree diagram using frequencies that show all the possible outcomes follows.




Figure 3.5.2 Total = $64 + 24 + 24 + 9 = 121$

The first set of branches represents the first draw. The second set of branches represents the second draw. Each of the outcomes is distinct. In fact, we can list each red ball as R_1 , R_2 , and R_3 and each blue ball as B_1 , B_2 , B_3 , B_4 , B_5 , B_6 , B_7 , and B_8 . Then the nine RR outcomes can be written as:

R_1R_1 ; R_1R_2 ; R_1R_3 ; R_2R_1 ; R_2R_2 ; R_2R_3 ; R_3R_1 ; R_3R_2 ; R_3R_3

The other outcomes are similar.

There are a total of 11 balls in the urn. Draw two balls, one at a time, with replacement. There are $11(11) = 121$ outcomes, the size of the **sample space**.

a. List the 24 BR outcomes: B_1R_1 , B_1R_2 , B_1R_3 , ...

Answer

Solution 3.24

a. B_1R_1 ; B_1R_2 ; B_1R_3 ; B_2R_1 ; B_2R_2 ; B_2R_3 ; B_3R_1 ; B_3R_2 ; B_3R_3 ; B_4R_1 ; B_4R_2 ; B_4R_3 ; B_5R_1 ; B_5R_2 ; B_5R_3 ; B_6R_1 ; B_6R_2 ; B_6R_3 ; B_7R_1 ; B_7R_2 ; B_7R_3 ; B_8R_1 ; B_8R_2 ; B_8R_3

b. Using the tree diagram, calculate $P(RR)$.

Answer

Solution 3.24

$$b. P(RR) = \left(\frac{3}{11}\right) \left(\frac{3}{11}\right) = \frac{9}{121}$$

c. Using the tree diagram, calculate $P(RB \cup BR)$.

Answer

Solution 3.24

$$c. P(RB \cup BR) = \left(\frac{3}{11}\right) \left(\frac{8}{11}\right) + \left(\frac{8}{11}\right) \left(\frac{3}{11}\right) = \frac{48}{121}$$

d. Using the tree diagram, calculate $P(\text{Ron1stdraw} \cap \text{Bon2nddraw})$.

Answer

Solution 3.24

$$d. P(\text{Ron1stdraw} \cap \text{Bon2nddraw}) = \left(\frac{3}{11}\right) \left(\frac{8}{11}\right) = \frac{24}{121}$$

e. Using the tree diagram, calculate $P(\text{R on 2nd draw} | \text{B on 1st draw})$.

Answer

Solution 3.24

$$e. P(\text{R on 2nd draw} | \text{B on 1st draw}) = P(\text{R on 2nd} | \text{B on 1st}) = \frac{24}{88} = \frac{3}{11}$$

This problem is a conditional one. The sample space has been reduced to those outcomes that already have a blue on the first draw. There are $24 + 64 = 88$ possible outcomes (24 BR and 64 BB). Twenty-four of the 88 possible outcomes are BR.

$$\frac{24}{88} = \frac{3}{11}.$$

f. Using the tree diagram, calculate $P(\text{BB})$.

Answer

Solution 3.24

$$f. P(\text{BB}) = \frac{64}{121}$$

g. Using the tree diagram, calculate $P(\text{B on the 2nd draw} | \text{R on the first draw})$.

Answer

Solution 3.24

$$g. P(\text{B on 2nd draw} | \text{R on 1st draw}) = \frac{8}{11}$$

There are $9 + 24$ outcomes that have R on the first draw (9 RR and 24 RB). The sample space is then $9 + 24 = 33$. 24 of the 33 outcomes have B on the second draw. The probability is then $\frac{24}{33}$.

Exercise 3.5.24

In a standard deck, there are 52 cards. 12 cards are face cards (event F) and 40 cards are not face cards (event N). Draw two cards, one at a time, with replacement. All possible outcomes are shown in the tree diagram as frequencies. Using the tree diagram, calculate $P(\text{FF})$.




Figure 3.5.3

Example 3.5.25

An urn has three red marbles and eight blue marbles in it. Draw two marbles, one at a time, this time without replacement, from the urn. "**Without replacement**" means that you do not put the first ball back before you select the second marble. Following is a tree diagram for this situation. The branches are labeled with probabilities instead of frequencies. The numbers at the ends of the branches are calculated by multiplying the numbers on the two corresponding branches, for example, $\left(\frac{3}{11}\right)\left(\frac{2}{10}\right) = \frac{6}{110}$.




Figure 3.5.4 Total = $\frac{56+24+24+6}{110} = \frac{110}{110} = 1$

NOTE

If you draw a red on the first draw from the three red possibilities, there are two red marbles left to draw on the second draw. You do not put back or replace the first marble after you have drawn it. You draw **without replacement**, so that on the second draw there are ten marbles left in the urn.

Calculate the following probabilities using the tree diagram.

a. $P(RR) = \underline{\hspace{2cm}}$

Answer

Solution 3.25

a. $P(RR) = \left(\frac{3}{11}\right)\left(\frac{2}{10}\right) = \frac{6}{110}$

b. Fill in the blanks:

$$P(RB \cup BR) = \left(\frac{3}{11}\right)\left(\frac{8}{10}\right) + (\underline{\hspace{1cm}})(\underline{\hspace{1cm}}) = \frac{48}{110}$$

Answer

Solution 3.25

b. $P(RB \cup BR) = \left(\frac{3}{11}\right)\left(\frac{8}{10}\right) + \left(\frac{8}{11}\right)\left(\frac{3}{10}\right) = \frac{48}{110}$

c. $P(R \text{ on 2nd} | B \text{ on 1st}) =$

Answer

Solution 3.25

c. $P(R \text{ on 2nd} | B \text{ on 1st}) = \frac{3}{10}$

d. Fill in the blanks.

$$P(Ron1st \cap Bon2nd) = (\underline{\quad})(\underline{\quad}) = \frac{24}{100}$$

Answer

Solution 3.25

$$d. P(R \text{ on } 1st \cap B \text{ on } 2nd) = \left(\frac{3}{11}\right) \left(\frac{8}{10}\right) = \frac{24}{110}$$

e. Find P(BB).

Answer

Solution 3.25

$$e. P(BB) = \left(\frac{8}{11}\right) \left(\frac{7}{10}\right)$$


f. Find P(B on 2nd|R on 1st).

Answer

Solution 3.25

f. Using the tree diagram, P(B on 2nd|R on 1st) = P(R|B) = $\frac{8}{10}$.

If we are using probabilities, we can label the tree in the following general way.



- P(R|R) here means P(R on 2nd|R on 1st)
- P(B|R) here means P(B on 2nd|R on 1st)
- P(R|B) here means P(R on 2nd|B on 1st)
- P(B|B) here means P(B on 2nd|B on 1st)

Exercise 3.5.25

In a standard deck, there are 52 cards. Twelve cards are face cards (F) and 40 cards are not face cards (N). Draw two cards, one at a time, without replacement. The tree diagram is labeled with all possible probabilities.





Figure 3.5.5

1. Find $P(FN \cup NF)$.

2. Find $P(N|F)$.
3. Find $P(\text{at most one face card})$.
Hint: "At most one face card" means zero or one face card.
4. Find $P(\text{at least on face card})$.
Hint: "At least one face card" means one or two face cards.

Example 3.5.26

A litter of kittens available for adoption at the Humane Society has four tabby kittens and five black kittens. A family comes in and randomly selects two kittens (without replacement) for adoption.



1. What is the probability that both kittens are tabby?
a. $\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)$ b. $\left(\frac{4}{9}\right)\left(\frac{4}{9}\right)$ c. $\left(\frac{4}{9}\right)\left(\frac{3}{8}\right)$ d. $\left(\frac{4}{9}\right)\left(\frac{5}{9}\right)$
2. What is the probability that one kitten of each coloring is selected?
a. $\left(\frac{4}{9}\right)\left(\frac{5}{9}\right)$ b. $\left(\frac{4}{9}\right)\left(\frac{5}{8}\right)$ c. $\left(\frac{4}{9}\right)\left(\frac{5}{9}\right) + \left(\frac{5}{9}\right)\left(\frac{4}{9}\right)$ d. $\left(\frac{4}{9}\right)\left(\frac{5}{8}\right) + \left(\frac{5}{9}\right)\left(\frac{4}{8}\right)$
3. What is the probability that a tabby is chosen as the second kitten when a black kitten was chosen as the first?
4. What is the probability of choosing two kittens of the same color?

Answer

Solution 3.26

a. c, b. d, c. $\frac{4}{8}$, d. $\frac{32}{72}$

Exercise 3.5.26

Suppose there are four red balls and three yellow balls in a box. Two balls are drawn from the box without replacement. What is the probability that one ball of each coloring is selected?

This page titled [3.5: Contingency Tables and Probability Trees](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.4: Contingency Tables and Probability Trees](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.6: Venn Diagrams

A **Venn diagram** is a picture that represents the outcomes of an experiment. It generally consists of a box that represents the sample space S together with circles or ovals. The circles or ovals represent events. Venn diagrams also help us to convert common English words into mathematical terms that help add precision.

Venn diagrams are named for their inventor, John Venn, a mathematics professor at Cambridge and an Anglican minister. His main work was conducted during the late 1870's and gave rise to a whole branch of mathematics and a new way to approach issues of logic. We will develop the probability rules just covered using this powerful way to demonstrate the probability postulates including the Addition Rule, Multiplication Rule, Complement Rule, Independence, and Conditional Probability.

Example 3.27

Suppose an experiment has the outcomes 1, 2, 3, ..., 12 where each outcome has an equal chance of occurring. Let event $A = \{1, 2, 3, 4, 5, 6\}$ and event $B = \{6, 7, 8, 9\}$. Then $A \cap B = \{6\}$ and $A \cup B = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$. The Venn diagram is as follows:




Figure 3.6

Figure 3.6 shows the most basic relationship among these numbers. First, the numbers are in groups called sets; set A and set B. Some number are in both sets; we say in set $A \cap$ in set B. The English word "and" means inclusive, meaning having the characteristics of both A and B, or in this case, being a part of both A and B. This condition is called the INTERSECTION of the two sets. All members that are part of both sets constitute the intersection of the two sets. The intersection is written as $A \cap B$ where \cap is the mathematical symbol for intersection. The statement $A \cap B$ is read as "A intersect B." You can remember this by thinking of the intersection of two streets.

There are also those numbers that form a group that, for membership, the number must be in either one or the other group. The number does not have to be in BOTH groups, but instead only in either one of the two. These numbers are called the UNION of the two sets and in this case they are the numbers 1-5 (from A exclusively), 7-9 (from set B exclusively) and also 6, which is in both sets A and B. The symbol for the UNION is \cup , thus $A \cup B =$ numbers 1-9, but excludes number 10, 11, and 12. The values 10, 11, and 12 are part of the universe, but are not in either of the two sets.

Translating the English word "AND" into the mathematical logic symbol \cap , intersection, and the word "OR" into the mathematical symbol \cup , union, provides a very precise way to discuss the issues of probability and logic. The general terminology for the three areas of the Venn diagram in Figure 3.6 is shown in Figure 3.7.

Exercise 3.27

Suppose an experiment has outcomes black, white, red, orange, yellow, green, blue, and purple, where each outcome has an equal chance of occurring. Let event $C = \{\text{green, blue, purple}\}$ and event $P = \{\text{red, yellow, blue}\}$. Then $C \cap P = \{\text{blue}\}$ and $C \cup P = \{\text{green, blue, purple, red, yellow}\}$. Draw a Venn diagram representing this situation.

Example 3.28

Flip two fair coins. Let $A =$ tails on the first coin. Let $B =$ tails on the second coin. Then $A = \{TT, TH\}$ and $B = \{TT, HT\}$. Therefore, $A \cap B = \{TT\}$. $A \cup B = \{TH, TT, HT\}$.

The sample space when you flip two fair coins is $X = \{HH, HT, TH, TT\}$. The outcome HH is in NEITHER A NOR B. The Venn diagram is as follows:




Figure 3.7

Exercise 3.28

Roll a fair, six-sided die. Let A = a prime number of dots is rolled. Let B = an odd number of dots is rolled. Then $A = \{2, 3, 5\}$ and $B = \{1, 3, 5\}$. Therefore, $A \cap B = \{3, 5\}$. $A \cup B = \{1, 2, 3, 5\}$. The sample space for rolling a fair die is $S = \{1, 2, 3, 4, 5, 6\}$. Draw a Venn diagram representing this situation.

Example 3.29

A person with type O blood and a negative Rh factor (Rh-) can donate blood to any person with any blood type. Four percent of African Americans have type O blood and a negative RH factor, 5–10% of African Americans have the Rh- factor, and 51% have type O blood.




Figure 3.8

The “O” circle represents the African Americans with type O blood. The “Rh-“ oval represents the African Americans with the Rh- factor.

We will take the average of 5% and 10% and use 7.5% as the percent of African Americans who have the Rh- factor. Let O = African American with Type O blood and R = African American with Rh- factor.

1. $P(O) =$ _____
2. $P(R) =$ _____
3. $P(O \cap R) =$ _____
4. $P(O \cup R) =$ _____
5. In the Venn Diagram, describe the overlapping area using a complete sentence.
6. In the Venn Diagram, describe the area in the rectangle but outside both the circle and the oval using a complete sentence.

Answer

Solution 3.29

- a. 0.51; b. 0.075; c. 0.04; d. 0.545; e. The area represents the African Americans that have type O blood and the Rh- factor.
- f. The area represents the African Americans that have neither type O blood nor the Rh- factor.

Example 3.30

Fifty percent of the workers at a factory work a second job, 25% have a spouse who also works, 5% work a second job and have a spouse who also works. Draw a Venn diagram showing the relationships. Let W = works a second job and S = spouse also works.

Answer

Forty percent of the students at a local college belong to a club and **50%** work part time. **Five percent** of the students work part time and belong to a club. Draw a Venn diagram showing the relationships. Let C = student belongs to a club and PT = student works part time.




Figure 3.9

If a student is selected at random, find

- the probability that the student belongs to a club. $P(C) = 0.40$
- the probability that the student works part time. $P(PT) = 0.50$
- the probability that the student belongs to a club AND works part time. $P(C \cap PT) = 0.05$
- the probability that the student belongs to a club given that the student works part time.

$$P(C|PT) = \frac{P(C \cap PT)}{P(PT)} = \frac{0.05}{0.50} = 0.1$$

- the probability that the student belongs to a club OR works part time.

$$P(C \cup PT) = P(C) + P(PT) - P(C \cap PT) = 0.40 + 0.50 - 0.05 = 0.85$$

In order to solve [Example 3.30](#) we had to draw upon the concept of conditional probability from the previous section. There we used tree diagrams to track the changes in the probabilities, because the sample space changed as we drew without replacement. In short, conditional probability is the chance that something will happen given that some other event has already happened. Put another way, the probability that something will happen conditioned upon the situation that something else is also true. In [Example 3.30](#) the probability $P(C|PT)$ is the conditional probability that the randomly drawn student is a member of the club, conditioned upon the fact that the student also is working part time. This allows us to see the relationship between Venn diagrams and the probability postulates.

Exercise 3.30

In a bookstore, the probability that the customer buys a novel is 0.6, and the probability that the customer buys a non-fiction book is 0.4. Suppose that the probability that the customer buys both is 0.2.

1. Draw a Venn diagram representing the situation.
2. Find the probability that the customer buys either a novel or a non-fiction book.
3. In the Venn diagram, describe the overlapping area using a complete sentence.
4. Suppose that some customers buy only compact disks. Draw an oval in your Venn diagram representing this event.

Example 3.31

A set of 20 German Shepherd dogs is observed. 12 are male, 8 are female, 10 have some brown coloring, and 5 have some white sections of fur. Answer the following using Venn Diagrams.

Draw a Venn diagram simply showing the sets of male and female dogs.

Answer

Solution 3.31

The Venn diagram below demonstrates the situation of mutually exclusive events where the outcomes are independent events. If a dog cannot be both male and female, then there is no intersection. Being male precludes being female and being female precludes being male: in this case, the characteristic gender is therefore mutually exclusive. A Venn diagram shows this as two sets with no intersection. The intersection is said to be the null set using the mathematical symbol \emptyset .




Figure 3.10

Draw a second Venn diagram illustrating that 10 of the male dogs have brown coloring.

Answer

Solution 3.31

The Venn diagram below shows the overlap between male and brown where the number 10 is placed in it. This represents $\text{Male} \cap \text{Brown}$: both male and brown. This is the intersection of these two characteristics. To get the union of Male and Brown, then it is simply the two circled areas minus the overlap. In proper terms, $\text{Male} \cup \text{Brown} = \text{Male} + \text{Brown} - \text{Male} \cap \text{Brown}$ will give us the number of dogs in the union of these two sets. If we did not subtract the intersection, we would have double counted some of the dogs.




Figure 3.11

Now draw a situation depicting a scenario in which the non-shaded region represents "No white fur and female," or *White fur* \cap Female. the prime above "fur" indicates "not white fur." The prime above a set means not in that set, e.g. A' means not A. Sometimes, the notation used is a line above the letter. For example, $\overline{A} = A'$.

Answer

Solution 3.31




Figure 3.12

The Addition Rule of Probability

We met the addition rule earlier but without the help of Venn diagrams. Venn diagrams help visualize the counting process that is inherent in the calculation of probability. To restate the Addition Rule of Probability:

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

Remember that probability is simply the proportion of the objects we are interested in relative to the total number of objects. This is why we can see the usefulness of the Venn diagrams. [Example 3.31](#) shows how we can use Venn diagrams to count the number of dogs in the union of brown and male by reminding us to subtract the intersection of brown and male. We can see the effect of this directly on probabilities in the addition rule.

Example 3.32

Let's sample 50 students who are in a statistics class. 20 are freshmen and 30 are sophomores. 15 students get a "B" in the course, and 5 students both get a "B" and are freshmen.

Find the probability of selecting a student who either earns a "B" OR is a freshmen. We are translating the word OR to the mathematical symbol for the addition rule, which is the union of the two sets.

Answer

Solution 3.32

We know that there are 50 students in our sample, so we know the denominator of our fraction to give us probability. We need only to find the number of students that meet the characteristics we are interested in, i.e. any freshman and any student who earned a grade of "B." With the Addition Rule of probability, we can skip directly to probabilities.

Let "A" = the number of freshmen, and let "B" = the grade of "B." Below we can see the process for using Venn diagrams to solve this.

The $P(A) = \frac{20}{50} = 0.40$, $P(B) = \frac{15}{50} = 0.30$, and $P(A \cap B) = \frac{5}{50} = 0.10$

Therefore, $P(A \cup B) = 0.40 + 0.30 - 0.10 = 0.60$




Figure 3.13

If two events are mutually exclusive, then, like the example where we diagram the male and female dogs, the addition rule is simplified to just $P(A \cup B) = P(A) + P(B) - 0$. This is true because, as we saw earlier, the union of mutually exclusive events is the null set, \emptyset . The diagrams below demonstrate this.

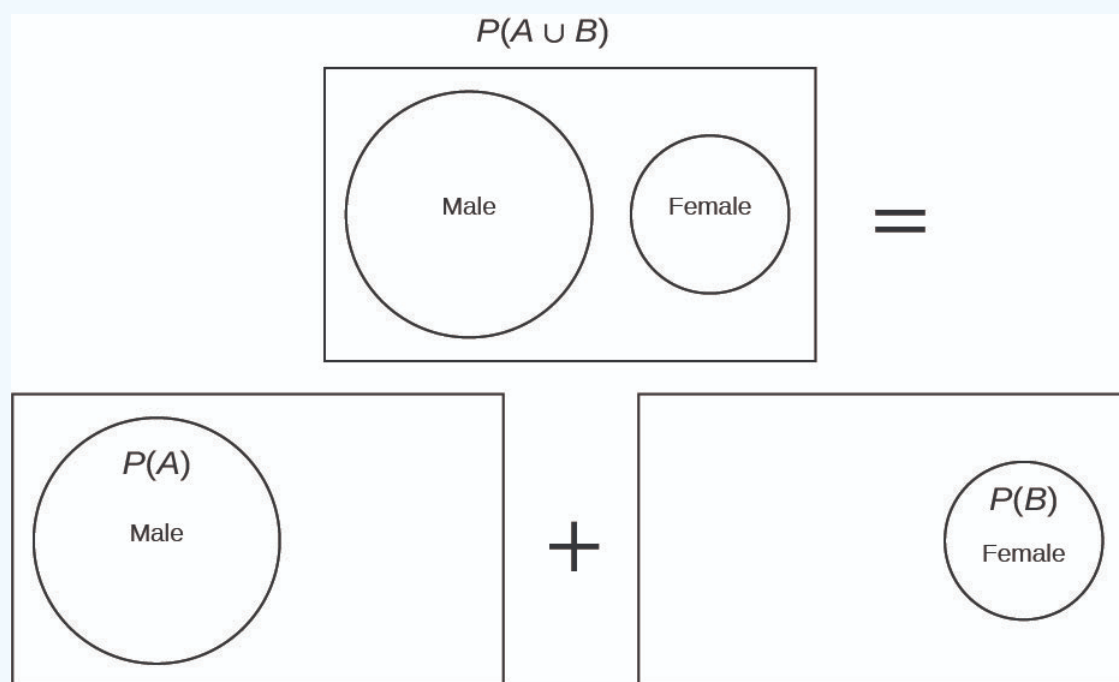


Figure 3.14

The Multiplication Rule of Probability

Restating the Multiplication Rule of Probability using the notation of Venn diagrams, we have:

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

The multiplication rule can be modified with a bit of algebra into the following conditional rule. Then Venn diagrams can then be used to demonstrate the process.

The conditional rule: $P(A|B) = \frac{P(A \cap B)}{P(B)}$

Using the same facts from [Example 3.32](#) above, find the probability that someone will earn a "B" if they are a "freshman."

$$P(A|B) = \frac{0.10}{0.30} = \frac{1}{3}$$




Figure 3.15

The multiplication rule must also be altered if the two events are independent. Independent events are defined as a situation where the conditional probability is simply the probability of the event of interest. Formally, independence of events is defined as $P(A|B) = P(A)$ or $P(B|A) = P(B)$. When flipping coins, the outcome of the second flip is independent of the outcome of the first flip; coins do not have memory. The Multiplication Rule of Probability for independent events thus becomes:

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

One easy way to remember this is to consider what we mean by the word "and." We see that the Multiplication Rule has translated the word "and" to the Venn notation for intersection. Therefore, the outcome must meet the two conditions of freshmen and grade of "B" in the above example. It is harder, less probable, to meet two conditions than just one or some other one. We can attempt to see the logic of the Multiplication Rule of probability due to the fact that fractions multiplied times each other become smaller.

The development of the Rules of Probability with the use of Venn diagrams can be shown to help as we wish to calculate probabilities from data arranged in a contingency table.

Example 3.33

[Table 3.11](#) is from a sample of 200 people who were asked how much education they completed. The columns represent the highest education they completed, and the rows separate the individuals by male and female.

Table 3.11

	Less than high school grad	High school grad	Some college	College grad	Total
Male	5	15	40	60	120
Female	8	12	30	30	80
Total	13	27	70	90	200

Now, we can use this table to answer probability questions. The following examples are designed to help understand the format above while connecting the knowledge to both Venn diagrams and the probability rules.

What is the probability that a selected person both finished college and is female?

Answer

Solution 3.33

This is a simple task of finding the value where the two characteristics intersect on the table, and then applying the postulate of probability, which states that the probability of an event is the proportion of outcomes that match the event in which we are interested as a proportion of all total possible outcomes.

$$P(\text{College Grad} \cap \text{Female}) = \frac{30}{200} = 0.15$$

What is the probability of selecting either a female or someone who finished college?

Answer

Solution 3.33

This task involves the use of the addition rule to solve for this probability.

$$P(\text{College Grad} \cup \text{Female}) = P(F) + P(CG) - P(F \cap CG)$$

$$P(\text{College Grad} \cup \text{Female}) = \frac{80}{200} + \frac{90}{200} - \frac{30}{200} = \frac{140}{200} = 0.70$$

What is the probability of selecting a high school graduate if we only select from the group of males?

Answer

Solution 3.33

Here we must use the conditional probability rule (the modified multiplication rule) to solve for this probability.

$$P(\text{HS Grad} \mid \text{Male}) = \frac{P(\text{HS Grad} \cap \text{Male})}{P(\text{Male})} = \frac{\left(\frac{15}{200}\right)}{\left(\frac{120}{200}\right)} = \frac{15}{120} = 0.125$$

Can we conclude that the level of education attained by these 200 people is independent of the gender of the person?

Answer

Solution 3.33

There are two ways to approach this test. The first method seeks to test if the intersection of two events equals the product of the events separately remembering that if two events are independent then $P(A) \cdot P(B) = P(A \cap B)$. For simplicity's sake, we can use calculated values from above.

Does $P(\text{College Grad} \cap \text{Female}) = P(CG) \cdot P(F)$?

$$\frac{30}{200} \neq \frac{90}{200} \cdot \frac{80}{200} \text{ because } 0.15 \neq 0.18.$$

Therefore, gender and education here are **not** independent.

The second method is to test if the conditional probability of A given B is equal to the probability of A. Again for simplicity, we can use an already calculated value from above.

Does $P(\text{HS Grad} \mid \text{Male}) = P(\text{HS Grad})$?

$$\frac{15}{120} \neq \frac{27}{200} \text{ because } 0.125 \neq 0.135.$$

Therefore, again gender and education here are **not** independent.

This page titled [3.6: Venn Diagrams](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.5: Venn Diagrams](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.7: Chapter Key Terms

Conditional Probability

the likelihood that an event will occur given that another event has already occurred

Contingency Table

the method of displaying a frequency distribution as a table with rows and columns to show how two variables may be dependent (contingent) upon each other; the table provides an easy way to calculate conditional probabilities.

Dependent Events

If two events are NOT independent, then we say that they are dependent.

Equally Likely

Each outcome of an experiment has the same probability.

Event

a subset of the set of all outcomes of an experiment; the set of all outcomes of an experiment is called a sample space and is usually denoted by S . An event is an arbitrary subset in S . It can contain one outcome, two outcomes, no outcomes (empty subset), the entire sample space, and the like. Standard notations for events are capital letters such as A , B , C , and so on.

Experiment

a planned activity carried out under controlled conditions

Independent Events

The occurrence of one event has no effect on the probability of the occurrence of another event. Events A and B are independent if one of the following is true:

1. $P(A|B) = P(A)$
2. $P(B|A) = P(B)$
3. $P(A \cap B) = P(A)P(B)$

Mutually Exclusive

Two events are mutually exclusive if the probability that they both happen at the same time is zero. If events A and B are mutually exclusive, then $P(A \cap B) = 0$.

Outcome

a particular result of an experiment

Probability

a number between zero and one, inclusive, that gives the likelihood that a specific event will occur; the foundation of statistics is given by the following 3 axioms (by A.N. Kolmogorov, 1930's): Let S denote the sample space and A and B are two events in S . Then:

- $0 \leq P(A) \leq 1$
- If A and B are any two mutually exclusive events, then $P(A \cup B) = P(A) + P(B)$.
- $P(S) = 1$

Sample Space

the set of all possible outcomes of an experiment

Sampling with Replacement

If each member of a population is replaced after it is picked, then that member has the possibility of being chosen more than once.

Sampling without Replacement

When sampling is done without replacement, each member of a population may be chosen only once.

The Complement Event

The complement of event A consists of all outcomes that are NOT in A .

The Conditional Probability of $A|B$

$P(A|B)$ is the probability that event A will occur given that the event B has already occurred.

The Intersection: the \cap Event

An outcome is in the event $(A \cap B)$ if the outcome is in both $A \cap B$ at the same time.

The Union: the \cup Event

An outcome is in the event $A \cup B$ if the outcome is in A or is in B or is in both A and B .

Tree Diagram

the useful visual representation of a sample space and events in the form of a “tree” with branches marked by possible outcomes together with associated probabilities (frequencies, relative frequencies)

Venn Diagram

the visual representation of a sample space and events in the form of circles or ovals showing their intersections

This page titled [3.7: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.8: Chapter Key Terms](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.8: Chapter Review

3.1 Terminology

In this module we learned the basic terminology of probability. The set of all possible outcomes of an experiment is called the sample space. Events are subsets of the sample space, and they are assigned a probability that is a number between zero and one, inclusive.

3.2 Independent and Mutually Exclusive Events

Two events A and B are independent if the knowledge that one occurred does not affect the chance the other occurs. If two events are not independent, then we say that they are dependent.

In sampling with replacement, each member of a population is replaced after it is picked, so that member has the possibility of being chosen more than once, and the events are considered to be independent. In sampling without replacement, each member of a population may be chosen only once, and the events are considered not to be independent. When events do not share outcomes, they are mutually exclusive of each other.

3.3 Two Basic Rules of Probability

The multiplication rule and the addition rule are used for computing the probability of A and B, as well as the probability of A or B for two given events A, B defined on the sample space. In sampling with replacement each member of a population is replaced after it is picked, so that member has the possibility of being chosen more than once, and the events are considered to be independent. In sampling without replacement, each member of a population may be chosen only once, and the events are considered to be not independent. The events A and B are mutually exclusive events when they do not have any outcomes in common.

3.4 Contingency Tables and Probability Trees

There are several tools you can use to help organize and sort data when calculating probabilities. Contingency tables help display data and are particularly useful when calculating probabilities that have multiple dependent variables.

A tree diagram use branches to show the different outcomes of experiments and makes complex probability questions easy to visualize.

3.5 Venn Diagrams

A Venn diagram is a picture that represents the outcomes of an experiment. It generally consists of a box that represents the sample space S or universe of the objects of interest together with circles or ovals. The circles or ovals represent groups of events called sets. A Venn diagram is especially helpful for visualizing the \cup event, the \cap event, and the complement of an event and for understanding conditional probabilities. A Venn diagram is especially helpful for visualizing an Intersection of two events, a Union of two events, or a Complement of one event. A system of Venn diagrams can also help to understand Conditional probabilities. Venn diagrams connect the brain and eyes by matching the literal arithmetic to a picture. It is important to note that more than one Venn diagram is needed to solve the probability rule formulas introduced in Section 3.3.

This page titled [3.8: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.12: Chapter Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.9: Chapter Formula Review

3.1 Terminology

A and B are events

$P(S) = 1$ where S is the sample space

$0 \leq P(A) \leq 1$

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

3.2 Independent and Mutually Exclusive Events

If A and B are independent, $P(A \cap B) = P(A)P(B)$, $P(A|B) = P(A)$ and $P(B|A) = P(B)$

If A and B are mutually exclusive, $P(A \cup B) = P(A) + P(B)$ and $P(A \cap B) = 0$

3.3 Two Basic Rules of Probability

The multiplication rule: $P(A \cap B) = P(A|B)P(B)$

The addition rule: $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

This page titled [3.9: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.6: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.10: Chapter Practice

3.1 Terminology

1.

In a particular college class, there are male and female students. Some students have long hair and some students have short hair. Write the **symbols** for the probabilities of the events for parts a through j. (Note that you cannot find numerical answers here. You were not given enough information to find any probability values yet; concentrate on understanding the symbols.)

- Use the following information to answer the next four exercises. A box is filled with several party favors. It contains 12 hats, 15 noisemakers, ten finger traps, and five bags of confetti.

Let H = the event of getting a hat.

Let N = the event of getting a noisemaker.

Let F = the event of getting a finger trap.

Let C = the event of getting a bag of confetti.

Find $P(H)$.

3.

Find $P(N)$.

4.

Find $P(F)$.

5.

Find $P(C)$.

6.

Find $P(B)$.

7.

Find $P(G)$.

8.

Find $P(P)$.

9.

Find $P(R)$.

10.

Find $P(Y)$.

11.

Find $P(O)$.

Use the following information to answer the next six exercises. There are 23 countries in North America, 12 countries in South America, 47 countries in Europe, 44 countries in Asia, 54 countries in Africa, and 14 in Oceania (Pacific Ocean region).

Let A = the event that a country is in Asia.

Let E = the event that a country is in Europe.

Let F = the event that a country is in Africa.

Let N = the event that a country is in North America.

Let O = the event that a country is in Oceania.

Let S = the event that a country is in South America.

12.

Find $P(A)$.

13.

Find $P(E)$.

14.

Find $P(F)$.

15.

Find $P(N)$.

16.

Find $P(O)$.

17.

Find $P(S)$.

18.

What is the probability of drawing a red card in a standard deck of 52 cards?

19.

What is the probability of drawing a club in a standard deck of 52 cards?

20.

What is the probability of rolling an even number of dots with a fair, six-sided die numbered one through six?

21.

What is the probability of rolling a prime number of dots with a fair, six-sided die numbered one through six?

Use the following information to answer the next two exercises. You see a game at a local fair. You have to throw a dart at a color wheel. Each section on the color wheel is equal in area.



Figure 3.10.16

Let B = the event of landing on blue.

Let R = the event of landing on red.

Let G = the event of landing on green.

Let Y = the event of landing on yellow.

22.

If you land on Y , you get the biggest prize. Find $P(Y)$.

23.

If you land on red, you don't get a prize. What is $P(R)$?

Use the following information to answer the next ten exercises. On a baseball team, there are infielders and outfielders. Some players are great hitters, and some players are not great hitters.

Let I = the event that a player is an infielder.

Let O = the event that a player is an outfielder.

Let H = the event that a player is a great hitter.
Let N = the event that a player is not a great hitter.

24.

Write the symbols for the probability that a player is not an outfielder.

25.

Write the symbols for the probability that a player is an outfielder or is a great hitter.

26.

Write the symbols for the probability that a player is an infielder and is not a great hitter.

27.

Write the symbols for the probability that a player is a great hitter, given that the player is an infielder.

28.

Write the symbols for the probability that a player is an infielder, given that the player is a great hitter.

29.

Write the symbols for the probability that of all the outfielders, a player is not a great hitter.

30.

Write the symbols for the probability that of all the great hitters, a player is an outfielder.

31.

Write the symbols for the probability that a player is an infielder or is not a great hitter.

32.

Write the symbols for the probability that a player is an outfielder and is a great hitter.

33.

Write the symbols for the probability that a player is an infielder.

34.

What is the word for the set of all possible outcomes?

35.

What is conditional probability?

36.

A shelf holds 12 books. Eight are fiction and the rest are nonfiction. Each is a different book with a unique title. The fiction books are numbered one to eight. The nonfiction books are numbered one to four. Randomly select one book

Let F = event that book is fiction

Let N = event that book is nonfiction

What is the sample space?

37.

What is the sum of the probabilities of an event and its complement?

Use the following information to answer the next two exercises. You are rolling a fair, six-sided number cube. Let E = the event that it lands on an even number. Let M = the event that it lands on a multiple of three.

38.

What does $P(E|M)$ mean in words?

39.

What does $P(E \cup M)$ mean in words?

3.2 Independent and Mutually Exclusive Events

40.

E and F are mutually exclusive events. $P(E) = 0.4$; $P(F) = 0.5$. Find $P(E|F)$

41.

J and K are independent events. $P(J|K) = 0.3$. Find $P(J)$

42.

U and V are mutually exclusive events. $P(U) = 0.26$; $P(V) = 0.37$. Find :

1. Use the following information to answer the next ten exercises. Forty-eight percent of all Californians registered voters prefer life in prison without parole over the death penalty for a person convicted of first degree murder. Among Latino California registered voters, 55% prefer life in prison without parole over the death penalty for a person convicted of first degree murder. 37.6% of all Californians are Latino.

In this problem, let:

- Suppose that one Californian is randomly selected.44.

Find $P(C)$.

45.

Find $P(L)$.

46.

Find $P(C|L)$.

47.

In words, what is $C|L$?

48.

Find $P(L \cap C)$.

49.

In words, what is $L \cap C$?

50.

Are L and C independent events? Show why or why not.

51.

Find $P(L \cup C)$.

52.

In words, what is $L \cup C$?

53.

Are L and C mutually exclusive events? Show why or why not.

3.5 Venn Diagrams

Table 3.10.12

Gender	Self-taught	Studied in school	Private instruction	Total
Female	12	38	22	72
Male	19	24	15	58
Total	31	62	37	130

54.

Find $P(\text{musician is a female})$.

55.

Find $P(\text{musician is a male} \cap \text{had private instruction})$.

56.

Find $P(\text{musician is a female} \cup \text{is self taught})$.

57.

Are the events “being a female musician” and “learning music in school” mutually exclusive events?

58.

The probability that a man develops some form of cancer in his lifetime is 0.4567. The probability that a man has at least one false positive test result (meaning the test comes back for cancer when the man does not have it) is 0.51. Let: C = a man develops cancer in his lifetime; P = man has at least one false positive. Construct a tree diagram of the situation.

This page titled [3.10: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.10: Chapter Practice](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.11: Chapter More Practice

Use the following information to answer the next seven exercises. An article in the *New England Journal of Medicine*, reported about a study of smokers in California and Hawaii. In one part of the report, the self-reported ethnicity and smoking levels per day were given. Of the people smoking at most ten cigarettes per day, there were 9,886 African Americans, 2,745 Native Hawaiians, 12,831 Latinos, 8,378 Japanese Americans, and 7,650 Whites. Of the people smoking 11 to 20 cigarettes per day, there were 6,514 African Americans, 3,062 Native Hawaiians, 4,932 Latinos, 10,680 Japanese Americans, and 9,877 Whites. Of the people smoking 21 to 30 cigarettes per day, there were 1,671 African Americans, 1,419 Native Hawaiians, 1,406 Latinos, 4,715 Japanese Americans, and 6,062 Whites. Of the people smoking at least 31 cigarettes per day, there were 759 African Americans, 788 Native Hawaiians, 800 Latinos, 2,305 Japanese Americans, and 3,970 Whites.

59.

Complete the table using the data provided. Suppose that one person from the study is randomly selected. Find the probability that person smoked 11 to 20 cigarettes per day.

Table 3.11.13 Smoking Levels by Ethnicity

Smoking level	African American	Native Hawaiian	Latino	Japanese Americans	White	TOTALS
1–10						
11–20						
21–30						
31+						
TOTALS						

60.

Suppose that one person from the study is randomly selected. Find the probability that person smoked 11 to 20 cigarettes per day.

61.

Find the probability that the person was Latino.

62.

In words, explain what it means to pick one person from the study who is “Japanese American **AND** smokes 21 to 30 cigarettes per day.” Also, find the probability.

63.

In words, explain what it means to pick one person from the study who is “Japanese American \cup smokes 21 to 30 cigarettes per day.” Also, find the probability.

64.

In words, explain what it means to pick one person from the study who is “Japanese American $|$ that person smokes 21 to 30 cigarettes per day.” Also, find the probability.

65.

Prove that smoking level/day and ethnicity are dependent events.

Use the following information to answer the next two exercises. Suppose that you have eight cards. Five are green and three are yellow. The cards are well shuffled.

66.

Suppose that you randomly draw two cards, one at a time, **with replacement**.

Let G_1 = first card is green

Let G_2 = second card is green

1. Use the following information to answer the next two exercises. The percent of licensed U.S. drivers (from a recent year) that are female is 48.60. Of the females, 5.03% are age 19 and under; 81.36% are age 20–64; 13.61% are age 65 or over. Of the

licensed U.S. male drivers, 5.04% are age 19 and under; 81.43% are age 20–64; 13.53% are age 65 or over. **68.**

Complete the following.

1. Construct a table or a tree diagram of the situation.
2. Find $P(\text{driver is female})$.
3. Find $P(\text{driver is age 65 or over} \mid \text{driver is female})$.
4. Find $P(\text{driver is age 65 or over} \cap \text{female})$.
5. In words, explain the difference between the probabilities in part c and part d.
6. Find $P(\text{driver is age 65 or over})$.
7. Are being age 65 or over and being female mutually exclusive events? How do you know?

69.

Suppose that 10,000 U.S. licensed drivers are randomly selected.

- a. How many would you expect to be male?
- b. Using the table or tree diagram, construct a contingency table of gender versus age group.
- c. Using the contingency table, find the probability that out of the age 20–64 group, a randomly selected driver is female.

70.

Approximately 86.5% of Americans commute to work by car, truck, or van. Out of that group, 84.6% drive alone and 15.4% drive in a carpool. Approximately 3.9% walk to work and approximately 5.3% take public transportation.

1. Construct a table or a tree diagram of the situation. Include a branch for all other modes of transportation to work.
2. Assuming that the walkers walk alone, what percent of all commuters travel alone to work?
3. Suppose that 1,000 workers are randomly selected. How many would you expect to travel alone to work?
4. Suppose that 1,000 workers are randomly selected. How many would you expect to drive in a carpool?

71.

When the Euro coin was introduced in 2002, two math professors had their statistics students test whether the Belgian one Euro coin was a fair coin. They spun the coin rather than tossing it and found that out of 250 spins, 140 showed a head (event H) while 110 showed a tail (event T). On that basis, they claimed that it is not a fair coin.

1. Based on the given data, find $P(H)$ and $P(T)$.
2. Use a tree to find the probabilities of each possible outcome for the experiment of tossing the coin twice.
3. Use the tree to find the probability of obtaining exactly one head in two tosses of the coin.
4. Use the tree to find the probability of obtaining at least one head.

This page titled [3.11: Chapter More Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **3.9: Chapter More Practice** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.12: Chapter Homework

3.1 Terminology

72.




Figure 3.12.17

The graph in **Figure 3.12.17** displays the sample sizes and percentages of people in different age and gender groups who were polled concerning their approval of Mayor Ford's actions in office. The total number in the sample of all the age groups is 1,045.

1. Define three events in the graph.
2. Describe in words what the entry 40 means.
3. Describe in words the complement of the entry in question 2.
4. Describe in words what the entry 30 means.
5. Out of the males and females, what percent are males?
6. Out of the females, what percent disapprove of Mayor Ford?
7. Out of all the age groups, what percent approve of Mayor Ford?
8. Find $P(\text{Approve}|\text{Male})$.
9. Out of the age groups, what percent are more than 44 years old?
10. Find $P(\text{Approve}|\text{Age} < 35)$.

73.

Explain what is wrong with the following statements. Use complete sentences.

- a. If there is a 60% chance of rain on Saturday and a 70% chance of rain on Sunday, then there is a 130% chance of rain over the weekend.
- b. The probability that a baseball player hits a home run is greater than the probability that he gets a successful hit.

3.2 Independent and Mutually Exclusive Events

Use the following information to answer the next 12 exercises. The graph shown is based on more than 170,000 interviews done by Gallup that took place from January through December 2012. The sample consists of employed Americans 18 years of age or older. The Emotional Health Index Scores are the sample space. We randomly sample one Emotional Health Index Score.




Figure 3.12.18

74.

Find the probability that an Emotional Health Index Score is 82.7.

75.

Find the probability that an Emotional Health Index Score is 81.0.

76.

Find the probability that an Emotional Health Index Score is more than 81?

77.

Find the probability that an Emotional Health Index Score is between 80.5 and 82?

78.

If we know an Emotional Health Index Score is 81.5 or more, what is the probability that it is 82.7?

79.

What is the probability that an Emotional Health Index Score is 80.7 or 82.7?

80.

What is the probability that an Emotional Health Index Score is less than 80.2 given that it is already less than 81.

81.

What occupation has the highest emotional index score?

82.

What occupation has the lowest emotional index score?

83.

What is the range of the data?

84.

Compute the average EHIS.

85.

If all occupations are equally likely for a certain individual, what is the probability that he or she will have an occupation with lower than average EHIS?

3.3 Two Basic Rules of Probability

86.

On February 28, 2013, a Field Poll Survey reported that 61% of California registered voters approved of allowing two people of the same gender to marry and have regular marriage laws apply to them. Among 18 to 39 year olds (California registered voters), the approval rating was 78%. Six in ten California registered voters said that the upcoming Supreme Court's ruling about the constitutionality of California's Proposition 8 was either very or somewhat important to them. Out of those CA registered voters who support same-sex marriage, 75% say the ruling is important to them.

In this problem, let:

- C = California registered voters who support same-sex marriage.
- B = California registered voters who say the Supreme Court's ruling about the constitutionality of California's Proposition 8 is very or somewhat important to them
- A = California registered voters who are 18 to 39 years old.

1. Find $P(C)$.
2. Find $P(B)$.
3. Find $P(C|A)$.
4. Find $P(B|C)$.
5. In words, what is $C|A$?

6. In words, what is $B|C$?
7. Find $P(C \cap B)$.
8. In words, what is $C \cap B$?
9. Find $P(C \cup B)$.
10. Are C and B mutually exclusive events? Show why or why not.

87.

After Rob Ford, the mayor of Toronto, announced his plans to cut budget costs in late 2011, the Forum Research polled 1,046 people to measure the mayor's popularity. Everyone polled expressed either approval or disapproval. These are the results their poll produced:

- In early 2011, 60 percent of the population approved of Mayor Ford's actions in office.
 - In mid-2011, 57 percent of the population approved of his actions.
 - In late 2011, the percentage of popular approval was measured at 42 percent.
- a. What is the sample size for this study?
 - b. What proportion in the poll disapproved of Mayor Ford, according to the results from late 2011?
 - c. How many people polled responded that they approved of Mayor Ford in late 2011?
 - d. What is the probability that a person supported Mayor Ford, based on the data collected in mid-2011?
 - e. What is the probability that a person supported Mayor Ford, based on the data collected in early 2011?

Use the following information to answer the next three exercises. The casino game, roulette, allows the gambler to bet on the probability of a ball, which spins in the roulette wheel, landing on a particular color, number, or range of numbers. The table used to place bets contains of 38 numbers, and each number is assigned to a color and a range.




Figure 3.12.19 (credit: film8ker/wikibooks)

88.

1. List the sample space of the 38 possible outcomes in roulette.
2. You bet on red. Find $P(\text{red})$.
3. You bet on -1st 12- (1st Dozen). Find $P(\text{-1st 12-})$.
4. You bet on an even number. Find $P(\text{even number})$.
5. Is getting an odd number the complement of getting an even number? Why?
6. Find two mutually exclusive events.
7. Are the events Even and 1st Dozen independent?

89.

Compute the probability of winning the following types of bets:

- a. Betting on two lines that touch each other on the table as in 1-2-3-4-5-6
- b. Betting on three numbers in a line, as in 1-2-3
- c. Betting on one number
- d. Betting on four numbers that touch each other to form a square, as in 10-11-13-14
- e. Betting on two numbers that touch each other on the table, as in 10-11 or 10-13
- f. Betting on 0-00-1-2-3
- g. Betting on 0-1-2; or 0-00-2; or 00-2-3

90.

Compute the probability of winning the following types of bets:

- Betting on a color
- Betting on one of the dozen groups
- Betting on the range of numbers from 1 to 18
- Betting on the range of numbers 19–36
- Betting on one of the columns
- Betting on an even or odd number (excluding zero)

91.

Suppose that you have eight cards. Five are green and three are yellow. The five green cards are numbered 1, 2, 3, 4, and 5. The three yellow cards are numbered 1, 2, and 3. The cards are well shuffled. You randomly draw one card.

- G = card drawn is green
- E = card drawn is even-numbered
 1. List the sample space.
 2. $P(G) = \underline{\hspace{2cm}}$
 3. $P(G|E) = \underline{\hspace{2cm}}$
 4. $P(G \cap E) = \underline{\hspace{2cm}}$
 5. $P(G \cup E) = \underline{\hspace{2cm}}$
 6. Are G and E mutually exclusive? Justify your answer numerically.

92.

Roll two fair dice separately. Each die has six faces.

1. List the sample space.
2. Let A be the event that either a three or four is rolled first, followed by an even number. Find $P(A)$.
3. Let B be the event that the sum of the two rolls is at most seven. Find $P(B)$.
4. In words, explain what " $P(A|B)$ " represents. Find $P(A|B)$.
5. Are A and B mutually exclusive events? Explain your answer in one to three complete sentences, including numerical justification.
6. Are A and B independent events? Explain your answer in one to three complete sentences, including numerical justification.

93.

A special deck of cards has ten cards. Four are green, three are blue, and three are red. When a card is picked, its color of it is recorded. An experiment consists of first picking a card and then tossing a coin.

1. List the sample space.
2. Let A be the event that a blue card is picked first, followed by landing a head on the coin toss. Find $P(A)$.
3. Let B be the event that a red or green is picked, followed by landing a head on the coin toss. Are the events A and B mutually exclusive? Explain your answer in one to three complete sentences, including numerical justification.
4. Let C be the event that a red or blue is picked, followed by landing a head on the coin toss. Are the events A and C mutually exclusive? Explain your answer in one to three complete sentences, including numerical justification.

94.

An experiment consists of first rolling a die and then tossing a coin.

1. List the sample space.
2. Let A be the event that either a three or a four is rolled first, followed by landing a head on the coin toss. Find $P(A)$.
3. Let B be the event that the first and second tosses land on heads. Are the events A and B mutually exclusive? Explain your answer in one to three complete sentences, including numerical justification.

95.

An experiment consists of tossing a nickel, a dime, and a quarter. Of interest is the side the coin lands on.

1. List the sample space.
2. Let A be the event that there are at least two tails. Find $P(A)$.

3. Let B be the event that the first and second tosses land on heads. Are the events A and B mutually exclusive? Explain your answer in one to three complete sentences, including justification.

96.

Consider the following scenario:

Let $P(C) = 0.4$.

Let $P(D) = 0.5$.

Let $P(C|D) = 0.6$.

1. Find $P(C \cap D)$.
2. Are C and D mutually exclusive? Why or why not?
3. Are C and D independent events? Why or why not?
4. Find $P(C \cup D)$.
5. Find $P(D|C)$.

97.

Y and Z are independent events.

1. Rewrite the basic Addition Rule $P(Y \cup Z) = P(Y) + P(Z) - P(Y \cap Z)$ using the information that Y and Z are independent events.
2. Use the rewritten rule to find $P(Z)$ if $P(Y \cup Z) = 0.71$ and $P(Y) = 0.42$.

98.

G and H are mutually exclusive events. $P(G) = 0.5$ $P(H) = 0.3$

1. Explain why the following statement MUST be false: $P(H|G) = 0.4$.
2. Find $P(H \cup G)$.
3. Are G and H independent or dependent events? Explain in a complete sentence.

99.

Approximately 281,000,000 people over age five live in the United States. Of these people, 55,000,000 speak a language other than English at home. Of those who speak another language at home, 62.3% speak Spanish.

Let: E = speaks English at home; E' = speaks another language at home; S = speaks Spanish;

Finish each probability statement by matching the correct answer.

Table 3.12.14

Probability Statements	Answers
a. $P(E')$ =	i. 0.8043
b. $P(E)$ =	ii. 0.623
c. $P(S \cap E')$ =	iii. 0.1957
d. $P(S E')$ =	iv. 0.1219

100.

1994, the U.S. government held a lottery to issue 55,000 Green Cards (permits for non-citizens to work legally in the U.S.). Renate Deutsch, from Germany, was one of approximately 6.5 million people who entered this lottery. Let G = won green card.

1. What was Renate's chance of winning a Green Card? Write your answer as a probability statement.
2. In the summer of 1994, Renate received a letter stating she was one of 110,000 finalists chosen. Once the finalists were chosen, assuming that each finalist had an equal chance to win, what was Renate's chance of winning a Green Card? Write your answer as a conditional probability statement. Let F = was a finalist.
3. Are G and F independent or dependent events? Justify your answer numerically and also explain why.
4. Are G and F mutually exclusive events? Justify your answer numerically and explain why.

101.

Three professors at George Washington University did an experiment to determine if economists are more selfish than other people. They dropped 64 stamped, addressed envelopes with \$10 cash in different classrooms on the George Washington campus. 44% were returned overall. From the economics classes 56% of the envelopes were returned. From the business, psychology, and history classes 31% were returned.

Let: R = money returned; E = economics classes; O = other classes

1. Write a probability statement for the overall percent of money returned.
2. Write a probability statement for the percent of money returned out of the economics classes.
3. Write a probability statement for the percent of money returned out of the other classes.
4. Is money being returned independent of the class? Justify your answer numerically and explain it.
5. Based upon this study, do you think that economists are more selfish than other people? Explain why or why not. Include numbers to justify your answer.

102.

The following table of data obtained from www.baseball-almanac.com shows hit information for four players. Suppose that one hit from the table is randomly selected.

Table 3.12.15

Name	Single	Double	Triple	Home run	Total hits
Babe Ruth	1,517	506	136	714	2,873
Jackie Robinson	1,054	273	54	137	1,518
Ty Cobb	3,603	174	295	114	4,189
Hank Aaron	2,294	624	98	755	3,771
Total	8,471	1,577	583	1,720	12,351

Are "the hit being made by Hank Aaron" and "the hit being a double" independent events?

1. Yes, because $P(\text{hit by Hank Aaron}|\text{hit is a double}) = P(\text{hit by Hank Aaron})$
2. No, because $P(\text{hit by Hank Aaron}|\text{hit is a double}) \neq P(\text{hit is a double})$
3. No, because $P(\text{hit is by Hank Aaron}|\text{hit is a double}) \neq P(\text{hit by Hank Aaron})$
4. Yes, because $P(\text{hit is by Hank Aaron}|\text{hit is a double}) = P(\text{hit is a double})$

103.

United Blood Services is a blood bank that serves more than 500 hospitals in 18 states. According to their website, a person with type O blood and a negative Rh factor (Rh-) can donate blood to any person with any bloodtype. Their data show that 43% of people have type O blood and 15% of people have Rh- factor; 52% of people have type O or Rh- factor.

1. Find the probability that a person has both type O blood and the Rh- factor.
2. Find the probability that a person does NOT have both type O blood and the Rh- factor.

104.

At a college, 72% of courses have final exams and 46% of courses require research papers. Suppose that 32% of courses have a research paper and a final exam. Let F be the event that a course has a final exam. Let R be the event that a course requires a research paper.

1. Find the probability that a course has a final exam or a research project.
2. Find the probability that a course has NEITHER of these two requirements.

105.

In a box of assorted cookies, 36% contain chocolate and 12% contain nuts. Of those, 8% contain both chocolate and nuts. Sean is allergic to both chocolate and nuts.

- a. Find the probability that a cookie contains chocolate or nuts (he can't eat it).
- b. Find the probability that a cookie does not contain chocolate or nuts (he can eat it).

106.

A college finds that 10% of students have taken a distance learning class and that 40% of students are part time students. Of the part time students, 20% have taken a distance learning class. Let D = event that a student takes a distance learning class and E = event that a student is a part time student

1. Find $P(D \cap E)$.
2. Find $P(E|D)$.
3. Find $P(D \cup E)$.
4. Using an appropriate test, show whether D and E are independent.
5. Using an appropriate test, show whether D and E are mutually exclusive.

3.5 Venn Diagrams

Use the information in the Table 3.12.16 to answer the next eight exercises. The table shows the political party affiliation of each of 67 members of the US Senate in June 2012, and when they are up for reelection.

Table 3.12.16

Up for reelection:	Democratic party	Republican party	Other	Total
November 2014	20	13	0	
November 2016	10	24	0	
Total				

107.

What is the probability that a randomly selected senator has an “Other” affiliation?

108.

What is the probability that a randomly selected senator is up for reelection in November 2016?

109.

What is the probability that a randomly selected senator is a Democrat and up for reelection in November 2016?

110.

What is the probability that a randomly selected senator is a Republican or is up for reelection in November 2014?

111.

Suppose that a member of the US Senate is randomly selected. Given that the randomly selected senator is up for reelection in November 2016, what is the probability that this senator is a Democrat?

112.

Suppose that a member of the US Senate is randomly selected. What is the probability that the senator is up for reelection in November 2014, knowing that this senator is a Republican?

113.

The events “Republican” and “Up for reelection in 2016” are _____

1. mutually exclusive.
2. independent.
3. both mutually exclusive and independent.
4. neither mutually exclusive nor independent.

114.

The events “Other” and “Up for reelection in November 2016” are _____

1. mutually exclusive.
2. independent.
3. both mutually exclusive and independent.
4. neither mutually exclusive nor independent.

115.

Table 3.12.17 gives the number of participants in the recent National Health Interview Survey who had been treated for cancer in the previous 12 months. The results are sorted by age, race (black or white), and sex. We are interested in possible relationships between age, race, and sex. We will let suicide victims be our population.

Table 3.12.17

Race and sex	15–24	25–40	41–65	Over 65	TOTALS
White, male	1,165	2,036	3,703		8,395
White, female	1,076	2,242	4,060		9,129
Black, male	142	194	384		824
Black, female	131	290	486		1,061
All others					
TOTALS	2,792	5,279	9,354		21,081

Do not include "all others" for parts f and g.

- Fill in the column for cancer treatment for individuals over age 65.
- Fill in the row for all other races.
- Find the probability that a randomly selected individual was a white male.
- Find the probability that a randomly selected individual was a black female.
- Find the probability that a randomly selected individual was black
- Find the probability that a randomly selected individual was male.
- Out of the individuals over age 65, find the probability that a randomly selected individual was a black or white male.

Use the following information to answer the next two exercises. The table of data obtained from www.baseball-almanac.com shows hit information for four well known baseball players. Suppose that one hit from the table is randomly selected.

Table 3.12.18

Name	Single	Double	Triple	Home run	TOTAL HITS
Babe Ruth	1,517	506	136	714	2,873
Jackie Robinson	1,054	273	54	137	1,518
Ty Cobb	3,603	174	295	114	4,189
Hank Aaron	2,294	624	98	755	3,771
TOTAL	8,471	1,577	583	1,720	12,351

116.

Find $P(\text{hit was made by Babe Ruth})$.

- $\frac{1518}{2873}$
- $\frac{2873}{12351}$
- $\frac{583}{12351}$
- $\frac{4189}{12351}$

117.

Find $P(\text{hit was made by Ty Cobb} | \text{The hit was a Home Run})$.

- $\frac{4189}{12351}$
- $\frac{114}{1720}$
- $\frac{1720}{4189}$
- $\frac{114}{12351}$

118.

Table 3.12.19 identifies a group of children by one of four hair colors, and by type of hair.

Table 3.12.19

Hair type	Brown	Blond	Black	Red	Totals
Wavy	20		15	3	43
Straight	80	15		12	
Totals		20			215

1. Complete the table.
2. What is the probability that a randomly selected child will have wavy hair?
3. What is the probability that a randomly selected child will have either brown or blond hair?
4. What is the probability that a randomly selected child will have wavy brown hair?
5. What is the probability that a randomly selected child will have red hair, given that he or she has straight hair?
6. If B is the event of a child having brown hair, find the probability of the complement of B.
7. In words, what does the complement of B represent?

119.

In a previous year, the weights of the members of the **San Francisco 49ers** and the **Dallas Cowboys** were published in the *San Jose Mercury News*. The factual data were compiled into the following table.

Table 3.12.20

Shirt #	≤ 210	211–250	251–290	> 290
1–33	21	5	0	290" class="lt-stats-5547">0
34–66	6	18	7	290" class="lt-stats-5547">4
66–99	6	12	22	290" class="lt-stats-5547">5

For the following, suppose that you randomly select one player from the 49ers or Cowboys.

1. Find the probability that his shirt number is from 1 to 33.
2. Find the probability that he weighs at most 210 pounds.
3. Find the probability that his shirt number is from 1 to 33 AND he weighs at most 210 pounds.
4. Find the probability that his shirt number is from 1 to 33 OR he weighs at most 210 pounds.
5. Find the probability that his shirt number is from 1 to 33 GIVEN that he weighs at most 210 pounds.

Use the following information to answer the next two exercises. This tree diagram shows the tossing of an unfair coin followed by drawing one bead from a cup containing three red (R), four yellow (Y) and five blue (B) beads. For the coin, $P(H) = \frac{2}{3}$ and $P(T) = \frac{1}{3}$ where H is heads and T is tails.




Figure 3.12.20

120.

Find P(tossing a Head on the coin AND a Red bead)

1. $\frac{2}{3}$
2. $\frac{5}{15}$
3. $\frac{6}{36}$
4. $\frac{5}{36}$

121.

Find P(Blue bead).

1. $\frac{15}{36}$
2. $\frac{10}{36}$
3. $\frac{10}{12}$
4. $\frac{6}{36}$

122.

A box of cookies contains three chocolate and seven butter cookies. Miguel randomly selects a cookie and eats it. Then he randomly selects another cookie and eats it. (How many cookies did he take?)

1. Draw the tree that represents the possibilities for the cookie selections. Write the probabilities along each branch of the tree.
2. Are the probabilities for the flavor of the SECOND cookie that Miguel selects independent of his first selection? Explain.
3. For each complete path through the tree, write the event it represents and find the probabilities.
4. Let S be the event that both cookies selected were the same flavor. Find P(S).
5. Let T be the event that the cookies selected were different flavors. Find P(T) by two different methods: by using the complement rule and by using the branches of the tree. Your answers should be the same with both methods.
6. Let U be the event that the second cookie selected is a butter cookie. Find P(U).

This page titled [3.12: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.7: Chapter Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.13: Chapter Reference

3.1 Terminology

“Countries List by Continent.” Worldatlas, 2013. Available online at <http://www.worldatlas.com/cntycont.htm> (accessed May 2, 2013).

3.2 Independent and Mutually Exclusive Events

Lopez, Shane, Preeti Sidhu. “U.S. Teachers Love Their Lives, but Struggle in the Workplace.” Gallup Wellbeing, 2013. <http://www.gallup.com/poll/161516/te...workplace.aspx> (accessed May 2, 2013).

Data from Gallup. Available online at www.gallup.com/ (accessed May 2, 2013).

3.3 Two Basic Rules of Probability

DiCamillo, Mark, Mervin Field. “The File Poll.” Field Research Corporation. Available online at www.field.com/fieldpollonline...rs/Rls2443.pdf (accessed May 2, 2013).

Rider, David, “Ford support plummeting, poll suggests,” The Star, September 14, 2011. Available online at http://www.thestar.com/news/gta/2011..._suggests.html (accessed May 2, 2013).

“Mayor’s Approval Down.” News Release by Forum Research Inc. Available online at www.forumresearch.com/forms/NewsArchives/NewsReleases/74209_TO_Issues_-_Mayoral_Approval_%28Forum_Research%29%2820130320%29.pdf (accessed May 2, 2013).

“Roulette.” Wikipedia. Available online at <http://en.Wikipedia.org/wiki/Roulette> (accessed May 2, 2013).

Shin, Hyon B., Robert A. Kominski. “Language Use in the United States: 2007.” United States Census Bureau. Available online at www.census.gov/hhes/socdemo/l...acs/ACS-12.pdf (accessed May 2, 2013).

Data from the Baseball-Almanac, 2013. Available online at www.baseball-almanac.com (accessed May 2, 2013).

Data from U.S. Census Bureau.

Data from the Wall Street Journal.

Data from The Roper Center: Public Opinion Archives at the University of Connecticut. Available online at www.ropercenter.uconn.edu/ (accessed May 2, 2013).

Data from Field Research Corporation. Available online at www.field.com/fieldpollonline (accessed May 2, 2013).

3.4 Contingency Tables and Probability Trees

“Blood Types.” American Red Cross, 2013. Available online at <http://www.redcrossblood.org/learn-a...od/blood-types> (accessed May 3, 2013).

Data from the National Center for Health Statistics, part of the United States Department of Health and Human Services.

Data from United States Senate. Available online at www.senate.gov (accessed May 2, 2013).

“Human Blood Types.” Unite Blood Services, 2011. Available online at <https://www.vitalant.org/Donate/Blood-Donation/Donate-Blood-Overview.aspx> (accessed May 2, 2013).

Haiman, Christopher A., Daniel O. Stram, Lynn R. Wilkens, Malcom C. Pike, Laurence N. Kolonel, Brien E. Henderson, and Loïc Le Marchand. “Ethnic and Racial Differences in the Smoking-Related Risk of Lung Cancer.” The New England Journal of Medicine, 2013. Available online at <http://www.nejm.org/doi/full/10.1056/NEJMoa033250> (accessed May 2, 2013).

Samuel, T. M. “Strange Facts about RH Negative Blood.” eHow Health, 2013. Available online at http://www.ehow.com/facts_5552003_st...ive-blood.html (accessed May 2, 2013).

“United States: Uniform Crime Report – State Statistics from 1960–2011.” The Disaster Center. Available online at <http://www.disastercenter.com/crime/> (accessed May 2, 2013).

Data from Clara County Public H.D.

Data from the American Cancer Society.

Data from The Data and Story Library, 1996. Available online at <http://lib.stat.cmu.edu/DASL/> (accessed May 2, 2013).

Data from the Federal Highway Administration, part of the United States Department of Transportation.

Data from the United States Census Bureau, part of the United States Department of Commerce.

Data from USA Today.

“Environment.” The World Bank, 2013. Available online at <http://data.worldbank.org/topic/environment> (accessed May 2, 2013).

“Search for Datasets.” Roper Center: Public Opinion Archives, University of Connecticut., 2013. Available online at <https://ropcenter.cornell.edu/?s=S...h+for+Datasets> (accessed February 6, 2019).

This page titled [3.13: Chapter Reference](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **3.11: Chapter Reference** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

3.14: Chapter Solution (Practice + Homework)

1.

1. $P(L') = P(S)$
2. $P(M \cup S)$
3. $P(F \cap L)$
4. $P(M|L)$
5. $P(L|M)$
6. $P(S|F)$
7. $P(F|L)$
8. $P(F \cup L)$
9. $P(M \cap S)$
10. $P(F)$

3.

$$P(N) = \frac{15}{42} = \frac{5}{14} = 0.36$$

5.

$$P(C) = \frac{5}{42} = 0.12$$

7.

$$P(G) = \frac{20}{150} = \frac{2}{15} = 0.13$$

9.

$$P(R) = \frac{22}{150} = \frac{11}{75} = 0.15$$

11.

$$P(O) = \frac{150 - 22 - 38 - 20 - 28 - 26}{150} = \frac{16}{150} = \frac{8}{75} = 0.11$$

13.

$$P(E) = \frac{47}{194} = 0.24$$

15.

$$P(N) = \frac{23}{194} = 0.12$$

17.

$$P(S) = \frac{12}{194} = \frac{6}{97} = 0.06$$

19.

$$\frac{13}{52} = \frac{1}{4} = 0.25$$

21.

$$\frac{3}{6} = \frac{1}{2} = 0.5$$

23.

$$P(R) = \frac{4}{8} = 0.5$$

25.

$$P(O \cup H)$$

27.

$$P(H|I)$$

29.

$$P(N|O)$$

31.

$$P(I \cup N)$$

33.

$$P(I)$$

35.

The likelihood that an event will occur given that another event has already occurred.

37.

1

39.

the probability of landing on an even number or a multiple of three

41.

$$P(J) = 0.3$$

43.

$$P(Q \cap R) = P(Q)P(R)$$

$$0.1 = (0.4)P(R)$$

$$P(R) = 0.25$$

45.

0.376

47.

$C|L$ means, given the person chosen is a Latino Californian, the person is a registered voter who prefers life in prison without parole for a person convicted of first degree murder.

49.

$L \setminus C$ is the event that the person chosen is a Latino California registered voter who prefers life without parole over the death penalty for a person convicted of first degree murder.

51.

0.6492

53.

No, because $P(L \setminus C)$ does not equal 0.

55.

$$P(\text{musician is a male} \cap \text{had private instruction}) = \frac{15}{130} = \frac{3}{26} = 0.12.$$

57.

The events are not mutually exclusive. It is possible to be a female musician who learned music in school.

58.




Figure 3.14.21

60.

$$\frac{35,065}{100,450}$$


62.

To pick one person from the study who is Japanese American AND smokes 21 to 30 cigarettes per day means that the person has to meet both criteria: both Japanese American and smokes 21 to 30 cigarettes. The sample space should include everyone in the study. The probability is $\frac{4,715}{100,450}$.

64.

To pick one person from the study who is Japanese American given that person smokes 21-30 cigarettes per day, means that the person must fulfill both criteria and the sample space is reduced to those who smoke 21-30 cigarettes per day. The probability is $\frac{4715}{15,273}$.

66.



1.

Figure 3.14.22

$$2. P(GG) = \left(\frac{5}{8}\right) \left(\frac{5}{8}\right) = \frac{25}{64}$$

$$3. P(\text{at least one green}) = P(GG) + P(GY) + P(YG) = \frac{25}{64} + \frac{15}{64} + \frac{15}{64} = \frac{55}{64}$$

$$4. P(G|G) = \frac{5}{8}$$

5. Yes, they are independent because the first card is placed back in the bag before the second card is drawn; the composition of cards in the bag remains the same from draw one to draw two.

68.

	<20>	20–64	>64	Totals
Female	" class="lt-stats-5549">0.0244	0.3954	64" class="lt-stats-5549">64">0.0661	0.486

	<20>	20–64	>64	Totals
Male	" class="lt-stats-5549">0.0259	0.4186	64" class="lt-stats-5549">64">0.0695	0.514
Totals	" class="lt-stats-5549">0.0503	0.8140	64" class="lt-stats-5549">64">0.1356	1

Table3.22

- $P(F) = 0.486$
- $P(> 64|F) = 0.1361$
- $P(> 64 \text{ and } F) = P(F)P(> 64|F) = (0.486)(0.1361) = 0.0661$
- $P(> 64|F)$ is the percentage of female drivers who are 65 or older and $P(>64 \cap F)$ is the percentage of drivers who are female and 65 or older.
- $P(> 64) = P(> 64 \cap F) + P(> 64 \cap M) = 0.1356$
- No, being female and 65 or older are not mutually exclusive because they can occur at the same time $P(> 64 \cap F) = 0.0661$.

70.

	Car, truck or van	Walk	Public transportation	Other	Totals
Alone	0.7318				
Not alone	0.1332				
Totals	0.8650	0.0390	0.0530	0.0430	1

Table3.23

- If we assume that all walkers are alone and that none from the other two groups travel alone (which is a big assumption) we have: $P(\text{Alone}) = 0.7318 + 0.0390 = 0.7708$
- Make the same assumptions as in (b) we have: $(0.7708)(1,000) = 771$
- $(0.1332)(1,000) = 133$

73.

- You can't calculate the joint probability knowing the probability of both events occurring, which is not in the information given; the probabilities should be multiplied, not added; and probability is never greater than 100%
- A home run by definition is a successful hit, so he has to have at least as many successful hits as home runs.

75.

0

77.

0.3571

79.

0.2142

81.

Physician (83.7)

83.

$83.7 - 79.6 = 4.1$

85.

$P(\text{Occupation} < 81.3) = 0.5$

87.

- The Forum Research surveyed 1,046 Torontonians.

2. 58%
3. 42% of 1,046 = 439 (rounding to the nearest integer)
4. 0.57
5. 0.60.

89.

1. $P(\text{ Betting on two line that touch each other on the table }) = \frac{6}{38}$.
2. $P(\text{ Betting on three numbers in a line }) = \frac{3}{38}$
3. $P(\text{ Betting on one number }) = \frac{1}{38}$
4. $P(\text{ Betting on four number that touch each other to form a square }) = \frac{4}{38}$.
5. $P(\text{ Betting on two number that touch each other on the table }) = \frac{2}{38}$
6. $P(\text{ Betting on } 0 - 00 - 1 - 2 - 3) = \frac{5}{38}$
7. $P(\text{ Betting on } 0 - 1 - 2; \text{ or } 0 - 00 - 2; \text{ or } 00 - 2 - 3) = \frac{3}{38}$

91.

1. $\{G1, G2, G3, G4, G5, Y1, Y2, Y3\}$
2. $\frac{5}{8}$
3. $\frac{2}{3}$
4. $\frac{12}{81}$
5. $\frac{6}{8}$
6. No, because $P(G \cap E)$ does not equal 0.

93.

NOTE

The coin toss is independent of the card picked first.

1. $\{(G, H)(G, T)(B, H)(B, T)(R, H)(R, T)\}$
2. $P(A) = P(\text{ blue })P(\text{ head }) = \left(\frac{3}{10}\right)\left(\frac{1}{2}\right) = \frac{3}{20}$
3. Yes, A and B are mutually exclusive because they cannot happen at the same time; you cannot pick a card that is both blue and also (red or green). $P(A \cap B) = 0$
4. No, A and C are not mutually exclusive because they can occur at the same time. In fact, C includes all of the outcomes of A; if the card chosen is blue it is also (red or blue). $P(A \cap C) = P(A) = \frac{3}{20}$

95.

1. $S = \{(HHH), (HHT), (HTH), (HTT), (THH), (THT), (TTH), (TTT)\}$
2. $\frac{4}{8}$
3. Yes, because if A has occurred, it is impossible to obtain two tails. In other words, $P(A \cap B) = 0$.

97.

1. If Y and Z are independent, then $P(Y \cap Z) = P(Y)P(Z)$, so $P(Y \cup Z) = P(Y) + P(Z) - P(Y)P(Z)$.
2. 0.5

99.

iii i iv ii

101.

1. $P(R) = 0.44$
2. $P(R|E) = 0.56$
3. $P(R|O) = 0.31$
4. No, whether the money is returned is not independent of which class the money was placed in. There are several ways to justify this mathematically, but one is that the money placed in economics classes is not returned at the same overall rate; $P(R|E) \neq P(R)$.

5. No, this study definitely does not support that notion; in fact, it suggests the opposite. The money placed in the economics classrooms was returned at a higher rate than the money place in all classes collectively; $P(R|E) > P(R)$.

103.

1. $P(\text{type O} \cup \text{Rh-}) = P(\text{type O}) + P(\text{Rh-}) - P(\text{type O} \cap \text{Rh-})$

$0.52 = 0.43 + 0.15 - P(\text{type O} \cap \text{Rh-})$; solve to find $P(\text{type O} \cap \text{Rh-}) = 0.06$

6% of people have type O, Rh- blood

2. $P(\text{NOT}(\text{type O} \cap \text{Rh-})) = 1 - P(\text{type O} \cap \text{Rh-}) = 1 - 0.06 = 0.94$

94% of people do not have type O, Rh- blood

105.

1. Let C = be the event that the cookie contains chocolate. Let N = the event that the cookie contains nuts.

2. $P(C \cup N) = P(C) + P(N) - P(C \cap N) = 0.36 + 0.12 - 0.08 = 0.40$

3. $P(\text{NEITHER chocolate NOR nuts}) = 1 - P(C \cup N) = 1 - 0.40 = 0.60$

107.

0

109.

$\frac{10}{67}$

111.

$\frac{10}{34}$

113.

d

115.

1.

Race and sex	1-14	15-24	25-64	Over 64	TOTALS
White, male	210	3,360	13,610	4,870	22,050
White, female	80	580	3,380	890	4,930
Black, male	10	460	1,060	140	1,670
Black, female	0	40	270	20	330
All others				100	
TOTALS	310	4,650	18,780	6,020	29,760

Table3.24

2.

Race and sex	1-14	15-24	25-64	Over 64	TOTALS
White, male	210	3,360	13,610	4,870	22,050
White, female	80	580	3,380	890	4,930
Black, male	10	460	1,060	140	1,670
Black, female	0	40	270	20	330
All others	10	210	460	100	780
TOTALS	310	4,650	18,780	6,020	29,760

Table3.25

3. $\frac{22,050}{29,760}$

4. $\frac{330}{29,760}$
5. $\frac{2,000}{29,760}$
6. $\frac{23,720}{29,760}$
7. $\frac{5,010}{6,020}$

117.

b

119.

1. $\frac{26}{106}$
2. $\frac{33}{106}$
3. $\frac{21}{106}$
4. $\left(\frac{26}{106}\right) + \left(\frac{33}{106}\right) - \left(\frac{21}{106}\right) = \left(\frac{38}{106}\right)$
5. $\frac{21}{33}$

121.

a

This page titled [3.14: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [3.13: Chapter Solution \(Practice + Homework\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

4: Discrete Random Variables

- [4.1: Introduction to Discrete Random Variables](#)
- [4.2: Hypergeometric Distribution](#)
- [4.3: Binomial Distribution](#)
- [4.4: Geometric Distribution](#)
- [4.5: Poisson Distribution](#)
- [4.6: Chapter Key Items](#)
- [4.7: Chapter Review](#)
- [4.8: Chapter Formula Review](#)
- [4.9: Chapter Practice](#)
- [4.10: Chapter Homework](#)
- [4.11: Chapter References](#)
- [4.12: Chapter Solution \(Practice + Homework\)](#)

This page titled [4: Discrete Random Variables](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

4.1: Introduction to Discrete Random Variables



Figure 4.1.1 You can use probability and discrete random variables to calculate the likelihood of lightning striking the ground five times during a half-hour thunderstorm. (Credit: Leszek Leszczynski)

A student takes a ten-question, true-false quiz. Because the student had such a busy schedule, he or she could not study and guesses randomly at each answer. What is the probability of the student passing the test with at least a 70%?

Small companies might be interested in the number of long-distance phone calls their employees make during the peak time of the day. Suppose the historical average is 20 calls. What is the probability that the employees make more than 20 long-distance phone calls during the peak time?

These two examples illustrate two different types of probability problems involving discrete random variables. Recall that discrete data are data that you can count, that is, the random variable can only take on whole number values. A **random variable** describes the outcomes of a statistical experiment in words. The values of a random variable can vary with each repetition of an experiment, often called a trial.

Random Variable Notation

The upper case letter X denotes a random variable. Lower case letters like x or y denote the value of a random variable. If X is a **random variable, then X is written in words, and x is given as a number.**

For example, let X = the number of heads you get when you toss three fair coins. The sample space for the toss of three fair coins is $TTT; THH; HTH; HHT; HTT; THT; TTH; HHH$. Then, $x = 0, 1, 2, 3$. X is in words and x is a number. Notice that for this example, the x values are countable outcomes. Because you can count the possible values as whole numbers that X can take on and the outcomes are random (the x values 0, 1, 2, 3), X is a discrete random variable.

Probability Density Functions (PDF) for a Random Variable

A **probability density function or probability distribution function** has two characteristics:

1. A probability density function is a mathematical formula that calculates probabilities for specific types of events, what we have been calling experiments. There is a sort of magic to a probability density function (Pdf) partially because the same formula often describes very different types of events. For example, the binomial Pdf will calculate probabilities for flipping coins, yes/no questions on an exam, opinions of voters in an up or down opinion poll, indeed any binary event. Other probability density functions will provide probabilities for the time until a part will fail, when a customer will arrive at the turnpike booth, the number of telephone calls arriving at a central switchboard, the growth rate of a bacterium, and on and on. There are whole families of probability density functions that are used in a wide variety of applications, including medicine, business and finance, physics and engineering, among others.

For our needs here we will concentrate on only a few probability density functions as we develop the tools of inferential statistics.

Counting Formulas and the Combinational Formula

As an equation this is:

$$P(A) = \frac{\text{number of ways to get A}}{\text{Total number of possible outcomes}} \quad (4.1.1)$$

When we looked at the sample space for flipping 3 coins we could easily write the full sample space and thus could easily count the number of events that met our desired result, e.g. $x = 1$, where X is the random variable defined as the number of heads.

As we have larger numbers of items in the sample space, such as a full deck of 52 cards, the ability to write out the sample space becomes impossible.

We see that probabilities are nothing more than counting the events in each group we are interested in and dividing by the number of elements in the universe, or sample space. This is easy enough if we are counting sophomores in a Stat class, but in more complicated cases listing all the possible outcomes may take a life time. There are, for example, 36 possible outcomes from throwing just two six-sided dice where the random variable is the sum of the number of spots on the up-facing sides. If there were four dice then the total number of possible outcomes would become 1,296. There are more than 2.5 MILLION possible 5 card poker hands in a standard deck of 52 cards. Obviously keeping track of all these possibilities and counting them to get at a single probability would be tedious at best.

An alternative to listing the complete sample space and counting the number of elements we are interested in, is to skip the step of listing the sample space, and simply figuring out the number of elements in it and doing the appropriate division. If we are after a probability we really do not need to see each and every element in the sample space, we only need to know how many elements are there. Counting formulas were invented to do just this. They tell us the number of unordered subsets of a certain size that can be created from a set of unique elements. By unordered it is meant that, for example, when dealing cards, it does not matter if you got {ace, ace, ace, ace, king} or {king, ace, ace, ace, ace} or {ace, king, ace, ace, ace} and so on. Each of these subsets are the same because they each have 4 aces and one king.

Combinational Formula

$$\binom{n}{x} = {}_n C_x = \frac{n!}{x!(n-x)!}$$

This is the formula that tells the number of unique unordered subsets of size x that can be created from n unique elements. The formula is read “ n combinatorial x ”. Sometimes it is read as “ n choose x .” The exclamation point “!” is called a factorial and tells us to take all the numbers from 1 through the number before the ! and multiply them together thus $4!$ is $1 \cdot 2 \cdot 3 \cdot 4 = 24$. By definition $0! = 1$. The formula is called the Combinatorial Formula. It is also called the Binomial Coefficient, for reasons that will be clear shortly. While this mathematical concept was understood long before 1653, Blaise Pascal is given major credit for his proof that he published in that year. Further, he developed a generalized method of calculating the values for combinatorials known to us as the Pascal Triangle. Pascal was one of the geniuses of an era of extraordinary intellectual advancement which included the work of Galileo, Rene Descartes, Isaac Newton, William Shakespeare and the refinement of the scientific method, the very rationale for the topic of this text.

Let’s find the hard way the total number of combinations of the four aces in a deck of cards if we were going to take them two at a time. The sample space would be:

$$S = \{(\text{Spade, Heart}), (\text{Spade, Diamond}), (\text{Spade, Club}), (\text{Diamond, Club}), (\text{Heart, Diamond}), (\text{Heart, Club})\}$$

There are 6 combinations; formally, six unique unordered subsets of size 2 that can be created from 4 unique elements. To use the combinatorial formula we would solve the formula as follows:

$$\binom{4}{2} = \frac{4!}{(4-2)!2!} = \frac{4 \cdot 3 \cdot 2 \cdot 1}{2 \cdot 1 \cdot 2 \cdot 1} = 6$$

If we wanted to know the number of unique 5 card poker hands that could be created from a 52 card deck we simply compute:

$$\binom{52}{5}$$

where 52 is the total number of unique elements from which we are drawing and 5 is the size group we are putting them into.

With the combinatorial formula we can count the number of elements in a sample space without having to write each one of them down, truly a lifetime’s work for just the number of 5 card hands from a deck of 52 cards. We can now apply this tool to a very important probability density function, the hypergeometric distribution.

Remember, a probability density function computes probabilities for us. We simply put the appropriate numbers in the formula and we get the probability of specific events. However, for these formulas to work they must be applied only to cases for which

they were designed.

This page titled [4.1: Introduction to Discrete Random Variables](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **4.0: Introduction to Discrete Random Variables** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.2: Hypergeometric Distribution

The simplest probability density function is the hypergeometric. This is the most basic one because it is created by combining our knowledge of probabilities from Venn diagrams, the addition and multiplication rules, and the combinatorial counting formula.

To find the number of ways to get 2 aces from the four in the deck we computed:

$$\binom{4}{2} = \frac{4!}{2!(4-2)!} = 6$$

And if we did not care what else we had in our hand for the other three cards we would compute:

$$\binom{48}{3} = \frac{48!}{3!45!} = 17,296$$

Putting this together, we can compute the probability of getting exactly two aces in a 5 card poker hand as:

$$\frac{\binom{4}{2} \binom{48}{3}}{\binom{52}{5}} = .0399$$

This solution is really just the probability distribution known as the Hypergeometric. The generalized formula is:

$$h(x) = \frac{\binom{A}{x} \binom{N-A}{n-x}}{\binom{N}{n}}$$

where x = the number we are interested in coming from the group with A objects.

$h(x)$ is the probability of x successes, in n attempts, when A successes (aces in this case) are in a population that contains N elements. The hypergeometric distribution is an example of a discrete probability distribution because there is no possibility of partial success, that is, there can be no poker hands with 2 1/2 aces. Said another way, a discrete random variable has to be a whole, or counting, number only. This probability distribution works in cases where the probability of a success changes with each draw. Another way of saying this is that the events are NOT independent. In using a deck of cards, we are sampling WITHOUT replacement. If we put each card back after it was drawn then the hypergeometric distribution be an inappropriate Pdf.

For the hypergeometric to work,

1. the population must be dividable into two and only two independent subsets (aces and non-aces in our example). The random variable X = the number of items from the group of interest.
2. the experiment must have changing probabilities of success with each experiment (the fact that cards are not replaced after the draw in our example makes this true in this case). Another way to say this is that you sample without replacement and therefore each pick is not independent.
3. the random variable must be discrete, rather than continuous.

Example 4.2.1

A candy dish contains 30 jelly beans and 20 gumdrops. Ten candies are picked at random. What is the probability that 5 of the 10 are gumdrops? The two groups are jelly beans and gumdrops. Since the probability question asks for the probability of picking gumdrops, the group of interest (first group A in the formula) is gumdrops. The size of the group of interest (first group) is 30. The size of the second group is 20. The size of the sample is 10 (jelly beans or gumdrops). Let X = the number of gumdrops in the sample of 10. X takes on the values $x = 0, 1, 2, \dots, 10$. a. What is the probability statement written mathematically? b. What is the hypergeometric probability density function written out to solve this problem? c. What is the answer to the question "What is the probability of drawing 5 gumdrops in 10 picks from the dish?"

Answer

a. $P(x = 5)$

b. $P(x = 5) = \frac{\binom{30}{5} \binom{20}{5}}{\binom{50}{10}}$

c. $P(x = 5) = 0.215$

Exercise 4.2.1

A bag contains letter tiles. Forty-four of the tiles are vowels, and 56 are consonants. Seven tiles are picked at random. You want to know the probability that four of the seven tiles are vowels. What is the group of interest, the size of the group of interest, and the size of the sample?

This page titled [4.2: Hypergeometric Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [4.1: Hypergeometric Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.3: Binomial Distribution

A more valuable probability density function with many applications is the binomial distribution. This distribution will compute probabilities for any binomial process. A binomial process, often called a Bernoulli process after the first person to fully develop its properties, is any case where there are only two possible outcomes in any one trial, called successes and failures. It gets its name from the binary number system where all numbers are reduced to either 1's or 0's, which is the basis for computer technology and CD music recordings.

Binomial Formula

$$b(x) = \binom{n}{x} p^x q^{n-x}$$

where $b(x)$ is the probability of X successes in n trials when the probability of a success in ANY ONE TRIAL is p . And of course $q = (1 - p)$ and is the probability of a failure in any one trial.

We can see now why the combinatorial formula is also called the binomial coefficient because it reappears here again in the binomial probability function. For the binomial formula to work, the probability of a success in any one trial must be the same from trial to trial, or in other words, the outcomes of each trial must be independent. Flipping a coin is a binomial process because the probability of getting a head in one flip does not depend upon what has happened in PREVIOUS flips. (At this time it should be noted that using p for the parameter of the binomial distribution is a violation of the rule that population parameters are designated with Greek letters. In many textbooks θ (pronounced theta) is used instead of p and this is how it should be.

Just like a set of data, a probability density function has a mean and a standard deviation that describes the data set. For the binomial distribution these are given by the formulas:

$$\begin{aligned}\mu &= np \\ \sigma &= \sqrt{npq}\end{aligned}$$

Notice that p is the only parameter in these equations. The binomial distribution is thus seen as coming from the one-parameter family of probability distributions. In short, we know all there is to know about the binomial once we know p , the probability of a success in any one trial.

In probability theory, under certain circumstances, one probability distribution can be used to approximate another. We say that one is the limiting distribution of the other. If a small number is to be drawn from a large population, even if there is no replacement, we can still use the binomial even though this is not a binomial process. If there is no replacement it violates the independence rule of the binomial. Nevertheless, we can use the binomial to approximate a probability that is really a hypergeometric distribution if we are drawing fewer than 10 percent of the population, i.e. n is less than 10 percent of N in the formula for the hypergeometric function. The rationale for this argument is that when drawing a small percentage of the population we do not alter the probability of a success from draw to draw in any meaningful way. Imagine drawing from not one deck of 52 cards but from 6 decks of cards. The probability of say drawing an ace does not change the conditional probability of what happens on a second draw in the same way it would if there were only 4 aces rather than the 24 aces now to draw from. This ability to use one probability distribution to estimate others will become very valuable to us later.

There are three characteristics of a binomial experiment.

1. There are a fixed number of trials. Think of trials as repetitions of an experiment. The letter n denotes the number of trials.
2. The random variable, x , number of successes, is discrete.
3. There are only two possible outcomes, called "success" and "failure," for each trial. The letter p denotes the probability of a success on any one trial, and q denotes the probability of a failure on any one trial. $p + q = 1$.
4. The n trials are independent and are repeated using identical conditions. Think of this as drawing WITH replacement. Because the n trials are independent, the outcome of one trial does not help in predicting the outcome of another trial. Another way of saying this is that for each individual trial, the probability, p , of a success and probability, q , of a failure remain the same. For example, randomly guessing at a true-false statistics question has only two outcomes. If a success is guessing correctly, then a failure is guessing incorrectly. Suppose Joe always guesses correctly on any statistics true-false question with a probability $p = 0.6$. Then, $q = 0.4$. This means that for every true-false statistics question Joe answers, his probability of success ($p = 0.6$) and his probability of failure ($q = 0.4$) remain the same.

The outcomes of a binomial experiment fit a **binomial probability distribution**. The random variable X = the number of successes obtained in the n independent trials.

The mean, μ , and variance, σ^2 , for the binomial probability distribution are $\mu = np$ and $\sigma^2 = npq$. The standard deviation, σ , is then $\sigma = \sqrt{npq}$.

Any experiment that has characteristics three and four and where $n = 1$ is called a **Bernoulli Trial** (named after Jacob Bernoulli who, in the late 1600s, studied them extensively). A binomial experiment takes place when the number of successes is counted in one or more Bernoulli Trials.

Example 4.3.2

Suppose you play a game that you can only either win or lose. The probability that you win any game is 55%, and the probability that you lose is 45%. Each game you play is independent. If you play the game 20 times, write the function that describes the probability that you win 15 of the 20 times. Here, if you define X as the number of wins, then X takes on the values 0, 1, 2, 3, ..., 20. The probability of a success is $p = 0.55$. The probability of a failure is $q = 0.45$. The number of trials is $n = 20$. The probability question can be stated mathematically as $P(x = 15)$

Exercise 4.3.2

A trainer is teaching a dolphin to do tricks. The probability that the dolphin successfully performs the trick is 35%, and the probability that the dolphin does not successfully perform the trick is 65%. Out of 20 attempts, you want to find the probability that the dolphin succeeds 12 times. Find the $P(X = 12)$ using the binomial Pdf

Example 4.3.3

A fair coin is flipped 15 times. Each flip is independent. What is the probability of getting more than ten heads? Let X = the number of heads in 15 flips of the fair coin. X takes on the values 0, 1, 2, 3, ..., 15. Since the coin is fair, $p = 0.5$ and $q = 0.5$. The number of trials is $n = 15$. State the probability question mathematically.

Answer

$$P(x > 10)$$

Example 4.3.4

Approximately 70% of statistics students do their homework in time for it to be collected and graded. Each student does homework independently. In a statistics class of 50 students, what is the probability that at least 40 will do their homework on time? Students are selected randomly.

a. This is a binomial problem because there is only a success or a _____, there are a fixed number of trials, and the probability of a success is 0.70 for each trial.

Answer

a. failure

b. If we are interested in the number of students who do their homework on time, then how do we define X ?

Answer

b. X = the number of statistics students who do their homework on time

c. What values does x take on?

Answer

c. 0, 1, 2, ..., 50

d. What is a "failure," in words?

Answer

d. Failure is defined as a student who does not complete his or her homework on time.

The probability of a success is $p = 0.70$. The number of trials is $n = 50$.

e. If $p + q = 1$, then what is q ?

Answer

e. $q = 0.30$

f. The words "at least" translate as what kind of inequality for the probability question $P(x \text{ ____ } 40)$.

Answer

f. greater than or equal to (\geq)

The probability question is $P(x \geq 40)$.

Exercise 4.3.4

Sixty-five percent of people pass the state driver's exam on the first try. A group of 50 individuals who have taken the driver's exam is randomly selected. Give two reasons why this is a binomial problem

Exercise 4.3.4

During the 2013 regular NBA season, DeAndre Jordan of the Los Angeles Clippers had the highest field goal completion rate in the league. DeAndre scored with 61.3% of his shots. Suppose you choose a random sample of 80 shots made by DeAndre during the 2013 season. Let X = the number of shots that scored points.

- What is the probability distribution for X ?
- Using the formulas, calculate the (i) mean and (ii) standard deviation of X .
- Find the probability that DeAndre scored with 60 of these shots.
- Find the probability that DeAndre scored with more than 50 of these shots.

This page titled [4.3: Binomial Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **4.2: Binomial Distribution** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.4: Geometric Distribution

The geometric probability density function builds upon what we have learned from the [binomial distribution](#). In this case the experiment continues until either a success or a failure occurs rather than for a set number of trials. There are three main characteristics of a geometric experiment.

1. There are one or more Bernoulli trials with all failures except the last one, which is a success. In other words, you keep repeating what you are doing until the first success. Then you stop. For example, you throw a dart at a bullseye until you hit the bullseye. The first time you hit the bullseye is a "success" so you stop throwing the dart. It might take six tries until you hit the bullseye. You can think of the trials as failure, failure, failure, failure, failure, success, STOP.
2. In theory, the number of trials could go on forever.
3. The probability, p , of a success and the probability, q , of a failure is the same for each trial. $p + q = 1$ and $q = 1 - p$. For example, the probability of rolling a three when you throw one fair die is $\frac{1}{6}$. This is true no matter how many times you roll the die. Suppose you want to know the probability of getting the first three on the fifth roll. On rolls one through four, you do not get a face with a three. The probability for each of the rolls is $q = \frac{5}{6}$, the probability of a failure. The probability of getting a three on the fifth roll is $(\frac{5}{6}) (\frac{5}{6}) (\frac{5}{6}) (\frac{5}{6}) (\frac{1}{6}) = 0.0804$
4. X = the number of independent trials until the first success.

Example 4.4.5

You play a game of chance that you can either win or lose (there are no other possibilities) **until** you lose. Your probability of losing is $p = 0.57$. What is the probability that it takes five games until you lose? Let X = the number of games you play until you lose (includes the losing game). Then X takes on the values 1, 2, 3, ... (could go on indefinitely). The probability question is $P(x = 5)$.

Exercise 4.4.5

You throw darts at a board until you hit the center area. Your probability of hitting the center area is $p = 0.17$. You want to find the probability that it takes eight throws until you hit the center. What values does X take on?

Example 4.4.6

A safety engineer feels that 35% of all industrial accidents in her plant are caused by failure of employees to follow instructions. She decides to look at the accident reports (selected randomly and replaced in the pile after reading) **until** she finds one that shows an accident caused by failure of employees to follow instructions. On average, how many reports would the safety engineer **expect** to look at until she finds a report showing an accident caused by employee failure to follow instructions? What is the probability that the safety engineer will have to examine at least three reports until she finds a report showing an accident caused by employee failure to follow instructions?

Let X = the number of accidents the safety engineer must examine **until** she finds a report showing an accident caused by employee failure to follow instructions. X takes on the values 1, 2, 3, The first question asks you to find the **expected value** or the mean. The second question asks you to find $P(x \geq 3)$. ("At least" translates to a "greater than or equal to" symbol).

Exercise 4.4.6

An instructor feels that 15% of students get below a C on their final exam. She decides to look at final exams (selected randomly and replaced in the pile after reading) until she finds one that shows a grade below a C. We want to know the probability that the instructor will have to examine at least ten exams until she finds one with a grade below a C. What is the probability question stated mathematically?

Example 4.4.7

Suppose that you are looking for a student at your college who lives within five miles of you. You know that 55% of the 25,000 students do live within five miles of you. You randomly contact students from the college **until** one says he or she lives within five miles of you. What is the probability that you need to contact four people?

This is a geometric problem because you may have a number of failures before you have the one success you desire. Also, the probability of a success stays approximately the same each time you ask a student if he or she lives within five miles of you. There is no definite number of trials (number of times you ask a student).

a. Let X = the number of _____ you must ask _____ one says yes.

Answer

a. Let X = the number of **students** you must ask **until** one says yes.

b. What values does X take on?

Answer

b. 1, 2, 3, ..., (total number of students)

c. What are p and q ?

Answer

c. $p = 0.55$; $q = 0.45$

d. The probability question is P (_____).

Answer

d. $P(x = 4)$

Notation for the Geometric: G = Geometric Probability Distribution Function

$$X \sim G(p)$$

Read this as " X is a random variable with a **geometric distribution**." The parameter is p ; p = the probability of a success for each trial.

The Geometric Pdf tells us the probability that the first occurrence of success requires x number of independent trials, each with success probability p . If the probability of success on each trial is p , then the probability that the x th trial (out of x trials) is the first success is:

$$P(X = x) = (1 - p)^{x-1} p$$

for $x = 1, 2, 3, \dots$

The expected value of X , the mean of this distribution, is $1/p$. This tells us how many trials we have to expect until we get the first success including in the count the trial that results in success. The above form of the Geometric distribution is used for modeling the number of trials until the first success. The number of trials includes the one that is a success: x = all trials including the one that is a success. This can be seen in the form of the formula. If X = number of trials including the success, then we must multiply the probability of failure, $(1 - p)$, times the number of failures, that is $X - 1$.

By contrast, the following form of the geometric distribution is used for modeling number of failures until the first success:

$$P(X = x) = (1 - p)^x p$$

for $x = 0, 1, 2, 3, \dots$

In this case the trial that is a success is not counted as a trial in the formula: x = number of failures. The expected value, mean, of

this distribution is $\mu = \frac{(1-p)}{p}$. This tells us how many failures to expect before we have a success. In either case, the sequence of probabilities is a geometric sequence.

Example 4.4.8

Assume that the probability of a defective computer component is 0.02. Components are randomly selected. Find the probability that the first defect is caused by the seventh component tested. How many components do you expect to test until one is found to be defective?

Let X = the number of computer components tested until the first defect is found.

X takes on the values 1, 2, 3, ... where $p = 0.02$. $X \sim G(0.02)$

Find $P(x = 7)$. Answer: $P(x = 7) = (1 - 0.02)^{7-1} \times 0.02 = 0.0177$.

The probability that the seventh component is the first defect is 0.0177.

The graph of $X \sim G(0.02)$ is:




Figure 4.4.2

The y -axis contains the probability of x , where X = the number of computer components tested. Notice that the probabilities decline by a common increment. This increment is the same ratio between each number and is called a geometric progression and thus the name for this probability density function.

The number of components that you would expect to test until you find the first defective component is the mean, $\mu = 50$.

The formula for the mean for the random variable defined as number of failures until first success is $\mu = \frac{1}{p} = \frac{1}{0.02} = 50$

See [Example 4.4.9](#) for an example where the geometric random variable is defined as number of trials until first success. The expected value of this formula for the geometric will be different from this version of the distribution.

The formula for the variance is $\sigma^2 = \left(\frac{1}{p}\right) \left(\frac{1}{p} - 1\right) = \left(\frac{1}{0.02}\right) \left(\frac{1}{0.02} - 1\right) = 2,450$

The standard deviation is $\sigma = \sqrt{\left(\frac{1}{p}\right) \left(\frac{1}{p} - 1\right)} = \sqrt{\left(\frac{1}{0.02}\right) \left(\frac{1}{0.02} - 1\right)} = 49.5$

The lifetime risk of developing pancreatic cancer is about one in 78 (1.28%). Let X = the number of people you ask before one says he or she has pancreatic cancer. The random variable X in this case includes only the number of trials that were failures and does not count the trial that was a success in finding a person who had the disease. The appropriate formula for this random variable is the second one presented above. Then X is a discrete random variable with a geometric distribution: $X \sim G\left(\frac{1}{78}\right)$ or $X \sim G(0.0128)$.

- What is the probability of that you ask 9 people before one says he or she has pancreatic cancer? This is asking, what is the probability that you ask 9 people unsuccessfully and the tenth person is a success?
- What is the probability that you must ask 20 people?
- Find the (i) mean and (ii) standard deviation of X .

Answer

a. $P(x = 9) = (1 - 0.0128)^9 \cdot 0.0128 = 0.0114$

$$b. P(x = 20) = (1 - 0.0128)^{19} \cdot 0.0128 = 0.01$$

$$i. \text{ Mean} = \mu = \frac{(1-p)}{p} = \frac{(1-0.0128)}{0.0128} = 77.12$$

$$ii. \text{ Standard Deviation} = \sigma = \sqrt{\frac{1-p}{p^2}} = \sqrt{\frac{1-0.0128}{0.0128^2}} \approx 77.62$$

Exercise 4.4.9

The literacy rate for a nation measures the proportion of people age 15 and over who can read and write. The literacy rate for women in The United Colonies of Independence is 12%. Let X = the number of women you ask until one says that she is literate.

- What is the probability distribution of X ?
- What is the probability that you ask five women before one says she is literate?
- What is the probability that you must ask ten women?

Example 4.4.10

A baseball player has a batting average of 0.320. This is the general probability that he gets a hit each time he is at bat. What is the probability that he gets his first hit in the third trip to bat?

Answer

$$P(x = 3) = (1 - 0.32)^{3-1} \times .32 = 0.1480$$

In this case the sequence is failure, failure success.

How many trips to bat do you expect the hitter to need before getting a hit?

Answer

$$\mu = \frac{1}{p} = \frac{1}{0.320} = 3.125 \approx 3$$

This is simply the expected value of successes and therefore the mean of the distribution.

Example 4.4.11

There is an 80% chance that a Dalmatian dog has 13 black spots. You go to a dog show and count the spots on Dalmatians. What is the probability that you will review the spots on 3 dogs before you find one that has 13 black spots?

Answer

$$P(x = 3) = (1 - 0.80)^3 \times 0.80 = 0.0064$$

Footnotes

1 "Prevalence of HIV, total (% of populations ages 15-49)," The World Bank, 2013. Available online at <http://data.worldbank.org/indicator/...last&sort=desc> (accessed May 15, 2013).

This page titled 4.4: Geometric Distribution is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by OpenStax via source content that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- 4.3: Geometric Distribution by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.5: Poisson Distribution

Another useful probability distribution is the Poisson distribution, or waiting time distribution. This distribution is used to determine how many checkout clerks are needed to keep the waiting time in line to specified levels, how many telephone lines are needed to keep the system from overloading, and many other practical applications. A modification of the Poisson, the Pascal, invented nearly four centuries ago, is used today by telecommunications companies worldwide for load factors, satellite hookup levels and Internet capacity problems. The distribution gets its name from Simeon Poisson who presented it in 1837 as an extension of the binomial distribution which we will see can be estimated with the Poisson.

There are two main characteristics of a Poisson experiment.

1. The **Poisson probability distribution** gives the probability of a number of events occurring in a **fixed interval** of time or space if these events happen with a known average rate.
2. The events are independently of the time since the last event. For example, a book editor might be interested in the number of words spelled incorrectly in a particular book. It might be that, on the average, there are five words spelled incorrectly in 100 pages. The interval is the 100 pages and it is assumed that there is no relationship between when misspellings occur.
3. The random variable X = the number of occurrences in the interval of interest.

Example 4.5.12

A bank expects to receive six bad checks per day, on average. What is the probability of the bank getting fewer than five bad checks on any given day? Of interest is the number of checks the bank receives in one day, so the time interval of interest is one day. Let X = the number of bad checks the bank receives in one day. If the bank expects to receive six bad checks per day then the average is six checks per day. Write a mathematical statement for the probability question.

Answer

$$P(x < 5)$$

Example 4.5.13

You notice that a news reporter says "uh," on average, two times per broadcast. What is the probability that the news reporter says "uh" more than two times per broadcast.

This is a Poisson problem because you are interested in knowing the number of times the news reporter says "uh" during a broadcast.

- a. What is the interval of interest?

Answer

- a. one broadcast measured in minutes

- b. What is the average number of times the news reporter says "uh" during one broadcast?

Answer

- b. 2

- c. Let X = _____. What values does X take on?

Answer

- c. Let X = the number of times the news reporter says "uh" during one broadcast.
 $x = 0, 1, 2, 3, \dots$

- d. The probability question is P (_____).

Answer

d. $P(x > 2)$

Notation for the Poisson: P = Poisson Probability Distribution Function

$$X \sim P(\mu)$$

Read this as "X is a random variable with a Poisson distribution." The parameter is μ (or λ); μ (or λ) = the mean for the interval of interest. The mean is the number of occurrences that occur on average during the interval period.

The formula for computing probabilities that are from a Poisson process is:

$$P(x) = \frac{\mu^x e^{-\mu}}{x!}$$

where $P(X)$ is the probability of X successes, μ is the expected number of successes based upon historical data, e is the natural logarithm approximately equal to 2.718, and X is the number of successes per unit, usually per unit of time.

In order to use the Poisson distribution, certain assumptions must hold. These are: the probability of a success, μ , is unchanged within the interval, there cannot be simultaneous successes within the interval, and finally, that the probability of a success among intervals is independent, the same assumption of the binomial distribution.

In a way, the Poisson distribution can be thought of as a clever way to convert a continuous random variable, usually time, into a discrete random variable by breaking up time into discrete independent intervals. This way of thinking about the Poisson helps us understand why it can be used to estimate the probability for the discrete random variable from the binomial distribution. The Poisson is asking for the probability of a number of successes during a period of time while the binomial is asking for the probability of a certain number of successes for a given number of trials.

Example 4.5.14

Leah's answering machine receives about six telephone calls between 8 a.m. and 10 a.m. What is the probability that Leah receives more than one call **in the next 15 minutes?**

Let X = the number of calls Leah receives in 15 minutes. (The **interval of interest** is 15 minutes or $\frac{1}{4}$ hour.)

$$x = 0, 1, 2, 3, \dots$$

If Leah receives, on the average, six telephone calls in two hours, and there are eight 15 minute intervals in two hours, then Leah receives

$(\frac{1}{8})(6) = 0.75$ calls in 15 minutes, on average. So, $\mu = 0.75$ for this problem.

$$X \sim P(0.75)$$

Find $P(x > 1)$. $P(x > 1) = 0.1734$

Probability that Leah receives more than one telephone call in the next 15 minutes is about 0.1734.

The graph of $X \sim P(0.75)$ is:




Figure 4.5.3

The y-axis contains the probability of x where X = the number of calls in 15 minutes.

Example 4.5.15

According to a survey a university professor gets, on average, 7 emails per day. Let X = the number of emails a professor receives per day. The discrete random variable X takes on the values $x = 0, 1, 2, \dots$. The random variable X has a Poisson distribution: $X \sim P(7)$. The mean is 7 emails.

- What is the probability that an email user receives exactly 2 emails per day?
- What is the probability that an email user receives at most 2 emails per day?
- What is the standard deviation?

Answer

- $P(x = 2) = \frac{\mu^x e^{-\mu}}{x!} = \frac{7^2 e^{-7}}{2!} = 0.022$
- $P(x \leq 2) = \frac{7^0 e^{-7}}{0!} + \frac{7^1 e^{-7}}{1!} + \frac{7^2 e^{-7}}{2!} = 0.029$
- Standard Deviation = $\sigma = \sqrt{\mu} = \sqrt{7} \approx 2.65$

Example 4.5.16

Text message users receive or send an average of 41.5 text messages per day.

- How many text messages does a text message user receive or send per hour?
- What is the probability that a text message user receives or sends two messages per hour?
- What is the probability that a text message user receives or sends more than two messages per hour?

Answer

- Let X = the number of texts that a user sends or receives in one hour. The average number of texts received per hour is $\frac{41.5}{24} \approx 1.7292$.
- $P(x = 2) = \frac{\mu^x e^{-\mu}}{x!} = \frac{1.7292^2 e^{-1.7292}}{2!} = 0.265$
- $P(x > 2) = 1 - P(x \leq 2) = 1 - \left[\frac{7^0 e^{-7}}{0!} + \frac{7^1 e^{-7}}{1!} + \frac{7^2 e^{-7}}{2!} \right] = 0.250$

Example 4.5.17

On May 13, 2013, starting at 4:30 PM, the probability of low seismic activity for the next 48 hours in Alaska was reported as about 1.02%. Use this information for the next 200 days to find the probability that there will be low seismic activity in ten of the next 200 days. Use both the binomial and Poisson distributions to calculate the probabilities. Are they close?

Answer

Let X = the number of days with low seismic activity.

Using the binomial distribution:

$$P(x = 10) = \frac{200!}{10!(200 - 10)!} \times .0102^{10} \times .9898^{190} = 0.000039$$

Using the Poisson distribution:

Calculate $\mu = np = 200(0.0102) \approx 2.04$

$$P(x = 10) = \frac{\mu^x e^{-\mu}}{x!} = \frac{2.04^{10} e^{-2.04}}{10!} = 0.000045$$

We expect the approximation to be good because n is large (greater than 20) and p is small (less than 0.05). The results are close—both probabilities reported are almost 0.

Estimating the Binomial Distribution with the Poisson Distribution

We found before that the binomial distribution provided an approximation for the hypergeometric distribution. Now we find that the Poisson distribution can provide an approximation for the binomial. We say that the binomial distribution approaches the Poisson. The binomial distribution approaches the Poisson distribution as n gets larger and p is small such that np becomes a constant value. There are several rules of thumb for when one can say they will use a Poisson to estimate a binomial. One suggests that np , the mean of the binomial, should be less than 25. Another author suggests that it should be less than 7. And another, noting that the mean and variance of the Poisson are both the same, suggests that np and npq , the mean and variance of the binomial, should be greater than 5. There is no one broadly accepted rule of thumb for when one can use the Poisson to estimate the binomial.

As we move through these probability distributions we are getting to more sophisticated distributions that, in a sense, contain the less sophisticated distributions within them. This proposition has been proven by mathematicians. This gets us to the highest level of sophistication in the next probability distribution which can be used as an approximation to all of those that we have discussed so far. This is the normal distribution.

Example 4.5.18

A survey of 500 seniors in the Price Business School yields the following information. 75% go straight to work after graduation. 15% go on to work on their MBA. 9% stay to get a minor in another program. 1% go on to get a Master's in Finance.

What is the probability that more than 2 seniors go to graduate school for their Master's in finance?

Answer

This is clearly a binomial probability distribution problem. The choices are binary when we define the results as "Graduate School in Finance" versus "all other options." The random variable is discrete, and the events are, we could assume, independent. Solving as a binomial problem, we have:

Binomial Solution

$$n \cdot p = 500 \cdot 0.01 = 5 = \mu$$

$$P(0) = \frac{500!}{0!(500-0)!} 0.01^0 (1-0.01)^{500-0} = 0.00657$$

$$P(1) = \frac{500!}{1!(500-1)!} 0.01^1 (1-0.01)^{500} = 0.03318$$

$$P(2) = \frac{500!}{2!(500-2)!} 0.01^2 (1-0.01)^{500} = 0.08363$$

Adding all 3 together = 0.12339

$$1 - 0.12339 = 0.87661$$

Poisson approximation

$$n \cdot p = 500 \cdot 0.01 = 5 = \mu$$

$$n \cdot p \cdot (1-p) = 500 \cdot 0.01 \cdot (0.99) \approx 5 = \sigma^2 = \mu$$

$$P(X) = \frac{e^{-np} (np)^x}{x!} = \left\{ P(0) = \frac{e^{-5} \cdot 5^0}{0!} \right\} + \left\{ P(1) = \frac{e^{-5} \cdot 5^1}{1!} \right\} + \left\{ P(2) = \frac{e^{-5} \cdot 5^2}{2!} \right\}$$

$$0.0067 + 0.0337 + 0.0842 = 0.1247$$

$$1 - 0.1247 = 0.8753$$

An approximation that is off by 1 one thousandth is certainly an acceptable approximation.

This page titled [4.5: Poisson Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **4.4: Poisson Distribution** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.6: Chapter Key Items

Bernoulli Trials

an experiment with the following characteristics:

1. There are only two possible outcomes called “success” and “failure” for each trial.
2. The probability p of a success is the same for any trial (so the probability $q = 1 - p$ of a failure is the same for any trial).

Binomial Experiment

a statistical experiment that satisfies the following three conditions:

1. There are a fixed number of trials, n .
2. There are only two possible outcomes, called "success" and, "failure," for each trial. The letter p denotes the probability of a success on one trial, and q denotes the probability of a failure on one trial.
3. The n trials are independent and are repeated using identical conditions.

Binomial Probability Distribution

a discrete random variable (RV) that arises from Bernoulli trials; there are a fixed number, n , of independent trials.

“Independent” means that the result of any trial (for example, trial one) does not affect the results of the following trials, and all trials are conducted under the same conditions. Under these circumstances the binomial RV X is defined as the number of successes in n trials. The mean is $\mu = np$ and the standard deviation is $\sigma = \sqrt{npq}$. The probability of exactly x successes in n

trials is $P(X = x) = \binom{n}{x} p^x q^{n-x}$.

Geometric Distribution

a discrete random variable (RV) that arises from the Bernoulli trials; the trials are repeated until the first success. The geometric variable X is defined as the number of trials until the first success. The mean is $\mu = \frac{1}{p}$ and the standard deviation is

$\sigma = \sqrt{\frac{1}{p} \left(\frac{1}{p} - 1 \right)}$. The probability of exactly x failures before the first success is given by the formula:

$P(X = x) = p(1 - p)^{x-1}$ where one wants to know probability for the number of trials until the first success: the x th trial is the first success.

An alternative formulation of the geometric distribution asks the question: what is the probability of x failures until the first success? In this formulation the trial that resulted in the first success is not counted. The formula for this presentation of the geometric is: $P(X = x) = p(1 - p)^x$

The expected value in this form of the geometric distribution is $\mu = \frac{1-p}{p}$

The easiest way to keep these two forms of the geometric distribution straight is to remember that p is the probability of success and $(1 - p)$ is the probability of failure. In the formula the exponents simply count the number of successes and number of failures of the desired outcome of the experiment. Of course the sum of these two numbers must add to the number of trials in the experiment.

Geometric Experiment

a statistical experiment with the following properties:

1. There are one or more Bernoulli trials with all failures except the last one, which is a success.
2. In theory, the number of trials could go on forever. There must be at least one trial.
3. The probability, p , of a success and the probability, q , of a failure do not change from trial to trial.

Hypergeometric Experiment

a statistical experiment with the following properties:

1. You take samples from two groups.
2. You are concerned with a group of interest, called the first group.
3. You sample without replacement from the combined groups.
4. Each pick is not independent, since sampling is without replacement.

Hypergeometric Probability

a discrete random variable (RV) that is characterized by:

1. A fixed number of trials.
2. The probability of success is not the same from trial to trial.

We sample from two groups of items when we are interested in only one group. X is defined as the number of successes out of the total number of items chosen.

Poisson Probability Distribution

a discrete random variable (RV) that counts the number of times a certain event will occur in a specific interval; characteristics of the variable:

- The probability that the event occurs in a given interval is the same for all intervals.
- The events occur with a known mean and independently of the time since the last event.

The distribution is defined by the mean μ of the event in the interval. The mean is $\mu = np$. The standard deviation is $\sigma = \sqrt{\mu}$.

The probability of having exactly x successes in r trials is $P(x) = \frac{\mu^x e^{-\mu}}{x!}$. The Poisson distribution is often used to approximate the binomial distribution, when n is “large” and p is “small” (a general rule is that np should be greater than or equal to 25 and p should be less than or equal to 0.01).

Probability Distribution Function (PDF)

a mathematical description of a discrete random variable (RV), given either in the form of an equation (formula) or in the form of a table listing all the possible outcomes of an experiment and the probability associated with each outcome.

Random Variable (RV)

a characteristic of interest in a population being studied; common notation for variables are upper case Latin letters X, Y, Z, \dots ; common notation for a specific value from the domain (set of all possible values of a variable) are lower case Latin letters x, y , and z . For example, if X is the number of children in a family, then x represents a specific integer 0, 1, 2, 3,.... Variables in statistics differ from variables in intermediate algebra in the two following ways.

- The domain of the random variable (RV) is not necessarily a numerical set; the domain may be expressed in words; for example, if $X =$ hair color then the domain is {black, blond, gray, green, orange}.
- We can tell what specific value x the random variable X takes only after performing the experiment.

This page titled [4.6: Chapter Key Items](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [4.7: Chapter Key Items](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.7: Chapter Review

Introduction

The characteristics of a probability distribution or density function (PDF) are as follows:

1. Each probability is between zero and one, inclusive (*inclusive* means to include zero and one).
2. The sum of the probabilities is one.

4.1 Hypergeometric Distribution

The combinatorial formula can provide the number of unique subsets of size x that can be created from n unique objects to help us calculate probabilities. The combinatorial formula is $\binom{n}{x} = {}_n C_x = \frac{n!}{x!(n-x)!}$

A **hypergeometric experiment** is a statistical experiment with the following properties:

1. You take samples from two groups.
2. You are concerned with a group of interest, called the first group.
3. You sample without replacement from the combined groups.
4. Each pick is not independent, since sampling is without replacement.

The outcomes of a hypergeometric experiment fit a hypergeometric probability distribution. The random variable X = the number

of items from the group of interest. $h(x) = \frac{\binom{A}{x} \binom{N-A}{n-x}}{\binom{N}{n}}$.

Binomial Distribution

A statistical experiment can be classified as a binomial experiment if the following conditions are met:

1. There are a fixed number of trials, n .
2. There are only two possible outcomes, called "success" and "failure" for each trial. The letter p denotes the probability of a success on one trial and q denotes the probability of a failure on one trial.
3. The n trials are independent and are repeated using identical conditions.

The outcomes of a binomial experiment fit a binomial probability distribution. The random variable X = the number of successes obtained in the n independent trials. The mean of X can be calculated using the formula $\mu = np$, and the standard deviation is given by the formula $\sigma = \sqrt{npq}$.

The formula for the Binomial probability density function is

$$P(x) = \frac{n!}{x!(n-x)!} \cdot p^x q^{(n-x)}$$

Geometric Distribution

There are three characteristics of a geometric experiment:

1. There are one or more Bernoulli trials with all failures except the last one, which is a success.
2. In theory, the number of trials could go on forever. There must be at least one trial.
3. The probability, p , of a success and the probability, q , of a failure are the same for each trial.

In a geometric experiment, define the discrete random variable X as the number of independent trials until the first success. We say that X has a geometric distribution and write $X \sim G(p)$ where p is the probability of success in a single trial.

The mean of the geometric distribution $X \sim G(p)$ is $\mu = 1/p$ where x = number of trials until first success for the formula $P(X = x) = (1-p)^{x-1}p$ where the number of trials is up and including the first success.

An alternative formulation of the geometric distribution asks the question: what is the probability of x failures until the first success? In this formulation the trial that resulted in the first success is not counted. The formula for this presentation of the geometric is:

$$P(X = x) = p(1 - p)^x$$

The expected value in this form of the geometric distribution is

$$\mu = \frac{1 - p}{p}$$

The easiest way to keep these two forms of the geometric distribution straight is to remember that p is the probability of success and $(1 - p)$ is the probability of failure. In the formula the exponents simply count the number of successes and number of failures of the desired outcome of the experiment. Of course the sum of these two numbers must add to the number of trials in the experiment.

Poisson Distribution

A **Poisson probability distribution** of a discrete random variable gives the probability of a number of events occurring in a fixed interval of time or space, if these events happen at a known average rate and independently of the time since the last event. The Poisson distribution may be used to approximate the binomial, if the probability of success is "small" (less than or equal to 0.01) and the number of trials is "large" (greater than or equal to 25). Other rules of thumb are also suggested by different authors, but all recognize that the Poisson distribution is the limiting distribution of the binomial as n increases and p approaches zero.

The formula for computing probabilities that are from a Poisson process is:

$$P(x) = \frac{\mu^x e^{-\mu}}{x!}$$

where $P(X)$ is the probability of successes, μ (pronounced mu) is the expected number of successes, e is the natural logarithm approximately equal to 2.718, and X is the number of successes per unit, usually per unit of time.

This page titled [4.7: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **4.10: Chapter Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.8: Chapter Formula Review

Hypergeometric Distribution

$$h(x) = \frac{\binom{A}{x} \binom{N-A}{n-x}}{\binom{N}{n}}$$

Binomial Distribution

$X \sim B(n, p)$ means that the discrete random variable X has a binomial probability distribution with n trials and probability of success p .

X = the number of successes in n independent trials

n = the number of independent trials

X takes on the values $x = 0, 1, 2, 3, \dots, n$

p = the probability of a success for any trial

q = the probability of a failure for any trial

$$p + q = 1$$

$$q = 1 - p$$

The mean of X is $\mu = np$. The standard deviation of X is $\sigma = \sqrt{npq}$.

$$P(x) = \frac{n!}{x!(n-x)!} \cdot p^x q^{(n-x)}$$

where $P(X)$ is the probability of X successes in n trials when the probability of a success in ANY ONE TRIAL is p .

Geometric Distribution

$$P(X = x) = p(1 - p)^{x-1}$$

$X \sim G(p)$ means that the discrete random variable X has a geometric probability distribution with probability of success in a single trial p .

X = the number of independent trials until the first success

X takes on the values $x = 1, 2, 3, \dots$

p = the probability of a success for any trial

q = the probability of a failure for any trial $p + q = 1$

$$q = 1 - p$$

The mean is $\mu = \frac{1}{p}$.

The standard deviation is $\sigma = \sqrt{\frac{1-p}{p^2}} = \sqrt{\frac{1}{p} \left(\frac{1}{p} - 1 \right)}$.

Poisson Distribution

$X \sim P(\mu)$ means that X has a Poisson probability distribution where X = the number of occurrences in the interval of interest.

X takes on the values $x = 0, 1, 2, 3, \dots$

The mean μ or λ is typically given.

The variance is $\sigma^2 = \mu$, and the standard deviation is

$$\sigma = \sqrt{\mu}.$$

When $P(\mu)$ is used to approximate a binomial distribution, $\mu = np$ where n represents the number of independent trials and p represents the probability of success in a single trial.

$$P(x) = \frac{\mu^x e^{-\mu}}{x!}$$

This page titled [4.8: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [4.5: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.9: Chapter Practice

Introduction

Use the following information to answer the next five exercises: A company wants to evaluate its attrition rate, in other words, how long new hires stay with the company. Over the years, they have established the following probability distribution.

Let X = the number of years a new hire will stay with the company.

Let $P(x)$ = the probability that a new hire will stay with the company x years.

1.

Complete Table 4.9.1 using the data provided.

Table 4.9.1

x	$P(x)$
0	0.12
1	0.18
2	0.30
3	0.15
4	
5	0.10
6	0.05

2.

$$P(x = 4) = \underline{\hspace{2cm}}$$

3.

$$P(x \geq 5) = \underline{\hspace{2cm}}$$

4.

On average, how long would you expect a new hire to stay with the company?

5.

What does the column " $P(x)$ " sum to?

Use the following information to answer the next six exercises: A baker is deciding how many batches of muffins to make to sell in his bakery. He wants to make enough to sell every one and no fewer. Through observation, the baker has established a probability distribution.

Table 4.9.2

x	$P(x)$
1	0.15
2	0.35
3	0.40
4	0.10

6.

Define the random variable X .

7.

What is the probability the baker will sell more than one batch? $P(x > 1) = \underline{\hspace{2cm}}$

8.

What is the probability the baker will sell exactly one batch? $P(x = 1) = \underline{\hspace{2cm}}$

9.

On average, how many batches should the baker make?

Use the following information to answer the next four exercises: Ellen has music practice three days a week. She practices for all of the three days 85% of the time, two days 8% of the time, one day 4% of the time, and no days 3% of the time. One week is selected at random.

10.

Define the random variable X .

11.

Construct a probability distribution table for the data.

12.

We know that for a probability distribution function to be discrete, it must have two characteristics. One is that the sum of the probabilities is one. What is the other characteristic?

Use the following information to answer the next five exercises: Javier volunteers in community events each month. He does not do more than five events in a month. He attends exactly five events 35% of the time, four events 25% of the time, three events 20% of the time, two events 10% of the time, one event 5% of the time, and no events 5% of the time.

13.

Define the random variable X .

14.

What values does x take on?

15.

Construct a PDF table.

16.

Find the probability that Javier volunteers for less than three events each month. $P(x < 3) = \underline{\hspace{2cm}}$

17.

Find the probability that Javier volunteers for at least one event each month. $P(x > 0) = \underline{\hspace{2cm}}$

4.1 Hypergeometric Distribution

Use the following information to answer the next five exercises: Suppose that a group of statistics students is divided into two groups: business majors and non-business majors. There are 16 business majors in the group and seven non-business majors in the group. A random sample of nine students is taken. We are interested in the number of business majors in the sample.

18.

In words, define the random variable X .

19.

What values does X take on?

4.2 Binomial Distribution

Use the following information to answer the next eight exercises: The Higher Education Research Institute at UCLA collected data from 203,967 incoming first-time, full-time freshmen from 270 four-year colleges and universities in the U.S. 71.3% of those students replied that, yes, they believe that same-sex couples should have the right to legal marital status. Suppose that you randomly pick eight first-time, full-time freshmen from the survey. You are interested in the number that believes that same-sex couples should have the right to legal marital status.

20.

In words, define the random variable X .

21.

$$X \sim \text{____}(\text{____}, \text{____})$$

22.

What values does the random variable X take on?

23.

Construct the probability distribution function (PDF).

Table 4.9.3

x	$P(x)$

24.

On average (μ), how many would you expect to answer yes?

25.

What is the standard deviation (σ)?

26.

What is the probability that at most five of the freshmen reply “yes”?

27.

What is the probability that at least two of the freshmen reply “yes”?

4.3 Geometric Distribution

Use the following information to answer the next six exercises: The Higher Education Research Institute at UCLA collected data from 203,967 incoming first-time, full-time freshmen from 270 four-year colleges and universities in the U.S. 71.3% of those students replied that, yes, they believe that same-sex couples should have the right to legal marital status. Suppose that you randomly select freshman from the study until you find one who replies “yes.” You are interested in the number of freshmen you must ask.

28.

In words, define the random variable X .

29.

$$X \sim \text{____}(\text{____}, \text{____})$$

30.

What values does the random variable X take on?

31.

Construct the probability distribution function (PDF). Stop at $x = 6$.

Table 4.9.4

x	$P(x)$
1	
2	
3	
4	
5	

x	$P(x)$
6	

32.

On average (μ), how many freshmen would you expect to have to ask until you found one who replies "yes?"

33.

What is the probability that you will need to ask fewer than three freshmen?

4.4 Poisson Distribution

Use the following information to answer the next six exercises: On average, a clothing store gets 120 customers per day.

34.

Assume the event occurs independently in any given day. Define the random variable X .

35.

What values does X take on?

36.

What is the probability of getting 150 customers in one day?

37.

What is the probability of getting 35 customers in the first four hours? Assume the store is open 12 hours each day.

38.

What is the probability that the store will have more than 12 customers in the first hour?

39.

What is the probability that the store will have fewer than 12 customers in the first two hours?

40.

Which type of distribution can the Poisson model be used to approximate? When would you do this?

Use the following information to answer the next six exercises: On average, eight teens in the U.S. die from motor vehicle injuries per day. As a result, states across the country are debating raising the driving age.

41.

Assume the event occurs independently in any given day. In words, define the random variable X .

42.

$X \sim \text{_____}(\text{_____, } \text{_____})$

43.

What values does X take on?

44.

For the given values of the random variable X , fill in the corresponding probabilities.

45.

Is it likely that there will be no teens killed from motor vehicle injuries on any given day in the U.S.? Justify your answer numerically.

46.

Is it likely that there will be more than 20 teens killed from motor vehicle injuries on any given day in the U.S.? Justify your answer numerically

This page titled [4.9: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **4.8: Chapter Practice** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.10: Chapter Homework

4.1 Hypergeometric Distribution

47.

A group of Martial Arts students is planning on participating in an upcoming demonstration. Six are students of Tae Kwon Do; seven are students of Shotokan Karate. Suppose that eight students are randomly picked to be in the first demonstration. We are interested in the number of Shotokan Karate students in that first demonstration.

- Suppose that 1,000 babies from healthy baby nurseries were randomly surveyed. Find the probability that exactly two babies were born deaf.

Use the following information to answer the next four exercises. Recently, a nurse commented that when a patient calls the medical advice line claiming to have the flu, the chance that he or she truly has the flu (and not just a nasty cold) is only about 4%. Of the next 25 patients calling in claiming to have the flu, we are interested in how many actually have the flu.

53.

Define the random variable and list its possible values.

54.

State the distribution of X .

55.

Find the probability that at least four of the 25 patients actually have the flu.

56.

On average, for every 25 patients calling in, how many do you expect to have the flu?

57.

People visiting video rental stores often rent more than one DVD at a time. The probability distribution for DVD rentals per customer at Video To Go is given Table 4.10.5. There is five-video limit per customer at this store, so nobody ever rents more than five DVDs.

Table 4.10.5

x	$P(x)$
0	0.03
1	0.50
2	0.24
3	
4	0.07
5	0.04

- Use the following information to answer the next two exercises: The probability that the San Jose Sharks will win any given game is 0.3694 based on a 13-year win history of 382 wins out of 1,034 games played (as of a certain date). An upcoming monthly schedule contains 12 games.**59.**

The expected number of wins for that upcoming month is:

- Let X = the number of games won in that upcoming month.**60.**

What is the probability that the San Jose Sharks win six games in that upcoming month?

- Use the following information to answer the next two exercises: The average number of times per week that Mrs. Plum's cats wake her up at night because they want to play is ten. We are interested in the number of times her cats wake her up each week.**93.**

In words, the random variable $X =$ _____

1. the number of times Mrs. Plum's cats wake her up each week.
2. the number of times Mrs. Plum's cats wake her up each hour.
3. the number of times Mrs. Plum's cats wake her up each night.
4. the number of times Mrs. Plum's cats wake her up.

94.

Find the probability that her cats will wake her up no more than five times next week.

1. 0.5000
2. 0.9329
3. 0.0378
4. 0.0671

This page titled [4.10: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [4.6: Chapter Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.11: Chapter References

Binomial Distribution

- “Access to electricity (% of population),” The World Bank, 2013. Available online at <http://data.worldbank.org/indicator/...first&sort=asc> (accessed May 15, 2015).
- “Distance Education.” Wikipedia. Available online at http://en.Wikipedia.org/wiki/Distance_education (accessed May 15, 2013).
- “NBA Statistics – 2013,” ESPN NBA, 2013. Available online at http://espn.go.com/nba/statistics/_/seasontype/2 (accessed May 15, 2013).
- Newport, Frank. “Americans Still Enjoy Saving Rather than Spending: Few demographic differences seen in these views other than by income,” GALLUP® Economy, 2013. Available online at <http://www.gallup.com/poll/162368/am...-spending.aspx> (accessed May 15, 2013).
- Pryor, John H., Linda DeAngelo, Laura Palucki Blake, Sylvia Hurtado, Serge Tran. *The American Freshman: National Norms Fall 2011*. Los Angeles: Cooperative Institutional Research Program at the Higher Education Research Institute at UCLA, 2011. Also available online at <http://heri.ucla.edu/PDFs/pubs/TFS/N...eshman2011.pdf> (accessed May 15, 2013).
- “The World FactBook,” Central Intelligence Agency. Available online at <https://www.cia.gov/library/publicat...k/geos/af.html> (accessed May 15, 2013).
- “What are the key statistics about pancreatic cancer?” American Cancer Society, 2013. Available online at <http://www.cancer.org/cancer/pancrea...key-statistics> (accessed May 15, 2013).

Geometric Distribution

- “Millennials: A Portrait of Generation Next,” PewResearchCenter. Available online at <http://www.pewsocialtrends.org/files...-to-change.pdf> (accessed May 15, 2013).
- “Millennials: Confident. Connected. Open to Change.” Executive Summary by PewResearch Social & Demographic Trends, 2013. Available online at <http://www.pewsocialtrends.org/2010/...pen-to-change/> (accessed May 15, 2013).
- “Prevalence of HIV, total (% of populations ages 15-49),” The World Bank, 2013. Available online at <http://data.worldbank.org/indicator/...last&sort=desc> (accessed May 15, 2013).
- Pryor, John H., Linda DeAngelo, Laura Palucki Blake, Sylvia Hurtado, Serge Tran. *The American Freshman: National Norms Fall 2011*. Los Angeles: Cooperative Institutional Research Program at the Higher Education Research Institute at UCLA, 2011. Also available online at <http://heri.ucla.edu/PDFs/pubs/TFS/N...eshman2011.pdf> (accessed May 15, 2013).
- “Summary of the National Risk and Vulnerability Assessment 2007/8: A profile of Afghanistan,” The European Union and ICON-Institute. Available online at ec.europa.eu/europeaid/where/...summary_en.pdf (accessed May 15, 2013).
- “The World FactBook,” Central Intelligence Agency. Available online at <https://www.cia.gov/library/publicat...k/geos/af.html> (accessed May 15, 2013).
- “UNICEF reports on Female Literacy Centers in Afghanistan established to teach women and girls basic resading [sic] and writing skills,” UNICEF Television. Video available online at <http://www.unicefusa.org/assets/vide...y-centers.html> (accessed May 15, 2013).

Poisson Distribution

- “ATL Fact Sheet,” Department of Aviation at the Hartsfield-Jackson Atlanta International Airport, 2013. Available online at www.atl.com/about-atl/atl-factsheet/ (accessed February 6, 2019).
- Center for Disease Control and Prevention. “Teen Drivers: Fact Sheet,” Injury Prevention & Control: Motor Vehicle Safety, October 2, 2012. Available online at <http://www.cdc.gov/Motorvehiclesafet...factsheet.html> (accessed May 15, 2013).
- “Children and Childrearing,” Ministry of Health, Labour, and Welfare. Available online at <http://www.mhlw.go.jp/english/policy...ing/index.html> (accessed May 15, 2013).
- “Eating Disorder Statistics,” South Carolina Department of Mental Health, 2006. Available online at <http://www.state.sc.us/dmh/anorexia/statistics.htm> (accessed May 15, 2013).
- “Giving Birth in Manila: The maternity ward at the Dr Jose Fabella Memorial Hospital in Manila, the busiest in the Philippines, where there is an average of 60 births a day,” theguardian, 2013. Available online at <http://www.theguardian.com/world/gal...471900&index=2> (accessed May 15, 2013).
- “How Americans Use Text Messaging,” Pew Internet, 2013. Available online at <http://pewinternet.org/Reports/2011/...in-Report.aspx> (accessed May 15, 2013).

- Lenhart, Amanda. “Teens, Smartphones & Texting: Texting volum is up while the frequency of voice calling is down. About one in four teens say they own smartphones,” Pew Internet, 2012. Available online at www.pewinternet.org/~media/F...nd_Texting.pdf (accessed May 15, 2013).
- “One born every minute: the maternity unit where mothers are THREE to a bed,” MailOnline. Available online at <http://www.dailymail.co.uk/news/arti...thers-bed.html> (accessed May 15, 2013).
- Vanderkam, Laura. “Stop Checking Your Email, Now.” CNNMoney, 2013. Available online at management.fortune.cnn.com/20...our-email-now/ (accessed May 15, 2013).
- “World Earthquakes: Live Earthquake News and Highlights,” World Earthquakes, 2012. www.world-earthquakes.com/ind...thq_prediction (accessed May 15, 2013).

This page titled [4.11: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [4.9: Chapter References](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

4.12: Chapter Solution (Practice + Homework)

1.

Table 4.12.6

x	$P(x)$
0	0.12
1	0.18
2	0.30
3	0.15
4	0.10
5	0.10
6	0.05

3.

$$0.10 + 0.05 = 0.15$$

5.

1

7.

$$0.35 + 0.40 + 0.10 = 0.85$$

9.

$$1(0.15) + 2(0.35) + 3(0.40) + 4(0.10) = 0.15 + 0.70 + 1.20 + 0.40 = 2.45$$

11.

Table 4.12.7

x	$P(x)$
0	0.03
1	0.04
2	0.08
3	0.85

13.

Let X = the number of events Javier volunteers for each month.

15.

Table 4.12.8

x	$P(x)$
0	0.05
1	0.05
2	0.10
3	0.20
4	0.25
5	0.35

17.

$$1 - 0.05 = 0.95$$

18.

X = the number of business majors in the sample.

19.

2, 3, 4, 5, 6, 7, 8, 9

20.

X = the number that reply "yes"

22.

0, 1, 2, 3, 4, 5, 6, 7, 8

24.

5.7

26.

0.4151

28.

X = the number of freshmen selected from the study until one replied "yes" that same-sex couples should have the right to legal marital status.

30.

1, 2, ...

32.

1.4

35.

0, 1, 2, 3, 4, ...

37.

0.0485

39.

0.0214

41.

X = the number of U.S. teens who die from motor vehicle injuries per day.

43.

0, 1, 2, 3, 4, ...

45.

No

48.

1. 50.

1. 53.

X = the number of patients calling in claiming to have the flu, who actually have the flu.

55.

0.0165

57.

1. 59.

4. 4.43

4

63.

▪ 65.

1. 67.

1. 69.

1. 71.

1. 73.

1. Figure 4.12.4

2. 75.

1. 77.

1. 79.

0, 1, 2, and 3

1. 82.

1. 84.

Let X = the number of defective bulbs in a string.

- Using the binomial distribution:
 - The Poisson approximation is very good—the difference between the probabilities is only 0.0026

86.

1. 88.

1. 90.

1. 92.

1. 94.

4

This page titled [4.12: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [4.11: Chapter Solution \(Practice + Homework\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

5: Continuous Random Variables

- [5.1: Prelude to Continuous Random Variables](#)
- [5.2: Properties of Continuous Probability Density Functions](#)
- [5.3: The Uniform Distribution](#)
- [5.4: The Exponential Distribution](#)
- [5.5: Chapter Key Terms](#)
- [5.6: Chapter Review](#)
- [5.7: Chapter Formula Review](#)
- [5.8: Chapter Practice](#)
- [5.9: Chapter Homework](#)
- [5.10: Chapter References](#)
- [5.11: Chapter Solution \(Practice + Homework\)](#)

This page titled [5: Continuous Random Variables](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

5.1: Prelude to Continuous Random Variables

Continuous random variables have many applications. Baseball batting averages, IQ scores, the length of time a long distance telephone call lasts, the amount of money a person carries, the length of time a computer chip lasts, rates of return from an investment, and SAT scores are just a few. The field of reliability depends on a variety of continuous random variables, as do all areas of risk analysis.



Figure 5.1.1 The heights of these radish plants are continuous random variables. (Credit: Rev Stan)

Note

The values of discrete and continuous random variables can be ambiguous. For example, if X is equal to the number of miles (to the nearest mile) you drive to work, then X is a discrete random variable. You count the miles. If X is the distance you drive to work, then you measure values of X and X is a continuous random variable. For a second example, if X is equal to the number of books in a backpack, then X is a discrete random variable. If X is the weight of a book, then X is a continuous random variable because weights are measured. How the random variable is defined is very important.

This page titled [5.1: Prelude to Continuous Random Variables](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.0: Prelude to Continuous Random Variables](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.2: Properties of Continuous Probability Density Functions

The graph of a continuous probability distribution is a curve. Probability is represented by area under the curve. We have already met this concept when we developed relative frequencies with histograms in [Chapter 2](#). The relative area for a range of values was the probability of drawing at random an observation in that group. Again with the Poisson distribution in [Chapter 4](#), the graph in [Example 5.2.14](#) used boxes to represent the probability of specific values of the random variable. In this case, we were being a bit casual because the random variables of a Poisson distribution are discrete, whole numbers, and a box has width. Notice that the horizontal axis, the random variable x , purposefully did not mark the points along the axis. The probability of a specific value of a continuous random variable will be zero because the area under a point is zero. Probability is area.

The curve is called the **probability density function** (abbreviated as **pdf**). We use the symbol $f(x)$ to represent the curve. $f(x)$ is the function that corresponds to the graph; we use the density function $f(x)$ to draw the graph of the probability distribution.

Area under the curve is given by a different function called the **cumulative distribution function** (abbreviated as **cdf**). The cumulative distribution function is used to evaluate probability as area. Mathematically, the cumulative probability density function is the integral of the pdf, and the probability between two values of a continuous random variable will be the integral of the pdf between these two values: the area under the curve between these values. Remember that the area under the pdf for all possible values of the random variable is one, certainty. Probability thus can be seen as the relative percent of certainty between the two values of interest.

- The outcomes are measured, not counted.
- The entire area under the curve and above the x-axis is equal to one.
- Probability is found for intervals of x values rather than for individual x values.
- $P(c < x < d)$ is the probability that the random variable X is in the interval between the values c and d . $P(c < x < d)$ is the area under the curve, above the x-axis, to the right of c and the left of d .
- $P(x = c) = 0$ The probability that x takes on any single individual value is zero. The area below the curve, above the x-axis, and between $x = c$ and $x = c$ has no width, and therefore no area (area = 0). Since the probability is equal to the area, the probability is also zero.
- $P(c < x < d)$ is the same as $P(c \leq x \leq d)$ because probability is equal to area.

We will find the area that represents probability by using geometry, formulas, technology, or probability tables. In general, integral calculus is needed to find the area under the curve for many probability density functions. When we use formulas to find the area in this textbook, the formulas were found by using the techniques of integral calculus.

There are many continuous probability distributions. When using a continuous probability distribution to model probability, the distribution used is selected to model and fit the particular situation in the best way.

In this chapter and the next, we will study the uniform distribution, the exponential distribution, and the normal distribution. The following graphs illustrate these distributions.




Figure 5.2.2 The graph shows a Uniform Distribution with the area between $x = 3$ and $x = 6$ shaded to represent the probability that the value of the random variable X is in the interval between three and six.




Figure 5.2.3 The graph shows an Exponential Distribution with the area between $x = 2$ and $x = 4$ shaded to represent the probability that the value of the random variable X is in the interval between two and four.




Figure 5.2.4 The graph shows the Standard Normal Distribution with the area between $x = 1$ and $x = 2$ shaded to represent the probability that the value of the random variable X is in the interval between one and two.

For continuous probability distributions, PROBABILITY = AREA.

Example 5.2.1

Consider the function $f(x) = \frac{1}{20}$ for $0 \leq x \leq 20$. x is a real number. The graph of $f(x) = \frac{1}{20}$ is a horizontal line. However, since $0 \leq x \leq 20$, $f(x)$ is restricted to the portion between $x = 0$ and $x = 20$, inclusive.




Figure 5.2.5

$$f(x) = \frac{1}{20} \text{ for } 0 \leq x \leq 20.$$

The graph of $f(x) = \frac{1}{20}$ is a horizontal line segment when $0 \leq x \leq 20$.

The area between $f(x) = \frac{1}{20}$ where $0 \leq x \leq 20$ and the x -axis is the area of a rectangle with base = 20 and height = $\frac{1}{20}$.

$$\text{AREA} = 20 \left(\frac{1}{20} \right) = 1$$

Suppose we want to find the area between $f(x) = \frac{1}{20}$ and the x -axis where $0 < x < 2$.




Figure 5.2.6

$$\text{AREA} = (2 - 0) \left(\frac{1}{20} \right) = 0.1$$

$$(2 - 0) = 2 = \text{base of rectangle}$$

REMINDER

area of a rectangle = (base)(height).

The area corresponds to a probability. The probability that x is between zero and two is 0.1, which can be written mathematically as $P(0 < x < 2) = P(x < 2) = 0.1$.

Suppose we want to find the area between $f(x) = \frac{1}{20}$ and the x -axis where $4 < x < 15$.




Figure 5.2.7

$$\text{AREA} = (15 - 4) \left(\frac{1}{20}\right) = 0.55$$

$(15 - 4) = 11 =$ the base of a rectangle

The area corresponds to the probability $P(4 < x < 15) = 0.55$.

Suppose we want to find $P(x = 15)$. On an x-y graph, $x = 15$ is a vertical line. A vertical line has no width (or zero width). Therefore, $P(x = 15) = (\text{base})(\text{height}) = (0) \left(\frac{1}{20}\right) = 0$




Figure 5.2.8

$P(X \leq x)$, which can also be written as $P(X < x)$ for continuous distributions, is called the cumulative distribution function or CDF. Notice the "less than or equal to" symbol. We can also use the CDF to calculate $P(X > x)$. The CDF gives "area to the left" and $P(X > x)$ gives "area to the right." We calculate $P(X > x)$ for continuous distributions as follows: $P(X > x) = 1 - P(X < x)$.




Figure 5.2.9

Label the graph with $f(x)$ and x . Scale the x and y axes with the maximum x and y values. $f(x) = \frac{1}{20}, 0 \leq x \leq 20$.

To calculate the probability that x is between two values, look at the following graph. Shade the region between $x = 2.3$ and $x = 12.7$. Then calculate the shaded area of a rectangle.




Figure 5.2.10

$$P(2.3 < x < 12.7) = (\text{base})(\text{height}) = (12.7 - 2.3) \left(\frac{1}{20}\right) = 0.52$$

Exercise 5.2.1

Consider the function $f(x) = \frac{1}{8}$ for $0 \leq x \leq 8$. Draw the graph of $f(x)$ and find $P(2.5 < x < 7.5)$.

This page titled [5.2: Properties of Continuous Probability Density Functions](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.1: Properties of Continuous Probability Density Functions](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.3: The Uniform Distribution

The uniform distribution is a continuous probability distribution and is concerned with events that are equally likely to occur. When working out problems that have a uniform distribution, be careful to note if the data is inclusive or exclusive of endpoints.

The mathematical statement of the uniform distribution is

$$f(x) = \frac{1}{b-a} \text{ for } a \leq x \leq b$$

where a = the lowest value of x and b = the highest value of x .

Formulas for the theoretical mean and standard deviation are

$$\mu = \frac{a+b}{2} \text{ and } \sigma = \sqrt{\frac{(b-a)^2}{12}}$$

Exercise 5.3.1

The data that follow are the number of passengers on 35 different charter fishing boats. The sample mean = 7.9 and the sample standard deviation = 4.33. The data follow a uniform distribution where all values between and including zero and 14 are equally likely. State the values of a and b . Write the distribution in proper notation, and calculate the theoretical mean and standard deviation.

1	12	4	10	4	14	11
7	11	4	13	2	4	6
3	10	0	12	6	9	10
5	13	4	10	14	12	11
6	10	11	0	11	13	2

Table 5.1

Example 5.3.2

The amount of time, in minutes, that a person must wait for a bus is uniformly distributed between zero and 15 minutes, inclusive.

a. What is the probability that a person waits fewer than 12.5 minutes?

Answer

a. Let X = the number of minutes a person must wait for a bus. $a = 0$ and $b = 15$. $X \sim U(0, 15)$. Write the probability density function. $f(x) = \frac{1}{15-0} = \frac{1}{15}$ for $0 \leq x \leq 15$.

Find $P(x < 12.5)$. Draw a graph.

$$P(x < k) = (\text{base})(\text{height}) = (12.5 - 0) \left(\frac{1}{15} \right) = 0.8333$$

The probability a person waits less than 12.5 minutes is 0.8333.




Figure 5.11

b. On the average, how long must a person wait? Find the mean, μ , and the standard deviation, σ .

Answer

b. $\mu = \frac{a+b}{2} = \frac{15+0}{2} = 7.5$. On the average, a person must wait 7.5 minutes.

$\sigma = \sqrt{\frac{(b-a)^2}{12}} = \sqrt{\frac{(15-0)^2}{12}} = 4.3$. The Standard deviation is 4.3 minutes.

c. Ninety percent of the time, the time a person must wait falls below what value?

Note

This asks for the 90th percentile.

Answer

c. Find the 90th percentile. Draw a graph. Let k = the 90th percentile.

$$P(x < k) = 0.90 = (k) \left(\frac{1}{15} \right)$$

$$k = (0.90)(15) = 13.5$$

The 90th percentile is 13.5 minutes. Ninety percent of the time, a person must wait at most 13.5 minutes.




Figure 5.3.12

Exercise 5.3.2

The total duration of baseball games in the major league in the 2011 season is uniformly distributed between 447 hours and 521 hours inclusive.

1. Find a and b and describe what they represent.
2. Write the distribution.
3. Find the mean and the standard deviation.
4. What is the probability that the duration of games for a team for the 2011 season is between 480 and 500 hours?

This page titled [5.3: The Uniform Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.2: The Uniform Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.4: The Exponential Distribution

The **exponential distribution** is often concerned with the amount of time until some specific event occurs. For example, the amount of time (beginning now) until an earthquake occurs has an exponential distribution. Other examples include the length of time, in minutes, of long distance business telephone calls, and the amount of time, in months, a car battery lasts. It can be shown, too, that the value of the change that you have in your pocket or purse approximately follows an exponential distribution.

Values for an exponential random variable occur in the following way. There are fewer large values and more small values. For example, marketing studies have shown that the amount of money customers spend in one trip to the supermarket follows an exponential distribution. There are more people who spend small amounts of money and fewer people who spend large amounts of money.

Exponential distributions are commonly used in calculations of product reliability, or the length of time a product lasts.

The random variable for the exponential distribution is continuous and often measures a passage of time, although it can be used in other applications. Typical questions may be, “what is the probability that some event will occur within the next x hours or days, or what is the probability that some event will occur between x_1 hours and x_2 hours, or what is the probability that the event will take more than x_1 hours to perform?” In short, the random variable X equals (a) the time between events or (b) the passage of time to complete an action, e.g. wait on a customer. The probability density function is given by:

$$f(x) = \frac{1}{\mu} e^{-\frac{1}{\mu}x}$$

where μ is the historical average waiting time.

and has a mean and standard deviation of $1/\mu$.

An alternative form of the exponential distribution formula recognizes what is often called the decay factor. The decay factor simply measures how rapidly the probability of an event declines as the random variable X increases. When the notation using the decay parameter m is used, the probability density function is presented as:

$$f(x) = me^{-mx}$$

where $m = \frac{1}{\mu}$

In order to calculate probabilities for specific probability density functions, the cumulative density function is used. The cumulative density function (cdf) is simply the integral of the pdf and is:

$$F(x) = \int_0^{\infty} \left[\frac{1}{\mu} e^{-\frac{x}{\mu}} \right] = 1 - e^{-\frac{x}{\mu}}$$

Example 5.4.3

Let X = amount of time (in minutes) a postal clerk spends with a customer. The time is known from historical data to have an average amount of time equal to four minutes.

It is given that $\mu = 4$ minutes, that is, the average time the clerk spends with a customer is 4 minutes. Remember that we are still doing probability and thus we have to be told the population parameters such as the mean. To do any calculations, we need to know the mean of the distribution: the historical time to provide a service, for example. Knowing the historical mean allows the calculation of the decay parameter, m .

$m = \frac{1}{\mu}$. Therefore, $m = \frac{1}{4} = 0.25$.

When the notation used the decay parameter, m , the probability density function is presented as $f(x) = me^{-mx}$, which is simply the original formula with m substituted for $\frac{1}{\mu}$, or $f(x) = \frac{1}{\mu} e^{-\frac{1}{\mu}x}$.

To calculate probabilities for an exponential probability density function, we need to use the cumulative density function. As shown below, the curve for the cumulative density function is:

$f(x) = 0.25e^{-0.25x}$ where x is at least zero and $m = 0.25$.

For example, $f(5) = 0.25e^{-(0.25)(5)} = 0.072$. In other words, the function has a value of .072 when $x = 5$.

The graph is as follows:




Figure 5.4.13

Notice the graph is a declining curve. When $x = 0$,

$f(x) = 0.25e^{(-0.25)(0)} = (0.25)(1) = 0.25 = m$. The maximum value on the y-axis is always m , one divided by the mean.

Exercise 5.4.3

The amount of time spouses shop for anniversary cards can be modeled by an exponential distribution with the average amount of time equal to eight minutes. Write the distribution, state the probability density function, and graph the distribution.

Example 5.4.4

a. Using the information in [Example 5.4.3](#), find the probability that a clerk spends four to five minutes with a randomly selected customer.

Answer

a. Find $P(4 < x < 5)$.

The cumulative distribution function (CDF) gives the area to the left.

$$P(x < x) = 1 - e^{-mx}$$

$$P(x < 5) = 1 - e^{(-0.25)(5)} = 0.7135 \text{ and } P(x < 4) = 1 - e^{(-0.25)(4)} = 0.6321$$

$$P(4 < x < 5) = 0.7135 - 0.6321 = 0.0814$$




Figure 5.14

Exercise 5.4.4

The number of days ahead travelers purchase their airline tickets can be modeled by an exponential distribution with the average amount of time equal to 15 days. Find the probability that a traveler will purchase a ticket fewer than ten days in advance. How many days do half of all travelers wait?

Example 5.4.5

On the average, a certain computer part lasts ten years. The length of time the computer part lasts is exponentially distributed.

a. What is the probability that a computer part lasts more than 7 years?

Answer

a. Let x = the amount of time (in years) a computer part lasts.

$$\mu = 10 \text{ so } m = \frac{1}{\mu} = \frac{1}{10} = 0.1$$

Find $P(x > 7)$. Draw the graph.

$$P(x > 7) = 1 - P(x < 7)$$

Since $P(X < x) = 1 - e^{-mx}$ then $P(X > x) = 1 - (1 - e^{-mx}) = e^{-mx}$

$P(x > 7) = e^{-(0.1)(7)} = 0.4966$. The probability that a computer part lasts more than seven years is 0.4966




Figure 5.4.15

b. On the average, how long would five computer parts last if they are used one after another?

Answer

b. On the average, one computer part lasts ten years. Therefore, five computer parts, if they are used one right after the other would last, on the average, $(5)(10) = 50$ years.

d. What is the probability that a computer part lasts between nine and 11 years?

Answer

d. Find $P(9 < x < 11)$. Draw the graph.




Figure 5.4.16

$P(9 < x < 11) = P(x < 11) - P(x < 9) = (1 - e^{-(0.1)(11)}) - (1 - e^{-(0.1)(9)}) = 0.6671 - 0.5934 = 0.0737$. The probability that a computer part lasts between nine and 11 years is 0.0737

Exercise 5.4.5

On average, a pair of running shoes can last 18 months if used every day. The length of time running shoes last is exponentially distributed. What is the probability that a pair of running shoes last more than 15 months? On average, how long would six pairs of running shoes last if they are used one after the other? Eighty percent of running shoes last at most how long if used every day?

Example 5.4.6

Suppose that the length of a phone call, in minutes, is an exponential random variable with decay parameter $\frac{1}{12}$. The decay parameter is another way to view $1/\lambda$. If another person arrives at a public telephone just before you, find the probability that you will have to wait more than five minutes. Let X = the length of a phone call, in minutes.

What is m , μ , and σ ? The probability that you must wait more than five minutes is _____.

Answer

$$m = \frac{1}{12}$$

$$\mu = 12$$

$$\sigma = 12$$

$$P(x > 5) = 0.6592$$

Example 5.4.7

The time spent waiting between events is often modeled using the exponential distribution. For example, suppose that an average of 30 customers per hour arrive at a store and the time between arrivals is exponentially distributed.

1. On average, how many minutes elapse between two successive arrivals?
2. When the store first opens, how long on average does it take for three customers to arrive?
3. After a customer arrives, find the probability that it takes less than one minute for the next customer to arrive.
4. After a customer arrives, find the probability that it takes more than five minutes for the next customer to arrive.
5. Is an exponential distribution reasonable for this situation?

Answer

a. Since we expect 30 customers to arrive per hour (60 minutes), we expect on average one customer to arrive every two minutes on average.

b. Since one customer arrives every two minutes on average, it will take six minutes on average for three customers to arrive.

c. Let X = the time between arrivals, in minutes. By part a, $\mu = 2$, so $m = \frac{1}{2} = 0.5$.

The cumulative distribution function is $P(X < x) = 1 - e^{(-0.5)(x)}$

Therefore $P(X < 1) = 1 - e^{(-0.5)(1)} = 0.3935$.




Figure 5.4.17

d. $P(X > 5) = 1 - P(X < 5) = 1 - (1 - e^{(-0.5)(5)}) = e^{-2.5} \approx 0.0821$.




Figure 5.4.18

This model assumes that a single customer arrives at a time, which may not be reasonable since people might shop in groups, leading to several customers arriving at the same time. It also assumes that the flow of customers does not change throughout the day, which is not valid if some times of the day are busier than others.

Memorylessness of the Exponential Distribution

Recall that the amount of time between customers for the postal clerk discussed earlier is exponentially distributed with a mean of two minutes. Suppose that five minutes have elapsed since the last customer arrived. Since an unusually long amount of time has now elapsed, it would seem to be more likely for a customer to arrive within the next minute. With the exponential distribution, this is not the case—the additional time spent waiting for the next customer does not depend on how much time has already elapsed since the last customer. This is referred to as the **memoryless property**. The exponential and geometric probability density functions are the only probability functions that have the memoryless property. Specifically, the **memoryless property** says that

$$P(X > r + t | X > r) = P(X > t) \quad \text{for all } r \geq 0 \text{ and } t \geq 0$$

For example, if five minutes have elapsed since the last customer arrived, then the probability that more than one minute will elapse before the next customer arrives is computed by using $r = 5$ and $t = 1$ in the foregoing equation.

$$P(X > 5 + 1 | X > 5) = P(X > 1) = e^{(-0.5)(1)} = 0.6065 .$$

This is the same probability as that of waiting more than one minute for a customer to arrive after the previous arrival.

The exponential distribution is often used to model the longevity of an electrical or mechanical device. In [Example 5.4.5](#), the lifetime of a certain computer part has the exponential distribution with a mean of ten years. The **memoryless property** says that knowledge of what has occurred in the past has no effect on future probabilities. In this case it means that an old part is not any more likely to break down at any particular time than a brand new part. In other words, the part stays as good as new until it suddenly breaks. For example, if the part has already lasted ten years, then the probability that it lasts another seven years is $P(X > 17 | X > 10) = P(X > 7) = 0.4966$, where the vertical line is read as "given".

Example 5.4.8

Refer back to the postal clerk again where the time a postal clerk spends with his or her customer has an exponential distribution with a mean of four minutes. Suppose a customer has spent four minutes with a postal clerk. What is the probability that he or she will spend at least an additional three minutes with the postal clerk?

The decay parameter of X is $m = \frac{1}{4} = 0.25$, so $X \sim Exp(0.25)$.

The cumulative distribution function is $P(X < x) = 1 - e^{-0.25x}$.

We want to find $P(X > 7 | X > 4)$. The **memoryless property** says that $P(X > 7 | X > 4) = P(X > 3)$, so we just need to find the probability that a customer spends more than three minutes with a postal clerk.

This is $P(X > 3) = 1 - P(X < 3) = 1 - (1 - e^{-0.25 \cdot 3}) = e^{-0.75} \approx 0.4724$.




Figure 5.4.19

Relationship between the Poisson and the Exponential Distribution

There is an interesting relationship between the exponential distribution and the Poisson distribution. Suppose that the time that elapses between two successive events follows the exponential distribution with a mean of μ units of time. Also assume that these times are independent, meaning that the time between events is not affected by the times between previous events. If these assumptions hold, then the number of events per unit time follows a Poisson distribution with mean μ . Recall that if X has the Poisson distribution with mean μ , then $P(X = x) = \frac{\mu^x e^{-\mu}}{x!}$.

The formula for the exponential distribution: $P(X = x) = me^{-mx} = \frac{1}{\mu}e^{-\frac{1}{\mu}x}$ Where m = the rate parameter, or μ = average time between occurrences.

We see that the exponential is the cousin of the Poisson distribution and they are linked through this formula. There are important differences that make each distribution relevant for different types of probability problems.

First, the Poisson has a discrete random variable, x , where time; a continuous variable is artificially broken into discrete pieces. We saw that the number of occurrences of an event in a given time interval, x , follows the Poisson distribution.

For example, the **number** of times the telephone rings per hour. By contrast, the time **between** occurrences follows the exponential distribution. For example. The telephone just rang, how long will it be until it rings again? We are measuring length of time of the interval, a continuous random variable, exponential, not events during an interval, Poisson.

The Exponential Distribution v. the Poisson Distribution

A visual way to show both the similarities and differences between these two distributions is with a time line.




Figure 5.20

The random variable for the Poisson distribution is discrete and thus counts events during a given time period, t_1 to t_2 on Figure 5.4.20, and calculates the probability of that number occurring. The number of events, four in the graph, is measured in counting numbers; therefore, the random variable of the Poisson is a discrete random variable.

The exponential probability distribution calculates probabilities of the passage of time, a continuous random variable. In Figure 5.4.20 this is shown as the bracket from t_1 to the next occurrence of the event marked with a triangle.

Classic Poisson distribution questions are "how many people will arrive at my checkout window in the next hour?".

Classic exponential distribution questions are "how long it will be until the next person arrives," or a variant, "how long will the person remain here once they have arrived?".

Again, the formula for the exponential distribution is:

$$f(x) = me^{-mx} \text{ or } f(x) = \frac{1}{\mu}e^{-\frac{1}{\mu}x}$$

We see immediately the similarity between the exponential formula and the Poisson formula.

$$P(x) = \frac{\mu^x e^{-\mu}}{x!}$$

Both probability density functions are based upon the relationship between time and exponential growth or decay. The "e" in the formula is a constant with the approximate value of 2.71828 and is the base of the natural logarithmic exponential growth formula. When people say that something has grown exponentially this is what they are talking about.

An example of the exponential and the Poisson will make clear the differences between the two. It will also show the interesting applications they have.

Poisson Distribution

Suppose that historically 10 customers arrive at the checkout lines each hour. Remember that this is still probability so we have to be told these historical values. We see this is a Poisson probability problem.

We can put this information into the Poisson probability density function and get a general formula that will calculate the probability of **any** specific number of customers arriving in the next hour.

The formula is for any value of the random variable we chose, and so the x is put into the formula. This is the formula:

$$f(x) = \frac{10^x e^{-10}}{x!}$$

As an example, the probability of 15 people arriving at the checkout counter in the next hour would be

$$P(x = 15) = \frac{10^{15} e^{-10}}{15!} = 0.0611$$

Here we have inserted $x = 15$ and calculated the probability that in the next hour 15 people will arrive is .061.

Exponential Distribution

If we keep the same historical facts that 10 customers arrive each hour, but we now are interested in the service time a person spends at the counter, then we would use the exponential distribution. The exponential probability function for any value of x , the random variable, for this particular checkout counter historical data is:

$$f(x) = \frac{1}{.1} e^{-x/.1} = 10e^{-10x}$$

To calculate μ , the historical average service time, we simply divide the number of people that arrive per hour, 10, into the time period, one hour, and have $\mu = 0.1$. Historically, people spend 0.1 of an hour at the checkout counter, or 6 minutes. This explains the .1 in the formula.

There is a natural confusion with μ in both the Poisson and exponential formulas. They have different meanings, although they have the same symbol. The mean of the exponential is one divided by the mean of the Poisson. If you are given the historical number of arrivals you have the mean of the Poisson. If you are given an historical length of time between events you have the mean of an exponential.

Continuing with our example at the checkout clerk; if we wanted to know the probability that a person would spend 9 minutes or less checking out, then we use this formula. First, we convert to the same time units which are parts of one hour. Nine minutes is 0.15 of one hour. Next we note that we are asking for a range of values. This is always the case for a continuous random variable. We write the probability question as:

$$p(x \leq 9) = 1 - 10e^{-10x}$$

We can now put the numbers into the formula and we have our result.

$$p(x = .15) = 1 - 10e^{-10(.15)} = 0.7769$$

The probability that a customer will spend 9 minutes or less checking out is 0.7769.

We see that we have a high probability of getting out in less than nine minutes and a tiny probability of having 15 customers arriving in the next hour.

This page titled [5.4: The Exponential Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.3: The Exponential Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.5: Chapter Key Terms

Conditional Probability

the likelihood that an event will occur given that another event has already occurred.

decay parameter

The decay parameter describes the rate at which probabilities decay to zero for increasing values of x . It is the value m in the probability density function $f(x) = me^{-mx}$ of an exponential random variable. It is also equal to $m = \frac{1}{\mu}$, where μ is the mean of the random variable.

Exponential Distribution

a continuous random variable (RV) that appears when we are interested in the intervals of time between some random events, for example, the length of time between emergency arrivals at a hospital. The mean is $\mu = \frac{1}{m}$ and the standard deviation is $\sigma = \frac{1}{m}$. The probability density function is $f(x) = me^{-mx}$ or $f(x) = \frac{1}{\mu}e^{-\frac{1}{\mu}x}$, $x \geq 0$ and the cumulative distribution function is $P(X \leq x) = 1 - e^{-mx}$ or $P(X \leq x) = 1 - e^{-\frac{1}{\mu}x}$.

memoryless property

For an exponential random variable X , the memoryless property is the statement that knowledge of what has occurred in the past has no effect on future probabilities. This means that the probability that X exceeds $x + t$, given that it has exceeded x , is the same as the probability that X would exceed t if we had no knowledge about it. In symbols we say that $P(X > x + t | X > x) = P(X > t)$.

Poisson distribution

If there is a known average of μ events occurring per unit time, and these events are independent of each other, then the number of events X occurring in one unit of time has the Poisson distribution. The probability of x events occurring in one unit time is equal to $P(X = x) = \frac{\mu^x e^{-\mu}}{x!}$.

Uniform Distribution

a continuous random variable (RV) that has equally likely outcomes over the domain, $a < x < b$; it is often referred as the rectangular distribution because the graph of the pdf has the form of a rectangle. The mean is $\mu = \frac{a+b}{2}$ and the standard deviation is $\sigma = \sqrt{\frac{(b-a)^2}{12}}$. The probability density function is $f(x) = \frac{1}{b-a}$ for $a < x < b$.

This page titled [5.5: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

5.6: Chapter Review

5.1 Properties of Continuous Probability Density Functions

The probability density function (pdf) is used to describe probabilities for continuous random variables. The area under the density curve between two points corresponds to the probability that the variable falls between those two values. In other words, the area under the density curve between points a and b is equal to $P(a < x < b)$. The cumulative distribution function (cdf) gives the probability as an area. If X is a continuous random variable, the probability density function (pdf), $f(x)$, is used to draw the graph of the probability distribution. The total area under the graph of $f(x)$ is one. The area under the graph of $f(x)$ and between values a and b gives the probability $P(a < x < b)$.




Figure 5.6.21

The cumulative distribution function (cdf) of X is defined by $P(X \leq x)$. It is a function of x that gives the probability that the random variable is less than or equal to x .

5.2 The Uniform Distribution

If X has a uniform distribution where $a < x < b$ or $a \leq x \leq b$, then X takes on values between a and b (may include a and b).

All values x are equally likely. We write $X \sim U(a, b)$. The mean of X is $\mu = \frac{a+b}{2}$. The standard deviation of X is $\sigma = \sqrt{\frac{(b-a)^2}{12}}$. The probability density function of X is $f(x) = \frac{1}{b-a}$ for $a \leq x \leq b$. The cumulative distribution function of X is $P(X \leq x) = \frac{x-a}{b-a}$. X is continuous.




Figure 5.6.22

The probability $P(c < X < d)$ may be found by computing the area under $f(x)$, between c and d . Since the corresponding area is a rectangle, the area may be found simply by multiplying the width and the height.

5.3 The Exponential Distribution

If X has an **exponential distribution** with mean μ , then the **decay parameter** is $m = \frac{1}{\mu}$. The probability density function of X is $f(x) = me^{-mx}$ (or equivalently $f(x) = \frac{1}{\mu}e^{-x/\mu}$). The cumulative distribution function of X is $P(X \leq x) = 1 - e^{-mx}$.

This page titled [5.6: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.9: Chapter Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.7: Chapter Formula Review

5.1 Properties of Continuous Probability Density Functions

Probability density function (pdf) $f(x)$:

- Cumulative distribution function (cdf): $P(X \leq x)$

5.2 The Uniform Distribution

$$X \sim U(a, b)$$

The mean is $\mu = \frac{a+b}{2}$

The standard deviation is $\sigma = \sqrt{\frac{(b-a)^2}{12}}$

Probability density function: $f(x) = \frac{1}{b-a}$ for $a \leq X \leq b$

Area to the Left of x : $\backslash P(X < x) >$

Area to the Right of x : $P(X > x) = (b - x) \left(\frac{1}{b-a} \right)$

Area Between c and d : $\backslash P(c < d) >$

5.3 The Exponential Distribution

- pdf: $f(x) = me^{-mx}$ where $x \geq 0$ and $m > 0$
- cdf: $P(X \leq x) = 1 - e^{-mx}$
- mean $\mu = \frac{1}{m}$
- standard deviation $\sigma = \mu$
- Additionally
 - $P(X > x) = e^{-mx}$
 - $P(a < X < b) = e^{-ma} - e^{-mb}$
- Poisson probability: $P(X = x) = \frac{\mu^x e^{-\mu}}{x!}$ with mean and variance of μ

This page titled [5.7: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.4: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.8: Chapter Practice

5.1 Properties of Continuous Probability Density Functions

1.

Which type of distribution does the graph illustrate?




Figure 5.8.23

2.

Which type of distribution does the graph illustrate?




Figure 5.8.24

3.

Which type of distribution does the graph illustrate?




Figure 5.8.25

4.

What does the shaded area represent? $P(__ < x < __)$




Figure 5.8.26

5.

What does the shaded area represent? $P(__ < x < __)$




Figure 5.8.27

6.

For a continuous probability distribution, $0 \leq x \leq 15$. What is $P(x > 15)$?

7.

What is the area under $f(x)$ if the function is a continuous probability density function?

8.

For a continuous probability distribution, $0 \leq x \leq 10$. What is $P(x = 7)$?

9.

A **continuous** probability function is restricted to the portion between $x = 0$ and 7. What is $P(x = 10)$?

10.

$f(x)$ for a continuous probability function is $\frac{1}{5}$, and the function is restricted to $0 \leq x \leq 5$. What is $P(x < 0)$?

11.

$f(x)$, a continuous probability function, is equal to $\frac{1}{12}$, and the function is restricted to $0 \leq x \leq 12$. What is $P(0 < x < 12)$?

12.

Find the probability that x falls in the shaded area.




Figure 5.8.28

13.

Find the probability that x falls in the shaded area.




Figure 5.8.29

14.

Find the probability that x falls in the shaded area.




Figure 5.8.30

15.

$f(x)$, a continuous probability function, is equal to $\frac{1}{3}$ and the function is restricted to $1 \leq x \leq 4$. Describe $P(x > \frac{3}{2})$.

5.2 The Uniform Distribution

Use the following information to answer the next ten questions. The data that follow are the square footage (in 1,000 feet squared) of 28 homes.

Table 5.8.2

1.5	2.4	3.6	2.6	1.6	2.4	2.0
3.5	2.5	1.8	2.4	2.5	3.5	4.0
2.6	1.6	2.2	1.8	3.8	2.5	1.5
2.8	1.8	4.5	1.9	1.9	3.1	1.6

The sample mean = 2.50 and the sample standard deviation = 0.8302.

The distribution can be written as $X \sim U(1.5, 4.5)$.

16.

What type of distribution is this?

17.

In this distribution, outcomes are equally likely. What does this mean?

18.

What is the height of $f(x)$ for the continuous probability distribution?

19.

What are the constraints for the values of x ?

20.

Graph $P(2 < x < 3)$.

21.

What is $P(2 < x < 3)$?

22.

What is $P(x < 3.5 | x < 4)$?

23.

What is $P(x = 1.5)$?

24.

Find the probability that a randomly selected home has more than 3,000 square feet given that you already know the house has more than 2,000 square feet.

Use the following information to answer the next eight exercises. A distribution is given as $X \sim U(0, 12)$.

25.

What is a ? What does it represent?

26.

What is b ? What does it represent?

27.

What is the probability density function?

28.

What is the theoretical mean?

29.

What is the theoretical standard deviation?

30.

Draw the graph of the distribution for $P(x > 9)$.

31.

Find $P(x > 9)$.

Use the following information to answer the next eleven exercises. The age of cars in the staff parking lot of a suburban college is uniformly distributed from six months (0.5 years) to 9.5 years.

32.

What is being measured here?

33.

In words, define the random variable X .

34.

Are the data discrete or continuous?

35.

The interval of values for x is _____.

36.

The distribution for X is _____.

37.

Write the probability density function.

38.

Graph the probability distribution.

1. Sketch the graph of the probability distribution.



Figure 5.8.31

2. Identify the following values:

- Lowest value for \bar{x} : _____
- Highest value for \bar{x} : _____
- Height of the rectangle: _____
- Label for x-axis (words): _____
- Label for y-axis (words): _____

39.

Find the average age of the cars in the lot.

40.

Find the probability that a randomly chosen car in the lot was less than four years old.

a. Sketch the graph, and shade the area of interest.



Figure 5.8.32

b. Find the probability. $P(x < 4) =$ _____

41.

Considering only the cars less than 7.5 years old, find the probability that a randomly chosen car in the lot was less than four years old.

1. Sketch the graph, shade the area of interest.




Figure 5.8.33

2. Find the probability. $P(x < 4 | x < 7.5) =$ _____

42.

What has changed in the previous two problems that made the solutions different?

43.

Find the third quartile of ages of cars in the lot. This means you will have to find the value such that $\frac{3}{4}$, or 75%, of the cars are at most (less than or equal to) that age.

1. Sketch the graph, and shade the area of interest.



Figure 5.8.34

2. Find the value k such that $P(x < k) = 0.75$.
3. The third quartile is _____

5.3 The Exponential Distribution

Use the following information to answer the next ten exercises. A customer service representative must spend different amounts of time with each customer to resolve various concerns. The amount of time spent with each customer can be modeled by the following distribution: $X \sim Exp(0.2)$

44.

What type of distribution is this?

45.

Are outcomes equally likely in this distribution? Why or why not?

46.

What is m ? What does it represent?

47.

What is the mean?

48.

What is the standard deviation?

49.

State the probability density function.

50.

Graph the distribution.

51.

Find $P(2 < x < 10)$.

52.

Find $P(x > 6)$.

53.

Find the 70th percentile.

Use the following information to answer the next seven exercises. A distribution is given as $X \sim Exp(0.75)$.

54.

What is m ?

55.

What is the probability density function?

56.

What is the cumulative distribution function?

57.

Draw the distribution.

58.

Find $P(x < 4)$.

59.

Find the 30th percentile.

60.

Find the median.

61.

Which is larger, the mean or the median?

Use the following information to answer the next 16 exercises. Carbon-14 is a radioactive element with a half-life of about 5,730 years. Carbon-14 is said to decay exponentially. The decay rate is 0.000121. We start with one gram of carbon-14. We are interested in the time (years) it takes to decay carbon-14.

62.

What is being measured here?

63.

Are the data discrete or continuous?

64.

In words, define the random variable X .

65.

What is the decay rate (m)?

66.

The distribution for X is _____.

67.

Find the amount (percent of one gram) of carbon-14 lasting less than 5,730 years. This means, find $P(x < 5,730)$.

1. Sketch the graph, and shade the area of interest.



Figure 5.8.35

2. Find the probability. $P(x < 5,730) =$ _____

68.

Find the percentage of carbon-14 lasting longer than 10,000 years.

1. Sketch the graph, and shade the area of interest.



Figure 5.8.36

2. Find the probability. $P(x > 10,000) =$ _____

69.

Thirty percent (30%) of carbon-14 will decay within how many years?

1. Sketch the graph, and shade the area of interest.



Figure 5.8.37

Find the value k such that $P(x < k) = 0.30$.

This page titled [5.8: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.7: Chapter Practice](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.9: Chapter Homework

5.1 Properties of Continuous Probability Density Functions

For each probability and percentile problem, draw the picture.

70.

Consider the following experiment. You are one of 100 people enlisted to take part in a study to determine the percent of nurses in America with an R.N. (registered nurse) degree. You ask nurses if they have an R.N. degree. The nurses answer “yes” or “no.” You then calculate the percentage of nurses with an R.N. degree. You give that percentage to your supervisor.

1. What part of the experiment will yield discrete data?
2. What part of the experiment will yield continuous data?

71.

When age is rounded to the nearest year, do the data stay continuous, or do they become discrete? Why?

5.2 The Uniform Distribution

For each probability and percentile problem, draw the picture.

72.

Births are approximately uniformly distributed between the 52 weeks of the year. They can be said to follow a uniform distribution from one to 53 (spread of 52 weeks).

1. Graph the probability distribution.
2. $f(x) =$ _____
3. $\mu =$ _____
4. $\sigma =$ _____
5. Find the probability that a person is born at the exact moment week 19 starts. That is, find $P(x = 19) =$ _____
6. $P(2 < x < 31) =$ _____
7. Find the probability that a person is born after week 40.
8. $P(12 < x | x < 28) =$ _____

73.

A random number generator picks a number from one to nine in a uniform manner.

1. Graph the probability distribution.
2. $f(x) =$ _____
3. $\mu =$ _____
4. $\sigma =$ _____
5. $P(3.5 < x < 7.25) =$ _____
6. $P(x > 5.67)$
7. $P(x > 5 | x > 3) =$ _____

74.

According to a study by Dr. John McDougall of his live-in weight loss program at St. Helena Hospital, the people who follow his program lose between six and 15 pounds a month until they approach trim body weight. Let's suppose that the weight loss is uniformly distributed. We are interested in the weight loss of a randomly selected individual following the program for one month.

1. Define the random variable. $X =$ _____
2. Graph the probability distribution.
3. $f(x) =$ _____
4. $\mu =$ _____
5. $\sigma =$ _____
6. Find the probability that the individual lost more than ten pounds in a month.

- Suppose it is known that the individual lost more than ten pounds in a month. Find the probability that he lost less than 12 pounds in the month.
- $P(7 < x < 13 | x > 9) = \underline{\hspace{2cm}}$. State this in a probability question, similarly to parts g and h, draw the picture, and find the probability.

75.

A subway train on the Red Line arrives every eight minutes during rush hour. We are interested in the length of time a commuter must wait for a train to arrive. The time follows a uniform distribution.

- Define the random variable. $X = \underline{\hspace{2cm}}$
- Graph the probability distribution.
- $f(x) = \underline{\hspace{2cm}}$
- $\mu = \underline{\hspace{2cm}}$
- $\sigma = \underline{\hspace{2cm}}$
- Find the probability that the commuter waits less than one minute.
- Find the probability that the commuter waits between three and four minutes.

76.

The age of a first grader on September 1 at Garden Elementary School is uniformly distributed from 5.8 to 6.8 years. We randomly select one first grader from the class.

- Define the random variable. $X = \underline{\hspace{2cm}}$
- Graph the probability distribution.
- $f(x) = \underline{\hspace{2cm}}$
- $\mu = \underline{\hspace{2cm}}$
- $\sigma = \underline{\hspace{2cm}}$
- Find the probability that she is over 6.5 years old.
- Find the probability that she is between four and six years old.

Use the following information to answer the next three exercises. The Sky Train from the terminal to the rental-car and long-term parking center is supposed to arrive every eight minutes. The waiting times for the train are known to follow a uniform distribution.

77.

What is the average waiting time (in minutes)?

- zero
- two
- three
- four

78.

The probability of waiting more than seven minutes given a person has waited more than four minutes is?

- 0.125
- 0.25
- 0.5
- 0.75

79.

The time (in minutes) until the next bus departs a major bus depot follows a distribution with $f(x) = \frac{120}{120 - x}$ where x goes from 25 to 45 minutes.

- Define the random variable. $X = \underline{\hspace{2cm}}$
- Graph the probability distribution.
- The distribution is $\underline{\hspace{2cm}}$ (name of distribution). It is $\underline{\hspace{2cm}}$ (discrete or continuous).
- $\mu = \underline{\hspace{2cm}}$
- $\sigma = \underline{\hspace{2cm}}$

6. Find the probability that the time is at most 30 minutes. Sketch and label a graph of the distribution. Shade the area of interest. Write the answer in a probability statement.
7. Find the probability that the time is between 30 and 40 minutes. Sketch and label a graph of the distribution. Shade the area of interest. Write the answer in a probability statement.
8. $P(25 < x < 55) = \underline{\hspace{2cm}}$. State this in a probability statement, similarly to parts g and h, draw the picture, and find the probability.

80.

Suppose that the value of a stock varies each day from \$16 to \$25 with a uniform distribution.

- a. Find the probability that the value of the stock is more than \$19.
- b. Find the probability that the value of the stock is between \$19 and \$22.
- c. Given that the stock is greater than \$18, find the probability that the stock is more than \$21.

81.

A fireworks show is designed so that the time between fireworks is between one and five seconds, and follows a uniform distribution.

1. Find the average time between fireworks.
2. Find probability that the time between fireworks is greater than four seconds.

82.

The number of miles driven by a truck driver falls between 300 and 700, and follows a uniform distribution.

1. Find the probability that the truck driver goes more than 650 miles in a day.
2. Find the probability that the truck drivers goes between 400 and 650 miles in a day.

5.3 The Exponential Distribution

83.

Suppose that the length of long distance phone calls, measured in minutes, is known to have an exponential distribution with the average length of a call equal to eight minutes.

1. Define the random variable. $X = \underline{\hspace{2cm}}$.
2. Is X continuous or discrete?
3. $\mu = \underline{\hspace{2cm}}$
4. $\sigma = \underline{\hspace{2cm}}$
5. Draw a graph of the probability distribution. Label the axes.
6. Find the probability that a phone call lasts less than nine minutes.
7. Find the probability that a phone call lasts more than nine minutes.
8. Find the probability that a phone call lasts between seven and nine minutes.
9. If 25 phone calls are made one after another, on average, what would you expect the total to be? Why?

84.

Suppose that the useful life of a particular car battery, measured in months, decays with parameter 0.025. We are interested in the life of the battery.

1. Define the random variable. $X = \underline{\hspace{2cm}}$.
2. Is X continuous or discrete?
3. On average, how long would you expect one car battery to last?
4. On average, how long would you expect nine car batteries to last, if they are used one after another?
5. Find the probability that a car battery lasts more than 36 months.
6. Seventy percent of the batteries last at least how long?

85.

The percent of persons (ages five and older) in each state who speak a language at home other than English is approximately exponentially distributed with a mean of 9.848. Suppose we randomly pick a state.

1. Define the random variable. $X =$ _____.
2. Is X continuous or discrete?
3. $\mu =$ _____
4. $\sigma =$ _____
5. Draw a graph of the probability distribution. Label the axes.
6. Find the probability that the percent is less than 12.
7. Find the probability that the percent is between eight and 14.
8. The percent of all individuals living in the United States who speak a language at home other than English is 13.8.
 - o Why is this number different from 9.848%?
 - o What would make this number higher than 9.848%?

86.

The time (in years) **after** reaching age 60 that it takes an individual to retire is approximately exponentially distributed with a mean of about five years. Suppose we randomly pick one retired individual. We are interested in the time after age 60 to retirement.

1. Define the random variable. $X =$ _____.
2. Is X continuous or discrete?
3. $\mu =$ _____
4. $\sigma =$ _____
5. Draw a graph of the probability distribution. Label the axes.
6. Find the probability that the person retired after age 70.
7. Do more people retire before age 65 or after age 65?
8. In a room of 1,000 people over age 80, how many do you expect will NOT have retired yet?

87.

The cost of all maintenance for a car during its first year is approximately exponentially distributed with a mean of \$150.

1. Define the random variable. $X =$ _____.
2. $\mu =$ _____
3. $\sigma =$ _____
4. Draw a graph of the probability distribution. Label the axes.
5. Find the probability that a car required over \$300 for maintenance during its first year.

Use the following information to answer the next three exercises. The average lifetime of a certain new cell phone is three years. The manufacturer will replace any cell phone failing within two years of the date of purchase. The lifetime of these cell phones is known to follow an exponential distribution.

88.

The decay rate is:

- a. 0.3333
- b. 0.5000
- c. 2
- d. 3

89.

What is the probability that a phone will fail within two years of the date of purchase?

- a. 0.8647
- b. 0.4866
- c. 0.2212
- d. 0.9997

90.

What is the median lifetime of these phones (in years)?

- a. 0.1941

- b. 1.3863
- c. 2.0794
- d. 5.5452

91.

At a 911 call center, calls come in at an average rate of one call every two minutes. Assume that the time that elapses from one call to the next has the exponential distribution.

1. On average, how much time occurs between five consecutive calls?
2. Find the probability that after a call is received, it takes more than three minutes for the next call to occur.
3. Ninety-percent of all calls occur within how many minutes of the previous call?
4. Suppose that two minutes have elapsed since the last call. Find the probability that the next call will occur within the next minute.
5. Find the probability that less than 20 calls occur within an hour.

92.

In major league baseball, a no-hitter is a game in which a pitcher, or pitchers, doesn't give up any hits throughout the game. No-hitters occur at a rate of about three per season. Assume that the duration of time between no-hitters is exponential.

1. What is the probability that an entire season elapses with a single no-hitter?
2. If an entire season elapses without any no-hitters, what is the probability that there are no no-hitters in the following season?
3. What is the probability that there are more than 3 no-hitters in a single season?

93.

During the years 1998–2012, a total of 29 earthquakes of magnitude greater than 6.5 have occurred in Papua New Guinea. Assume that the time spent waiting between earthquakes is exponential.

1. What is the probability that the next earthquake occurs within the next three months?
2. Given that six months has passed without an earthquake in Papua New Guinea, what is the probability that the next three months will be free of earthquakes?
3. What is the probability of zero earthquakes occurring in 2014?
4. What is the probability that at least two earthquakes will occur in 2014?

94.

According to the American Red Cross, about one out of nine people in the U.S. have Type B blood. Suppose the blood types of people arriving at a blood drive are independent. In this case, the number of Type B blood types that arrive roughly follows the Poisson distribution.

1. If 100 people arrive, how many on average would be expected to have Type B blood?
2. What is the probability that over 10 people out of these 100 have type B blood?
3. What is the probability that more than 20 people arrive before a person with type B blood is found?

95.

A web site experiences traffic during normal working hours at a rate of 12 visits per hour. Assume that the duration between visits has the exponential distribution.

1. Find the probability that the duration between two successive visits to the web site is more than ten minutes.
2. The top 25% of durations between visits are at least how long?
3. Suppose that 20 minutes have passed since the last visit to the web site. What is the probability that the next visit will occur within the next 5 minutes?
4. Find the probability that less than 7 visits occur within a one-hour period.

96.

At an urgent care facility, patients arrive at an average rate of one patient every seven minutes. Assume that the duration between arrivals is exponentially distributed.

1. Find the probability that the time between two successive visits to the urgent care facility is less than 2 minutes.
2. Find the probability that the time between two successive visits to the urgent care facility is more than 15 minutes.

3. If 10 minutes have passed since the last arrival, what is the probability that the next person will arrive within the next five minutes?
 4. Find the probability that more than eight patients arrive during a half-hour period.
-

This page titled [5.9: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.5: Chapter Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.10: Chapter References

5.2 The Uniform Distribution

McDougall, John A. The McDougall Program for Maximum Weight Loss. Plume, 1995.

5.3 The Exponential Distribution

Data from the United States Census Bureau.

Data from World Earthquakes, 2013. Available online at <http://www.world-earthquakes.com/> (accessed June 11, 2013).

“No-hitter.” Baseball-Reference.com, 2013. Available online at <http://www.baseball-reference.com/bullpen/No-hitter> (accessed June 11, 2013).

Zhou, Rick. “Exponential Distribution lecture slides.” Available online at www.public.iastate.edu/~riczw/stat330s11/lecture/lec13.pdf (accessed June 11, 2013).

This page titled [5.10: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.8: Chapter References](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

5.11: Chapter Solution (Practice + Homework)




Figure 5.11.38

Figure 5.11.41

This page titled [5.11: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [5.10: Chapter Solution \(Practice + Homework\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

6: The Normal Distribution

- 6.1: Introduction to Normal Distribution
- 6.2: The Standard Normal Distribution
- 6.3: Using the Normal Distribution
- 6.4: Estimating the Binomial with the Normal Distribution
- 6.5: Chapter Key Items
- 6.6: Chapter Review
- 6.7: Chapter Formula Review
- 6.8: Chapter Practice
- 6.9: Chapter Homework
- 6.10: Chapter References
- 6.11: Chapter Solution (Practice + Homework)

This page titled [6: The Normal Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

6.1: Introduction to Normal Distribution

The normal probability density function, a continuous distribution, is the most important of all the distributions. It is widely used and even more widely abused. Its graph is bell-shaped. You see the bell curve in almost all disciplines. Some of these include psychology, business, economics, the sciences, nursing, and, of course, mathematics. Some of your instructors may use the normal distribution to help determine your grade. Most IQ scores are normally distributed. Often real-estate prices fit a normal distribution.



Figure 6.1.1 If you ask enough people about their shoe size, you will find that your graphed data is shaped like a bell curve and can be described as normally distributed. (credit: Ömer Ünü)

The normal distribution is extremely important, but it cannot be applied to everything in the real world. Remember here that we are still talking about the distribution of population data. This is a discussion of probability and thus it is the population data that may be normally distributed, and if it is, then this is how we can find probabilities of specific events just as we did for population data that may be binomially distributed or Poisson distributed. This caution is here because in the next chapter we will see that the normal distribution describes something very different from raw data and forms the foundation of inferential statistics.

The normal distribution has two parameters (two numerical descriptive measures): the mean (μ) and the standard deviation (σ). If X is a quantity to be measured that has a normal distribution with mean (μ) and standard deviation (σ), we designate this by writing the following formula of the normal probability density function:




Figure 6.1.2

The probability density function is a rather complicated function. **Do not memorize it.** It is not necessary.

$$f(x) = \frac{1}{\sigma \cdot \sqrt{2 \cdot \pi}} \cdot e^{-\frac{1}{2} \cdot \left(\frac{x-\mu}{\sigma}\right)^2}$$

The curve is symmetric about a vertical line drawn through the mean, μ . The mean is the same as the median, which is the same as the mode, because the graph is symmetric about μ . As the notation indicates, the normal distribution depends only on the mean and the standard deviation. Note that this is unlike several probability density functions we have already studied, such as the Poisson, where the mean is equal to μ and the standard deviation simply the square root of the mean, or the binomial, where p is used to determine both the mean and standard deviation. Since the area under the curve must equal one, a change in the standard deviation, σ , causes a change in the shape of the normal curve; the curve becomes fatter and wider or skinnier and taller depending on σ . A change in μ causes the graph to shift to the left or right. This means there are an infinite number of normal probability distributions. One of special interest is called the **standard normal distribution**.

This page titled [6.1: Introduction to Normal Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **6.0: Introduction to Normal Distribution** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.2: The Standard Normal Distribution

The **standard normal distribution** is a normal distribution of **standardized values called z-scores**. A z-score is measured in **units of the standard deviation**.

The mean for the standard normal distribution is zero, and the standard deviation is one. What this does is dramatically simplify the mathematical calculation of probabilities. Take a moment and substitute zero and one in the appropriate places in the above formula and you can see that the equation collapses into one that can be much more easily solved using integral calculus. The transformation $z = \frac{x-\mu}{\sigma}$ produces the distribution $Z \sim N(0, 1)$. The value x in the given equation comes from a known normal distribution with known mean μ and known standard deviation σ . The z-score tells how many standard deviations a particular x is away from the mean.

Z-Scores

If X is a normally distributed random variable and $X \sim N(\mu, \sigma)$, then the z-score for a particular x is:

$$z = \frac{x - \mu}{\sigma}$$

The z-score tells you how many standard deviations the value x is above (to the right of) or below (to the left of) the mean, μ . Values of x that are larger than the mean have positive z-scores, and values of x that are smaller than the mean have negative z-scores. If x equals the mean, then x has a z-score of zero.

Example 6.2.1

Suppose $X \sim N(5, 6)$. This says that X is a normally distributed random variable with mean $\mu = 5$ and standard deviation $\sigma = 6$. Suppose $x = 17$. Then:

$$z = \frac{x - \mu}{\sigma} = \frac{17 - 5}{6} = 2$$

This means that $x = 17$ is **two standard deviations** (2σ) above or to the right of the mean $\mu = 5$.

Now suppose $x = 1$. Then: $z = \frac{x - \mu}{\sigma} = \frac{1 - 5}{6} = -0.67$ (rounded to two decimal places)

This means that $x = 1$ is 0.67 standard deviations (-0.67σ) below or to the left of the mean $\mu = 5$.

The Empirical Rule

If X is a random variable and has a normal distribution with mean μ and standard deviation σ , then **the Empirical Rule** states the following:

- About 68% of the x values lie between -1σ and $+1\sigma$ of the mean μ (within one standard deviation of the mean).
- About 95% of the x values lie between -2σ and $+2\sigma$ of the mean μ (within two standard deviations of the mean).
- About 99.7% of the x values lie between -3σ and $+3\sigma$ of the mean μ (within three standard deviations of the mean). Notice that almost all the x values lie within three standard deviations of the mean.
- The z-scores for $+1\sigma$ and -1σ are $+1$ and -1 , respectively.
- The z-scores for $+2\sigma$ and -2σ are $+2$ and -2 , respectively.
- The z-scores for $+3\sigma$ and -3σ are $+3$ and -3 respectively.




Figure 6.2.1

Example 6.2.1

Suppose x has a normal distribution with mean 50 and standard deviation 6.

- About 68% of the x values lie within one standard deviation of the mean. Therefore, about 68% of the x values lie between $-1\sigma = (-1)(6) = -6$ and $1\sigma = (1)(6) = 6$ of the mean 50. The values $50 - 6 = 44$ and $50 + 6 = 56$ are within one standard deviation from the mean 50. The z-scores are -1 and $+1$ for 44 and 56, respectively.
- About 95% of the x values lie within two standard deviations of the mean. Therefore, about 95% of the x values lie between $-2\sigma = (-2)(6) = -12$ and $2\sigma = (2)(6) = 12$. The values $50 - 12 = 38$ and $50 + 12 = 62$ are within two standard deviations from the mean 50. The z-scores are -2 and $+2$ for 38 and 62, respectively.
- About 99.7% of the x values lie within three standard deviations of the mean. Therefore, about 99.7% of the x values lie between $-3\sigma = (-3)(6) = -18$ and $3\sigma = (3)(6) = 18$ of the mean 50. The values $50 - 18 = 32$ and $50 + 18 = 68$ are within three standard deviations from the mean 50. The z-scores are -3 and $+3$ for 32 and 68, respectively.

This page titled [6.2: The Standard Normal Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **6.1: The Standard Normal Distribution** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.3: Using the Normal Distribution

The shaded area in the following graph indicates the area to the right of x . This area is represented by the probability $P(X > x)$. Normal tables provide the probability between the mean, zero for the standard normal distribution, and a specific value such as x_1 . This is the unshaded part of the graph from the mean to x_1 .




Figure 6.3.1

Because the normal distribution is symmetrical, if x_1 were the same distance to the left of the mean the area, probability, in the left tail, would be the same as the shaded area in the right tail. Also, bear in mind that because of the symmetry of this distribution, one-half of the probability is to the right of the mean and one-half is to the left of the mean.

Calculations of Probabilities

To find the probability for probability density functions with a continuous random variable we need to calculate the area under the function across the values of X we are interested in. For the normal distribution this seems a difficult task given the complexity of the formula. There is, however, a simply way to get what we want. Here again is the formula for the normal distribution:

$$f(x) = \frac{1}{\sigma \cdot \sqrt{2 \cdot \pi}} \cdot e^{-\frac{1}{2} \cdot \left(\frac{x-\mu}{\sigma}\right)^2}$$

Looking at the formula for the normal distribution it is not clear just how we are going to solve for the probability doing it the same way we did it with the previous probability functions. There we put the data into the formula and did the math.

To solve this puzzle we start knowing that the area under a probability density function is the probability.




Figure 6.3.2

This shows that the area between X_1 and X_2 is the probability as stated in the formula: $P(X_1 \leq X \leq X_2)$

The mathematical tool needed to find the area under a curve is integral calculus. The integral of the normal probability density function between the two points x_1 and x_2 is the area under the curve between these two points and is the probability between these two points.

Doing these integrals is no fun and can be very time consuming. But now, remembering that there are an infinite number of normal distributions out there, we can consider the one with a mean of zero and a standard deviation of 1. This particular normal distribution is given the name Standard Normal Distribution. Putting these values into the formula it reduces to a very simple equation. We can now quite easily calculate all probabilities for any value of x , for this particular normal distribution, that has a mean of zero and a standard deviation of 1. These have been produced and are available here in the appendix to the text or everywhere on the web. They are presented in various ways. The table in this text is the most common presentation and is set up with probabilities for one-half the distribution beginning with zero, the mean, and moving outward. The shaded area in the graph at the top of the table in [Statistical Tables](#) represents the probability from zero to the specific Z value noted on the horizontal axis, Z .

The only problem is that even with this table, it would be a ridiculous coincidence that our data had a mean of zero and a standard deviation of one. The solution is to convert the distribution we have with its mean and standard deviation to this new Standard Normal Distribution. The Standard Normal has a random variable called Z .

Using the standard normal table, typically called the normal table, to find the probability of one standard deviation, go to the Z column, reading down to 1.0 and then read at column 0. That number, 0.3413 is the probability from zero to 1 standard deviation. At the top of the table is the shaded area in the distribution which is the probability for one standard deviation. The table has solved our integral calculus problem. But only if our data has a mean of zero and a standard deviation of 1.

However, the essential point here is, the probability for one standard deviation on one normal distribution is the same on every normal distribution. If the population data set has a mean of 10 and a standard deviation of 5 then the probability from 10 to 15, one standard deviation, is the same as from zero to 1,

one standard deviation on the standard normal distribution. To compute probabilities, areas, for any normal distribution, we need only to convert the particular normal distribution to the standard normal distribution and look up the answer in the tables. As review, here again is the **standardizing formula**:

$$Z = \frac{x - \mu}{\sigma}$$

where Z is the value on the standard normal distribution, X is the value from a normal distribution one wishes to convert to the standard normal, μ and σ are, respectively, the mean and standard deviation of that population. Note that the equation uses μ and σ which denotes population parameters. This is still dealing with probability so we always are dealing with the population, with **known** parameter values and a **known** distribution. It is also important to note that because the normal distribution is symmetrical it does not matter if the z-score is positive or negative when calculating a probability. One standard deviation to the left (negative Z-score) covers the same area as one standard deviation to the right (positive Z-score). This fact is why the Standard Normal tables do not provide areas for the left side of the distribution. Because of this symmetry, the Z-score formula is sometimes written as:

$$Z = \frac{|x - \mu|}{\sigma}$$

Where the vertical lines in the equation means the absolute value of the number.

What the standardizing formula is really doing is computing the number of standard deviations X is from the mean of its own distribution. The standardizing formula and the concept of counting standard deviations from the mean is the secret of all that we will do in this statistics class. The reason this is true is that **all** of statistics boils down to variation, and the counting of standard deviations is a measure of variation.

This formula, in many disguises, will reappear over and over throughout this course.

Example 6.3.1

The final exam scores in a statistics class were normally distributed with a mean of 63 and a standard deviation of five.

- Find the probability that a randomly selected student scored more than 65 on the exam.
- Find the probability that a randomly selected student scored less than 85.

Answer a

Let X = a score on the final exam. $X \sim N(63, 5)$, where $\mu = 63$ and $\sigma = 5$.

Draw a graph.

Then, find $P(x > 65)$.

$$P(x > 65) = 0.3446$$




Figure 6.3.3

$$Z_1 = \frac{x_1 - \mu}{\sigma} = \frac{65 - 63}{5} = 0.4$$

$$P(x \geq x_1) = P(Z \geq Z_1) = 0.3446$$

The probability that any student selected at random scores more than 65 is 0.3446. Here is how we found this answer.

Answer b

The normal table provides probabilities from zero to the value Z_1 . For this problem the question can be written as $P(X \geq 65) = P(Z \geq Z_1)$, which is the area in the tail. To find this area the formula would be $0.5 - P(X \leq 65)$. One half of the probability is above the mean value because this is a symmetrical distribution. The graph shows how to find the area in the tail by subtracting that portion from the mean, zero, to the Z_1 value. The final answer is: $P(X \geq 63) = P(Z \geq 0.4) = 0.3446$

$$z = \frac{65 - 63}{5} = 0.4$$

Area to the left of Z_1 to the mean of zero is 0.1554

$$P(x > 65) = P(z > 0.4) = 0.5 - 0.1554 = 0.3446$$

$Z = \frac{x - \mu}{\sigma} = \frac{85 - 63}{5} = 4.4$ which is larger than the maximum value on the Standard Normal Table. Therefore, the probability that one student scores less than 85 is approximately one or 100%.

A score of 85 is 4.4 standard deviations from the mean of 63 which is beyond the range of the standard normal table. Therefore, the probability that one student scores less than 85 is approximately one (or 100%).

Exercise 6.3.1

The golf scores for a school team were normally distributed with a mean of 68 and a standard deviation of three. Find the probability that a randomly selected golfer scored less than 65.

Example 6.3.2A

A personal computer is used for office work at home, research, communication, personal finances, education, entertainment, social networking, and a myriad of other things. Suppose that the average number of hours a household personal computer is used for entertainment is two hours per day. Assume the times for entertainment are normally distributed and the standard deviation for the times is half an hour.

a. Find the probability that a household personal computer is used for entertainment between 1.8 and 2.75 hours per day.

Answer

a. Let X = the amount of time (in hours) a household personal computer is used for entertainment. $X \sim N(2, 0.5)$ where $\mu = 2$ and $\sigma = 0.5$.

Find $P(1.8 < X < 2.75)$.

The probability for which you are looking is the area **between** $X = 1.8$ and $X = 2.75$. $P(1.8 < X < 2.75) = 0.5886$




Figure 6.3.4

$$P(1.8 \leq X \leq 2.75) = P(Z_1 \leq Z \leq Z_2)$$

The probability that a household personal computer is used between 1.8 and 2.75 hours per day for entertainment is 0.5886.

Example 6.3.2B

b. Find the maximum number of hours per day that the bottom quartile of households uses a personal computer for entertainment.

Answer

Solution 6.4

b. To find the maximum number of hours per day that the bottom quartile of households uses a personal computer for entertainment, **find the 25th percentile, k** , where $P(x < k) = 0.25$.




Figure 6.3.5

$f(Z) = 0.5 - 0.25 = 0.25$, therefore $Z \approx -0.675$ (or just 0.67 using the table) $Z = \frac{x-\mu}{\sigma} = \frac{x-2}{0.5} = -0.675$, therefore $x = -0.675 * 0.5 + 2 = 1.66$

The maximum number of hours per day that the bottom quartile of households uses a personal computer for entertainment is 1.66 hours.

Exercise 6.3.2

The golf scores for a school team were normally distributed with a mean of 68 and a standard deviation of three. Find the probability that a golfer scored between 66 and 70.

Example 6.3.3

In the United States the ages 13 to 55+ of smartphone users approximately follow a normal distribution with approximate mean and standard deviation of 36.9 years and 13.9 years, respectively.

a. Determine the probability that a random smartphone user in the age range 13 to 55+ is between 23 and 64.7 years old.

Answer

Answer

- a. 0.8186
- b. 0.8413

Example 6.3.4

A citrus farmer who grows mandarin oranges finds that the diameters of mandarin oranges harvested on his farm follow a normal distribution with a mean diameter of 5.85 cm and a standard deviation of 0.24 cm.

a. Find the probability that a randomly selected mandarin orange from this farm has a diameter larger than 6.0 cm. Sketch the graph.

Answer




Figure 6.3.6

$$Z_1 = \frac{6 - 5.85}{.24} = .625$$

$$P(x \geq 6) = P(z \geq 0.625) = 0.2670$$

b. The middle 20% of mandarin oranges from this farm have diameters between _____ and _____.

$$f(Z) = \frac{0.20}{2} = 0.10, \text{ therefore } Z \approx \pm 0.25$$

$$Z = \frac{x - \mu}{\sigma} = \frac{x - 5.85}{0.24} = \pm 0.25 \rightarrow \pm 0.25 \cdot 0.24 + 5.85 = (5.79, 5.91)$$

This page titled [6.3: Using the Normal Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [6.2: Using the Normal Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.4: Estimating the Binomial with the Normal Distribution

We found earlier that various probability density functions are the limiting distributions of others; thus, we can estimate one with another under certain circumstances. We will find here that the normal distribution can be used to estimate a binomial process. The Poisson was used to estimate the binomial previously, and the binomial was used to estimate the hypergeometric distribution.

In the case of the relationship between the hypergeometric distribution and the binomial, we had to recognize that a binomial process assumes that the probability of a success remains constant from trial to trial: a head on the last flip cannot have an effect on the probability of a head on the next flip. In the hypergeometric distribution this is the essence of the question because the experiment assumes that any "draw" is without replacement. If one draws without replacement, then all subsequent "draws" are conditional probabilities. We found that if the hypergeometric experiment draws only a small percentage of the total objects, then we can ignore the impact on the probability from draw to draw.

Imagine that there are 312 cards in a deck comprised of 6 normal decks. If the experiment called for drawing only 10 cards, less than 5% of the total, then we will accept the binomial estimate of the probability, even though this is actually a hypergeometric distribution because the cards are presumably drawn without replacement.

The Poisson likewise was considered an appropriate estimate of the binomial under certain circumstances. In Figure 6.4.11 shows a symmetrical normal distribution transposed on a graph of a binomial distribution where $p = 0.2$ and $n = 5$. The discrepancy between the estimated probability using a normal distribution and the probability of the original binomial distribution is apparent. The criteria for using a normal distribution to estimate a binomial thus addresses this problem by requiring BOTH np AND $n(1 - p)$ are greater than five. Again, this is a rule of thumb, but is effective and results in acceptable estimates of the binomial probability.




Figure 6.4.11

$$1 - [p(X = 0) + p(X = 1) + p(X = 2) + \dots + p(X = 16)] = p(X > 16) = p(Z > 2) = 0.0228$$

This page titled [6.4: Estimating the Binomial with the Normal Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [6.3: Estimating the Binomial with the Normal Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.5: Chapter Key Items

Normal Distribution

a continuous random variable (RV) with pdf $f(x) =$

$$\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

, where μ is the mean of the distribution and σ is the standard deviation; notation: $X \sim N(\mu, \sigma)$. If $\mu = 0$ and $\sigma = 1$, the RV , Z , is called the **standard normal distribution**.

Standard Normal Distribution

a continuous random variable (RV) $X \sim N(0, 1)$; when X follows the standard normal distribution, it is often noted as $Z \sim N(0, 1)$.

z-score

the linear transformation of the form $z = \frac{x-\mu}{\sigma}$ or written as $z = \frac{|x-\mu|}{\sigma}$; if this transformation is applied to any normal distribution $X \sim N(\mu, \sigma)$ the result is the standard normal distribution $Z \sim N(0, 1)$. If this transformation is applied to any specific value x of the RV with mean μ and standard deviation σ , the result is called the z-score of x . The z-score allows us to compare data that are normally distributed but scaled differently. A z-score is the number of standard deviations a particular x is away from its mean value.

This page titled [6.5: Chapter Key Items](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **6.6: Chapter Key Items** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.6: Chapter Review

6.1 The Standard Normal Distribution

A z-score is a standardized value. Its distribution is the standard normal, $Z \sim N(0, 1)$. The mean of the z-scores is zero and the standard deviation is one. If z is the z-score for a value x from the normal distribution $N(\mu, \sigma)$ then z tells you how many standard deviations x is above (greater than) or below (less than) μ .

6.3 Estimating the Binomial with the Normal Distribution

The normal distribution, which is continuous, is the most important of all the probability distributions. Its graph is bell-shaped. This bell-shaped curve is used in almost all disciplines. Since it is a continuous distribution, the total area under the curve is one. The parameters of the normal are the mean μ and the standard deviation σ . A special normal distribution, called the standard normal distribution is the distribution of z-scores. Its mean is zero, and its standard deviation is one.

This page titled [6.6: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [6.9: Chapter Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.7: Chapter Formula Review

Introduction

$$X \sim N(\mu, \sigma)$$

μ = the mean; σ = the standard deviation

The Standard Normal Distribution

$$Z \sim N(0, 1)$$

$z = a$ standardized value (z-score)

mean = 0; standard deviation = 1

To find the k^{th} percentile of X when the z-scores is known:

$$k = \mu + (z)\sigma$$

$$\text{z-score: } z = \frac{x - \mu}{\sigma} \text{ or } z = \frac{|x - \mu|}{\sigma}$$

Z = the random variable for z-scores

$$Z \sim N(0, 1)$$

Estimating the Binomial with the Normal Distribution

Normal Distribution: $X \sim N(\mu, \sigma)$ where μ is the mean and σ is the standard deviation.

Standard Normal Distribution: $Z \sim N(0, 1)$.

This page titled [6.7: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [6.4: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.8: Chapter Practice

6.1 The Standard Normal Distribution

1.

A bottle of water contains 12.05 fluid ounces with a standard deviation of 0.01 ounces. Define the random variable X in words. $X =$ _____.

2.

A normal distribution has a mean of 61 and a standard deviation of 15. What is the median?

3.

$$X \sim N(1, 2)$$

$$\sigma = \underline{\hspace{2cm}}$$

4.

A company manufactures rubber balls. The mean diameter of a ball is 12 cm with a standard deviation of 0.2 cm. Define the random variable X in words. $X =$ _____.

5.

$$X \sim N(-4, 1)$$

What is the median?

6.

$$X \sim N(3, 5)$$

$$\sigma = \underline{\hspace{2cm}}$$

7.

$$X \sim N(-2, 1)$$

$$\mu = \underline{\hspace{2cm}}$$

8.

What does a z-score measure?

9.

What does standardizing a normal distribution do to the mean?

10.

Is $X \sim N(0, 1)$ a standardized normal distribution? Why or why not?

11.

What is the z-score of $x = 12$, if it is two standard deviations to the right of the mean?

12.

What is the z-score of $x = 9$, if it is 1.5 standard deviations to the left of the mean?

13.

What is the z-score of $x = -2$, if it is 2.78 standard deviations to the right of the mean?

14.

What is the z-score of $x = 7$, if it is 0.133 standard deviations to the left of the mean?

15.

Suppose $X \sim N(2, 6)$. What value of x has a z-score of three?

16.

Suppose $X \sim N(8, 1)$. What value of x has a z-score of -2.25 ?

17.

Suppose $X \sim N(9, 5)$. What value of x has a z-score of -0.5 ?

18.

Suppose $X \sim N(2, 3)$. What value of x has a z-score of -0.67 ?

19.

Suppose $X \sim N(4, 2)$. What value of x is 1.5 standard deviations to the left of the mean?

20.

Suppose $X \sim N(4, 2)$. What value of x is two standard deviations to the right of the mean?

21.

Suppose $X \sim N(8, 9)$. What value of x is 0.67 standard deviations to the left of the mean?

22.

Suppose $X \sim N(-1, 2)$. What is the z-score of $x = 2$?

23.

Suppose $X \sim N(12, 6)$. What is the z-score of $x = 2$?

24.

Suppose $X \sim N(9, 3)$. What is the z-score of $x = 9$?

25.

Suppose a normal distribution has a mean of six and a standard deviation of 1.5. What is the z-score of $x = 5.5$?

26.

In a normal distribution, $x = 5$ and $z = -1.25$. This tells you that $x = 5$ is ____ standard deviations to the ____ (right or left) of the mean.

27.

In a normal distribution, $x = 3$ and $z = 0.67$. This tells you that $x = 3$ is ____ standard deviations to the ____ (right or left) of the mean.

28.

In a normal distribution, $x = -2$ and $z = 6$. This tells you that $x = -2$ is ____ standard deviations to the ____ (right or left) of the mean.

29.

In a normal distribution, $x = -5$ and $z = -3.14$. This tells you that $x = -5$ is ____ standard deviations to the ____ (right or left) of the mean.

30.

In a normal distribution, $x = 6$ and $z = -1.7$. This tells you that $x = 6$ is ____ standard deviations to the ____ (right or left) of the mean.

31.

About what percent of x values from a normal distribution lie within one standard deviation (left and right) of the mean of that distribution?

32.

About what percent of the x values from a normal distribution lie within two standard deviations (left and right) of the mean of that distribution?

33.

About what percent of x values lie between the second and third standard deviations (both sides)?

34.

Suppose $X \sim N(15, 3)$. Between what x values does 68.27% of the data lie? The range of x values is centered at the mean of the distribution (i.e., 15).

35.

Suppose $X \sim N(-3, 1)$. Between what x values does 95.45% of the data lie? The range of x values is centered at the mean of the distribution (i.e., -3).

36.

Suppose $X \sim N(-3, 1)$. Between what x values does 34.14% of the data lie?

37.

About what percent of x values lie between the mean and three standard deviations?

38.

About what percent of x values lie between the mean and one standard deviation?

39.

About what percent of x values lie between the first and second standard deviations from the mean (both sides)?

40.

About what percent of x values lie between the first and third standard deviations (both sides)?

Use the following information to answer the next two exercises: The life of Sunshine CD players is normally distributed with mean of 4.1 years and a standard deviation of 1.3 years. A CD player is guaranteed for three years. We are interested in the length of time a CD player lasts.

41.

Define the random variable X in words. $X =$ _____.

42.

$X \sim$ _____(_____, _____)

6.3 Estimating the Binomial with the Normal Distribution

43.

How would you represent the area to the left of one in a probability statement?




Figure 6.8.13

44.

What is the area to the right of one?




Figure 6.8.14

45.

Is $P(x < 1)$ equal to $P(x \leq 1)$? Why?

46.

How would you represent the area to the left of three in a probability statement?




Figure 6.8.15

47.

What is the area to the right of three?




Figure 6.8.16

48.

If the area to the left of x in a normal distribution is 0.123, what is the area to the right of x ?

49.

If the area to the right of x in a normal distribution is 0.543, what is the area to the left of x ?

Use the following information to answer the next four exercises:

$$X \sim N(54, 8)$$

50.

Find the probability that $x > 56$.

51.

Find the probability that $x < 30$.

52.

$$X \sim N(6, 2)$$

Find the probability that x is between three and nine.

53.

$$X \sim N(-3, 4)$$

Find the probability that x is between one and four.

54.

$$X \sim N(4, 5)$$

Find the maximum of x in the bottom quartile.

55.

Use the following information to answer the next three exercise: The life of Sunshine CD players is normally distributed with a mean of 4.1 years and a standard deviation of 1.3 years. A CD player is guaranteed for three years. We are interested in the length of time a CD player lasts. Find the probability that a CD player will break down during the guarantee period.

1. Sketch the situation. Label and scale the axes. Shade the region corresponding to the probability.



Figure 6.8.17

2. $P(0 < x < \underline{\hspace{2cm}}) = \underline{\hspace{2cm}}$ (Use zero for the minimum value of x .)

56.

Find the probability that a CD player will last between 2.8 and six years.

1. Sketch the situation. Label and scale the axes. Shade the region corresponding to the probability.



Figure 6.8.18

2. $P(\underline{\hspace{2cm}} < x < \underline{\hspace{2cm}}) = \underline{\hspace{2cm}}$

57.

An experiment with a probability of success given as 0.40 is repeated 100 times. Use the normal distribution to approximate the binomial distribution, and find the probability the experiment will have at least 45 successes.

58.

An experiment with a probability of success given as 0.30 is repeated 90 times. Use the normal distribution to approximate the binomial distribution, and find the probability the experiment will have at least 22 successes.

59.

An experiment with a probability of success given as 0.40 is repeated 100 times. Use the normal distribution to approximate the binomial distribution, and find the probability the experiment will have from 35 to 45 successes.

60.

An experiment with a probability of success given as 0.30 is repeated 90 times. Use the normal distribution to approximate the binomial distribution, and find the probability the experiment will have from 26 to 30 successes.

61.

An experiment with a probability of success given as 0.40 is repeated 100 times. Use the normal distribution to approximate the binomial distribution, and find the probability the experiment will have at most 34 successes.

62.

An experiment with a probability of success given as 0.30 is repeated 90 times. Use the normal distribution to approximate the binomial distribution, and find the probability the experiment will have at most 34 successes.

63.

A multiple choice test has a probability any question will be guesses correctly of 0.25. There are 100 questions, and a student guesses at all of them. Use the normal distribution to approximate the binomial distribution, and determine the probability at least 30, but no more than 32, questions will be guessed correctly.

64.

A multiple choice test has a probability any question will be guesses correctly of 0.25. There are 100 questions, and a student guesses at all of them. Use the normal distribution to approximate the binomial distribution, and determine the probability at least 24, but no more than 28, questions will be guessed correctly.

This page titled [6.8: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [6.7: Chapter Practice](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.9: Chapter Homework

6.1 The Standard Normal Distribution

Use the following information to answer the next two exercises: The patient recovery time from a particular surgical procedure is normally distributed with a mean of 5.3 days and a standard deviation of 2.1 days.

65.

What is the median recovery time?

- a. 2.7
- b. 5.3
- c. 7.4
- d. 2.1

66.

What is the z-score for a patient who takes ten days to recover?

- a. 1.5
- b. 0.2
- c. 2.2
- d. 7.3

67.

The length of time to find a parking space at 9 A.M. follows a normal distribution with a mean of five minutes and a standard deviation of two minutes. If the mean is significantly greater than the standard deviation, which of the following statements is true?

- I. The data cannot follow the uniform distribution.
- II. The data cannot follow the exponential distribution..
- III. The data cannot follow the normal distribution.

- a. I only
- b. II only
- c. III only
- d. I, II, and III

68.

The heights of the 430 National Basketball Association players were listed on team rosters at the start of the 2005–2006 season. The heights of basketball players have an approximate normal distribution with mean, $\mu = 79$ inches and a standard deviation, $\sigma = 3.89$ inches. For each of the following heights, calculate the z-score and interpret it using complete sentences.

- 1. 77 inches
- 2. 85 inches
- 3. If an NBA player reported his height had a z-score of 3.5, would you believe him? Explain your answer.

69.

The systolic blood pressure (given in millimeters) of males has an approximately normal distribution with mean $\mu = 125$ and standard deviation $\sigma = 14$. Systolic blood pressure for males follows a normal distribution.

- 1. Calculate the z-scores for the male systolic blood pressures 100 and 150 millimeters.
- 2. If a male friend of yours said he thought his systolic blood pressure was 2.5 standard deviations below the mean, but that he believed his blood pressure was between 100 and 150 millimeters, what would you say to him?

70.

Kyle's doctor told him that the z-score for his systolic blood pressure is 1.75. Which of the following is the best interpretation of this standardized score? The systolic blood pressure (given in millimeters) of males has an approximately normal distribution with mean $\mu = 125$ and standard deviation $\sigma = 14$. If $X =$ a systolic blood pressure score then $X \sim N(125, 14)$.

1. Which answer(s) is/are correct?
 - o Kyle's systolic blood pressure is 175.
 - o Kyle's systolic blood pressure is 1.75 times the average blood pressure of men his age.
 - o Kyle's systolic blood pressure is 1.75 above the average systolic blood pressure of men his age.
 - o Kyles's systolic blood pressure is 1.75 standard deviations above the average systolic blood pressure for men.
2. Calculate Kyle's blood pressure.

71.

Height and weight are two measurements used to track a child's development. The World Health Organization measures child development by comparing the weights of children who are the same height and the same gender. In 2009, weights for all 80 cm girls in the reference population had a mean $\mu = 10.2$ kg and standard deviation $\sigma = 0.8$ kg. Weights are normally distributed. $X \sim N(10.2, 0.8)$. Calculate the z-scores that correspond to the following weights and interpret them.

1. 11 kg
2. 7.9 kg
3. 12.2 kg

72.

In 2005, 1,475,623 students heading to college took the SAT. The distribution of scores in the math section of the SAT follows a normal distribution with mean $\mu = 520$ and standard deviation $\sigma = 115$.

1. Calculate the z-score for an SAT score of 720. Interpret it using a complete sentence.
2. What math SAT score is 1.5 standard deviations above the mean? What can you say about this SAT score?
3. For 2012, the SAT math test had a mean of 514 and standard deviation 117. The ACT math test is an alternate to the SAT and is approximately normally distributed with mean 21 and standard deviation 5.3. If one person took the SAT math test and scored 700 and a second person took the ACT math test and scored 30, who did better with respect to the test they took?

6.3 Estimating the Binomial with the Normal Distribution

Use the following information to answer the next two exercises: The patient recovery time from a particular surgical procedure is normally distributed with a mean of 5.3 days and a standard deviation of 2.1 days.

73.

What is the probability of spending more than two days in recovery?

- a. 0.0580
- b. 0.8447
- c. 0.0553
- d. 0.9420

Use the following information to answer the next three exercises: The length of time it takes to find a parking space at 9 A.M. follows a normal distribution with a mean of five minutes and a standard deviation of two minutes.

74.

Based upon the given information and numerically justified, would you be surprised if it took less than one minute to find a parking space?

- a. Yes
- b. No
- c. Unable to determine

75.

Find the probability that it takes at least eight minutes to find a parking space.

- a. 0.0001
- b. 0.9270
- c. 0.1862
- d. 0.0668

76.

Seventy percent of the time, it takes more than how many minutes to find a parking space?

- 1.24
- 2.41
- 3.95
- 6.05

77.

According to a study done by De Anza students, the height for Asian adult males is normally distributed with an average of 66 inches and a standard deviation of 2.5 inches. Suppose one Asian adult male is randomly chosen. Let X = height of the individual.

1. $X \sim \text{_____}(\text{_____,} \text{_____})$
2. Find the probability that the person is between 65 and 69 inches. Include a sketch of the graph, and write a probability statement.
3. Would you expect to meet many Asian adult males over 72 inches? Explain why or why not, and justify your answer numerically.
4. The middle 40% of heights fall between what two values? Sketch the graph, and write the probability statement.

78.

IQ is normally distributed with a mean of 100 and a standard deviation of 15. Suppose one individual is randomly chosen. Let X = IQ of an individual.

1. $X \sim \text{_____}(\text{_____,} \text{_____})$
2. Find the probability that the person has an IQ greater than 120. Include a sketch of the graph, and write a probability statement.
3. MENSA is an organization whose members have the top 2% of all IQs. Find the minimum IQ needed to qualify for the MENSA organization. Sketch the graph, and write the probability statement.

79.

The percent of fat calories that a person in America consumes each day is normally distributed with a mean of about 36 and a standard deviation of 10. Suppose that one individual is randomly chosen. Let X = percent of fat calories.

1. $X \sim \text{_____}(\text{_____,} \text{_____})$
2. Find the probability that the percent of fat calories a person consumes is more than 40. Graph the situation. Shade in the area to be determined.
3. Find the maximum number for the lower quarter of percent of fat calories. Sketch the graph and write the probability statement.

80.

Suppose that the distance of fly balls hit to the outfield (in baseball) is normally distributed with a mean of 250 feet and a standard deviation of 50 feet.

1. If X = distance in feet for a fly ball, then $X \sim \text{_____}(\text{_____,} \text{_____})$
2. If one fly ball is randomly chosen from this distribution, what is the probability that this ball traveled fewer than 220 feet? Sketch the graph. Scale the horizontal axis X . Shade the region corresponding to the probability. Find the probability.

81.

In China, four-year-olds average three hours a day unsupervised. Most of the unsupervised children live in rural areas, considered safe. Suppose that the standard deviation is 1.5 hours and the amount of time spent alone is normally distributed. We randomly select one Chinese four-year-old living in a rural area. We are interested in the amount of time the child spends alone per day.

1. In words, define the random variable X .
2. $X \sim \text{_____}(\text{_____,} \text{_____})$
3. Find the probability that the child spends less than one hour per day unsupervised. Sketch the graph, and write the probability statement.
4. What percent of the children spend over ten hours per day unsupervised?
5. Seventy percent of the children spend at least how long per day unsupervised?

82.

In the 1992 presidential election, Alaska's 40 election districts averaged 1,956.8 votes per district for President Clinton. The standard deviation was 572.3. (There are only 40 election districts in Alaska.) The distribution of the votes per district for President Clinton was bell-shaped. Let X = number of votes for President Clinton for an election district.

1. State the approximate distribution of X .
2. Is 1,956.8 a population mean or a sample mean? How do you know?
3. Find the probability that a randomly selected district had fewer than 1,600 votes for President Clinton. Sketch the graph and write the probability statement.
4. Find the probability that a randomly selected district had between 1,800 and 2,000 votes for President Clinton.
5. Find the third quartile for votes for President Clinton.

83.

Suppose that the duration of a particular type of criminal trial is known to be normally distributed with a mean of 21 days and a standard deviation of seven days.

1. In words, define the random variable X .
2. $X \sim \text{_____}(\text{_____,} \text{_____})$
3. If one of the trials is randomly chosen, find the probability that it lasted at least 24 days. Sketch the graph and write the probability statement.
4. Sixty percent of all trials of this type are completed within how many days?

84.

Terri Vogel, an amateur motorcycle racer, averages 129.71 seconds per 2.5 mile lap (in a seven-lap race) with a standard deviation of 2.28 seconds. The distribution of her race times is normally distributed. We are interested in one of her randomly selected laps.

1. In words, define the random variable X .
2. $X \sim \text{_____}(\text{_____,} \text{_____})$
3. Find the percent of her laps that are completed in less than 130 seconds.
4. The fastest 3% of her laps are under _____.
5. The middle 80% of her laps are from _____ seconds to _____ seconds.

85.

Thuy Dau, Ngoc Bui, Sam Su, and Lan Young conducted a survey as to how long customers at Lucky claimed to wait in the checkout line until their turn. Let X = time in line. [Table 6.9.1](#) displays the ordered real data (in minutes):

Table 6.9.1

0.50	4.25	5	6	7.25
1.75	4.25	5.25	6	7.25
2	4.25	5.25	6.25	7.25
2.25	4.25	5.5	6.25	7.75
2.25	4.5	5.5	6.5	8
2.5	4.75	5.5	6.5	8.25
2.75	4.75	5.75	6.5	9.5
3.25	4.75	5.75	6.75	9.5
3.75	5	6	6.75	9.75
3.75	5	6	6.75	10.75

1. Calculate the sample mean and the sample standard deviation.
2. Construct a histogram.
3. Draw a smooth curve through the midpoints of the tops of the bars.
4. In words, describe the shape of your histogram and smooth curve.

5. Let the sample mean approximate μ and the sample standard deviation approximate σ . The distribution of X can then be approximated by $X \sim \text{_____}(\text{_____,} \text{_____})$
6. Use the distribution in part e to calculate the probability that a person will wait fewer than 6.1 minutes.
7. Determine the cumulative relative frequency for waiting less than 6.1 minutes.
8. Why aren't the answers to part 6 and part 7 exactly the same?
9. Why are the answers to part 6 and part 7 as close as they are?
10. If only ten customers has been surveyed rather than 50, do you think the answers to part f and part g would have been closer together or farther apart? Explain your conclusion.

86.

Suppose that Ricardo and Anita attend different colleges. Ricardo's GPA is the same as the average GPA at his school. Anita's GPA is 0.70 standard deviations above her school average. In complete sentences, explain why each of the following statements may be false.

1. Ricardo's actual GPA is lower than Anita's actual GPA.
2. Ricardo is not passing because his z-score is zero.
3. Anita is in the 70th percentile of students at her college.

87.

An expert witness for a paternity lawsuit testifies that the length of a pregnancy is normally distributed with a mean of 280 days and a standard deviation of 13 days. An alleged father was out of the country from 240 to 306 days before the birth of the child, so the pregnancy would have been less than 240 days or more than 306 days long if he was the father. The birth was uncomplicated, and the child needed no medical intervention. What is the probability that he was NOT the father? What is the probability that he could be the father? Calculate the z-scores first, and then use those to calculate the probability.

88.

A NUMMI assembly line, which has been operating since 1984, has built an average of 6,000 cars and trucks a week. Generally, 10% of the cars were defective coming off the assembly line. Suppose we draw a random sample of $n = 100$ cars. Let X represent the number of defective cars in the sample. What can we say about X in regard to the 68-95-99.7 empirical rule (one standard deviation, two standard deviations and three standard deviations from the mean are being referred to)? Assume a normal distribution for the defective cars in the sample.

89.

We flip a coin 100 times ($n = 100$) and note that it only comes up heads 20% ($p = 0.20$) of the time. The mean and standard deviation for the number of times the coin lands on heads is $\mu = 20$ and $\sigma = 4$ (verify the mean and standard deviation). Solve the following:

1. There is about a 68% chance that the number of heads will be somewhere between ___ and ___.
2. There is about a ___ chance that the number of heads will be somewhere between 12 and 28.
3. There is about a ___ chance that the number of heads will be somewhere between eight and 32.

90.

A \$1 scratch off lotto ticket will be a winner one out of five times. Out of a shipment of $n = 190$ lotto tickets, find the probability for the lotto tickets that there are

1. somewhere between 34 and 54 prizes.
2. somewhere between 54 and 64 prizes.
3. more than 64 prizes.

91.

Facebook provides a variety of statistics on its Web site that detail the growth and popularity of the site.

On average, 28 percent of 18 to 34 year olds check their Facebook profiles before getting out of bed in the morning. Suppose this percentage follows a normal distribution with a standard deviation of five percent.

92.

A hospital has 49 births in a year. It is considered equally likely that a birth be a boy as it is the birth be a girl.

1. What is the mean?
2. What is the standard deviation?
3. Can this binomial distribution be approximated with a normal distribution?
4. If so, use the normal distribution to find the probability that at least 23 of the 49 births were boys.

93.

Historically, a final exam in a course is passed with a probability of 0.9. The exam is given to a group of 70 students.

1. What is the mean of the binomial distribution?
2. What is the standard deviation?
3. Can this binomial distribution be approximate with a normal distribution?
4. If so, use the normal distribution to find the probability that at least 60 of the students pass the exam?

94.

A tree in an orchard has 200 oranges. Of the oranges, 40 are not ripe. Use the normal distribution to approximate the binomial distribution, and determine the probability a box containing 35 oranges has at most two oranges that are not ripe.

95.

In a large city one in ten fire hydrants are in need of repair. If a crew examines 100 fire hydrants in a week, what is the probability they will find nine or fewer fire hydrants that need repair? Use the normal distribution to approximate the binomial distribution.

96.

On an assembly line it is determined 85% of the assembled products have no defects. If one day 50 items are assembled, what is the probability at least 4 and no more than 8 are defective. Use the normal distribution to approximate the binomial distribution.

This page titled [6.9: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [6.5: Chapter Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.10: Chapter References

The Standard Normal Distribution

- “Blood Pressure of Males and Females.” StatCrunch, 2013. Available online at <http://www.statcrunch.com/5.0/viewre...reportid=11960> (accessed May 14, 2013).
- “The Use of Epidemiological Tools in Conflict-affected populations: Open-access educational resources for policy-makers: Calculation of z-scores.” London School of Hygiene and Tropical Medicine, 2009. Available online at http://conflict.lshtm.ac.uk/page_125.htm (accessed May 14, 2013).
- “2012 College-Bound Seniors Total Group Profile Report.” CollegeBoard, 2012. Available online at <http://media.collegeboard.com/digita...Group-2012.pdf> (accessed May 14, 2013).
- “Digest of Education Statistics: ACT score average and standard deviations by sex and race/ethnicity and percentage of ACT test takers, by selected composite score ranges and planned fields of study: Selected years, 1995 through 2009.” National Center for Education Statistics. Available online at http://nces.ed.gov/programs/digest/d...s/dt09_147.asp (accessed May 14, 2013).
- Data from the *San Jose Mercury News*.
- Data from *The World Almanac and Book of Facts*.
- “List of stadiums by capacity.” Wikipedia. Available online at https://en.Wikipedia.org/wiki/List_o...ms_by_capacity (accessed May 14, 2013).
- Data from the National Basketball Association. Available online at www.nba.com (accessed May 14, 2013).

Using the Normal Distribution

- “Naegele’s rule.” Wikipedia. Available online at http://en.Wikipedia.org/wiki/Naegele's_rule (accessed May 14, 2013).
- “403: NUMMI.” Chicago Public Media & Ira Glass, 2013. Available online at <http://www.thisamericanlife.org/radi...sode/403/nummi> (accessed May 14, 2013).
- “Scratch-Off Lottery Ticket Playing Tips.” WinAtTheLottery.com, 2013. Available online at www.winatthelottery.com/publi...partment40.cfm (accessed May 14, 2013).
- “Smart Phone Users, By The Numbers.” Visual.ly, 2013. Available online at <http://visual.ly/smart-phone-users-numbers> (accessed May 14, 2013).
- “Facebook Statistics.” Statistics Brain. Available online at <http://www.statisticbrain.com/facebo...tics/> (accessed May 14, 2013).

This page titled [6.10: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **6.8: Chapter References** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

6.11: Chapter Solution (Practice + Homework)

1.

ounces of water in a bottle

3.

2

5.

-4

7.

-2

9.

The mean becomes zero.

11.

$$z = 2$$

13.

$$z = 2.78$$

15.

$$x = 20$$

17.

$$x = 6.5$$

19.

$$x = 1$$

21.

$$x = 1.97$$

23.

$$z = -1.67$$

25.

$$z \approx -0.33$$

27.

0.67, right

29.

3.14, left

31.

about 68%

33.

about 4%

35.

between -5 and -1

37.

about 50%

39.

about 27%

41.

The lifetime of a Sunshine CD player measured in years.

43.

$P(x < 1)$

45.

Yes, because they are the same in a continuous distribution: $P(x = 1) = 0$

47.

$1 - P(x < 3)$ or $P(x > 3)$

49.

$1 - 0.543 = 0.457$

51.

0.0013

53.

0.1186

55.

1. 57.

0.154

0.874

59.

0.693

60.

0.346

61.

0.110

62.

0.946

63.

0.071

64.

0.347

66.

c

68.

1. 70.

1. 72.

Let X = an SAT math score and Y = an ACT math score.

1. 75.

d

1. 79.

1. 81.

1. 83.

1. 85.

1. 88.

▪ 90.

▪ 92.

1. 93.

1. 94.

0.02

0.37

96.

0.50

This page titled [6.11: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **6.10: Chapter Solution (Practice + Homework)** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

7: The Central Limit Theorem

- [7.1: Introduction to the Central Limit Theorem](#)
- [7.2: The Central Limit Theorem for Sample Means](#)
- [7.3: Using the Central Limit Theorem](#)
- [7.4: The Central Limit Theorem for Proportions](#)
- [7.5: Finite Population Correction Factor](#)
- [7.6: Chapter Key Terms](#)
- [7.7: Chapter Review](#)
- [7.8: Chapter Formula Review](#)
- [7.9: Chapter Practice](#)
- [7.10: Chapter Homework](#)
- [7.11: Chapter References](#)
- [7.12: Chapter Solution \(Practice + Homework\)](#)

This page titled [7: The Central Limit Theorem](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

7.1: Introduction to the Central Limit Theorem

Why are we so concerned with means? Two reasons are: they give us a middle ground for comparison, and they are easy to calculate. In this chapter, you will study means and the **Central Limit Theorem**.

The **Central Limit Theorem** is one of the most powerful and useful ideas in all of statistics. The Central Limit Theorem is a theorem which means that it is NOT a theory or just somebody's idea of the way things work. As a theorem it ranks with the Pythagorean Theorem, or the theorem that tells us that the sum of the angles of a triangle must add to 180. These are facts of the ways of the world rigorously demonstrated with mathematical precision and logic. As we will see this powerful theorem will determine just what we can, and cannot say, in inferential statistics. The Central Limit Theorem is concerned with drawing finite samples of size n from a population with a known mean, μ , and a known standard deviation, σ . The conclusion is that if we collect samples of size n with a "large enough n ," calculate each sample's mean, and create a histogram (distribution) of those means, then the resulting distribution will tend to have an approximate normal distribution.

The astounding result is that it does not matter what the distribution of the original population is, or whether you even need to know it. The important fact is that the distribution of sample means tend to follow the normal distribution.



Figure 7.1.1 If you want to figure out the distribution of the change people carry in their pockets, using the Central Limit Theorem and assuming your sample is large enough, you will find that the distribution is the normal probability density function. (credit: John Lodder)

The size of the sample, n , that is required in order to be "large enough" depends on the original population from which the samples are drawn (the sample size should be at least 30 or the data should come from a normal distribution). If the original population is far from normal, then more observations are needed for the sample means. **Sampling is done randomly and with replacement in the theoretical model.**

This page titled [7.1: Introduction to the Central Limit Theorem](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.0: Introduction to the Central Limit Theorem](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.2: The Central Limit Theorem for Sample Means

The sampling distribution is a theoretical distribution. It is created by taking many many samples of size n from a population. Each sample mean is then treated like a single observation of this new distribution, the sampling distribution. The genius of thinking this way is that it recognizes that when we sample we are creating an observation and that observation must come from some particular distribution. The Central Limit Theorem answers the question: from what distribution did a sample mean come? If this is discovered, then we can treat a sample mean just like any other observation and calculate probabilities about what values it might take on. We have effectively moved from the world of statistics where we know only what we have from the sample, to the world of probability where we know the distribution from which the sample mean came and the parameters of that distribution.

The reasons that one samples a population are obvious. The time and expense of checking every invoice to determine its validity or every shipment to see if it contains all the items may well exceed the cost of errors in billing or shipping. For some products, sampling would require destroying them, called destructive sampling. One such example is measuring the ability of a metal to withstand saltwater corrosion for parts on ocean going vessels.

Sampling thus raises an important question; just which sample was drawn. Even if the sample were randomly drawn, there are theoretically an almost infinite number of samples. With just 100 items, there are more than 75 million unique samples of size five that can be drawn. If six are in the sample, the number of possible samples increases to just more than one billion. Of the 75 million possible samples, then, which one did you get? If there is variation in the items to be sampled, there will be variation in the samples. One could draw an "unlucky" sample and make very wrong conclusions concerning the population. This recognition that any sample we draw is really only one from a distribution of samples provides us with what is probably the single most important theorem in statistics: **the Central Limit Theorem**. Without the Central Limit Theorem it would be impossible to proceed to inferential statistics from simple probability theory. In its most basic form, the Central Limit Theorem states that **regardless** of the underlying probability density function of the population data, the theoretical distribution of the means of samples from the population will be normally distributed. In essence, this says that the mean of a sample should be treated like an observation drawn from a normal distribution. The Central Limit Theorem only holds if the sample size is "large enough" which has been shown to be only 30 observations or more.

Figure 7.2 graphically displays this very important proposition.




Figure 7.2

Notice that the horizontal axis in the top panel is labeled X . These are the individual observations of the population. This is the **unknown** distribution of the population values. The graph is purposefully drawn all squiggly to show that it does not matter just how odd ball it really is. Remember, we will never know what this distribution looks like, or its mean or standard deviation for that matter.

The horizontal axis in the bottom panel is labeled \bar{X} 's. This is the theoretical distribution called the sampling distribution of the means. Each observation on this distribution is a sample mean. All these sample means were calculated from individual samples with the same sample size. The theoretical sampling distribution contains all of the sample mean values from all the possible samples that could have been taken from the population. Of course, no one would ever actually take all of these samples, but if they did this is how they would look. And the Central Limit Theorem says that they will be normally distributed.

The Central Limit Theorem goes even further and tells us the mean and standard deviation of this theoretical distribution.

Table 7.1

Parameter	Population distribution	Sample	Sampling distribution of \bar{X} 's
Mean	μ	\bar{X}	$\mu_{\bar{x}}$ and $E(\mu_{\bar{x}}) = \mu$
Standard deviation	σ	s	$\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$

The practical significance of The Central Limit Theorem is that now we can compute probabilities for drawing a sample mean, \bar{X} , in just the same way as we did for drawing specific observations, X 's, when we knew the population mean and standard deviation and that the population data were normally distributed.. The standardizing formula has to be amended to recognize that the mean and standard deviation of the sampling distribution, sometimes, called the standard error of the mean, are different from those of the population distribution, but otherwise nothing has changed. The new standardizing formula is

$$Z = \frac{\bar{X} - \mu_{\bar{X}}}{\sigma_{\bar{X}}} = \frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$$

Notice that $\mu_{\bar{X}}$ in the first formula has been changed to simply μ in the second version. The reason is that mathematically it can be shown that the expected value of $\mu_{\bar{X}}$ is equal to μ . This was stated in Table 7.1 above. Mathematically, the $E(x)$ symbol read the "expected value of x ". This formula will be used in the next unit to provide estimates of the **unknown** population parameter μ .

This page titled [7.2: The Central Limit Theorem for Sample Means](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.1: The Central Limit Theorem for Sample Means](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.3: Using the Central Limit Theorem

Examples of the Central Limit Theorem

Law of Large Numbers

The **law of large numbers** says that if you take samples of larger and larger size from any population, then the mean of the sampling distribution, $\mu_{\bar{x}}$ tends to get closer and closer to the true population mean, μ . From the Central Limit Theorem, we know that as n gets larger and larger, the sample means follow a normal distribution. The larger n gets, the smaller the standard deviation of the sampling distribution gets. (Remember that the standard deviation for the sampling distribution of \bar{X} is $\frac{\sigma}{\sqrt{n}}$.) This means that the sample mean \bar{x} must be closer to the population mean μ as n increases. We can say that μ is the value that the sample means approach as n gets larger. The Central Limit Theorem illustrates the law of large numbers.


This concept is so important and plays such a critical role in what follows it deserves to be developed further. Indeed, there are two critical issues that flow from the Central Limit Theorem and the application of the Law of Large numbers to it. These are

1. The probability density function of the sampling distribution of means is normally distributed **regardless** of the underlying distribution of the population observations and
2. standard deviation of the sampling distribution decreases as the size of the samples that were used to calculate the means for the sampling distribution increases.


Taking these in order. It would seem counterintuitive that the population may have **any** distribution and the distribution of means coming from it would be normally distributed. With the use of computers, experiments can be simulated that show the process by which the sampling distribution changes as the sample size is increased. These simulations show visually the results of the mathematical proof of the Central Limit Theorem.

Here are three examples of very different population distributions and the evolution of the sampling distribution to a normal distribution as the sample size increases. The top panel in these cases represents the histogram for the original data. The three panels show the histograms for 1,000 randomly drawn samples for different sample sizes: $n = 10$, $n = 25$ and $n = 50$. As the sample size increases, and the number of samples taken remains constant, the distribution of the 1,000 sample means becomes closer to the smooth line that represents the normal distribution.


Figure 7.3.3 is for a normal distribution of individual observations and we would expect the sampling distribution to converge on the normal quickly. The results show this and show that even at a very small sample size the distribution is close to the normal distribution.



Sample Size $n = 10$



Sample Size $n = 25$



Sample Size $n = 50$






Figure 7.3.3


Figure 7.3.4 is a uniform distribution which, a bit amazingly, quickly approached the normal distribution even with only a sample of 10.



Distribution of Sample means with $n = 10$
Histogram (with Normal Curve) of Mean



Distribution of Sample means with $n = 25$
Histogram (with Normal Curve) of Mean



Distribution of Sample means with $n = 50$
Histogram (with Normal Curve) of Mean






Figure 7.3.4

Figure 7.3.5 is a skewed distribution. This last one could be an exponential, geometric, or binomial with a small probability of success creating the skew in the distribution. For skewed distributions our intuition would say that this will take larger sample sizes to move to a normal distribution and indeed that is what we observe from the simulation. Nevertheless, at a sample size of 50, not considered a very large sample, the distribution of sample means has very decidedly gained the shape of the normal distribution.


Distribution of Random Variable



Sample Size $n = 10$



Sample Size $n = 25$



Sample Size $n = 50$




Figure 7.3.5

The Central Limit Theorem provides more than the proof that the sampling distribution of means is normally distributed. It also provides us with the mean and standard deviation of this distribution. Further, as discussed above, the expected value of the mean, $\mu_{\bar{x}}$, is equal to the mean of the population of the original data which is what we are interested in estimating from the sample we took. We have already inserted this conclusion of the Central Limit Theorem into the formula we use for standardizing from the sampling distribution to the standard normal distribution. And finally, the Central Limit Theorem has also provided the standard deviation of the sampling distribution, $\sigma_{\bar{x}} = \frac{\sigma}{\sqrt{n}}$, and this is critical to have to calculate probabilities of values of the new random variable, \bar{x} .

Figure 7.3.6 shows a sampling distribution. The mean has been marked on the horizontal axis of the \bar{X} 's and the standard deviation has been written to the right above the distribution. Notice that the standard deviation of the sampling distribution is the original standard deviation of the population, divided by the sample size. We have already seen that as the sample size increases the sampling distribution becomes closer and closer to the normal distribution. As this happens, the standard deviation of the sampling distribution changes in another way; the standard deviation decreases as n increases. At very very large n , the standard deviation of the sampling distribution becomes very small and at infinity it collapses on top of the population mean. This is what it means that the expected value of $\mu_{\bar{x}}$ is the population mean, μ .




Figure 7.3.6

At non-extreme values of n , this relationship between the standard deviation of the sampling distribution and the sample size plays a very important part in our ability to estimate the parameters we are interested in.

Figure 7.3.7 shows three sampling distributions. The only change that was made is the sample size that was used to get the sample means for each distribution. As the sample size increases, n goes from 10 to 30 to 50, the standard deviations of the respective sampling distributions decrease because the sample size is in the denominator of the standard deviations of the sampling distributions.




Figure 7.3.7

The implications for this are very important. Figure 7.3.8 shows the effect of the sample size on the confidence we will have in our estimates. These are two sampling distributions from the same population. One sampling distribution was created with samples of size 10 and the other with samples of size 50. All other things constant, the sampling distribution with sample size 50 has a smaller standard deviation that causes the graph to be higher and narrower. The important effect of this is that for the same probability of one standard deviation from the mean, this distribution covers much less of a range of possible values than the other distribution. One standard deviation is marked on the \bar{X} axis for each distribution. This is shown by the two arrows that are plus or minus one standard deviation for each distribution. If the probability that the true mean is one standard deviation away from the mean, then for the sampling distribution with the smaller sample size, the possible range of values is much greater. A simple question is, would you rather have a sample mean from the narrow, tight distribution, or the flat, wide distribution as the estimate of the population mean? Your answer tells us why people intuitively will always choose data from a large sample rather than a small sample. The sample mean they are getting is coming from a more compact distribution. This concept will be the foundation for what will be called level of confidence in the next unit.




Figure 7.3.8

This page titled [7.3: Using the Central Limit Theorem](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.2: Using the Central Limit Theorem](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.4: The Central Limit Theorem for Proportions

The Central Limit Theorem tells us that the point estimate for the sample mean, \bar{x} , comes from a normal distribution of \bar{x} 's. This theoretical distribution is called the sampling distribution of \bar{x} 's. We now investigate the sampling distribution for another important parameter we wish to estimate; p from the binomial probability density function.

If the random variable is discrete, such as for categorical data, then the parameter we wish to estimate is the population proportion. This is, of course, the probability of drawing a success in any one random draw. Unlike the case just discussed for a continuous random variable where we did not know the population distribution of X 's, here we actually know the underlying probability density function for these data; it is the binomial. The random variable is X = the number of successes and the parameter we wish to know is p , the probability of drawing a success which is of course the proportion of successes in the population. The question at issue is: from what distribution was the sample proportion, $p' = \frac{x}{n}$ drawn? The sample size is n and X is the number of successes found in that sample. This is a parallel question that was just answered by the Central Limit Theorem: from what distribution was the sample mean, \bar{x} , drawn? We saw that once we knew that the distribution was the Normal distribution then we were able to create confidence intervals for the population parameter, μ . We will also use this same information to test hypotheses about the population mean later. We wish now to be able to develop confidence intervals for the population parameter " p " from the binomial probability density function.

In order to find the distribution from which sample proportions come we need to develop the sampling distribution of sample proportions just as we did for sample means. So again imagine that we randomly sample say 50 people and ask them if they support the new school bond issue. From this we find a sample proportion, p' , and graph it on the axis of p 's. We do this again and again etc., etc. until we have the theoretical distribution of p 's. Some sample proportions will show high favorability toward the bond issue and others will show low favorability because random sampling will reflect the variation of views within the population. What we have done can be seen in Figure 7.4.9. The top panel is the population distributions of probabilities for each possible value of the random variable X . While we do not know what the specific distribution looks like because we do not know p , the population parameter, we do know that it must look something like this. In reality, we do not know either the mean or the standard deviation of this population distribution, the same difficulty we faced when analyzing the X 's previously.




Figure 7.4.9

Figure 7.4.9 places the mean on the distribution of population probabilities as $\mu = np$ but of course we do not actually know the population mean because we do not know the population probability of success, p . Below the distribution of the population values is the sampling distribution of p 's. Again the Central Limit Theorem tells us that this distribution is normally distributed just like the case of the sampling distribution for \bar{x} 's. This sampling distribution also has a mean, the mean of the p 's, and a standard deviation, $\sigma_{p'}$.

Importantly, in the case of the analysis of the distribution of sample means, the Central Limit Theorem told us the expected value of the mean of the sample means in the sampling distribution, and the standard deviation of the sampling distribution. Again the Central Limit Theorem provides this information for the sampling distribution for proportions. The answers are:

1. The expected value of the mean of sampling distribution of sample proportions, $\mu_{p'}$, is the population proportion, p .
2. The standard deviation of the sampling distribution of sample proportions, $\sigma_{p'}$, is the population standard deviation divided by the square root of the sample size, n .

Both these conclusions are the same as we found for the sampling distribution for sample means. However in this case, because the mean and standard deviation of the binomial distribution both rely upon pp , the formula for the standard deviation of the sampling distribution requires algebraic manipulation to be useful. We will take that up in the next chapter. The proof of these important conclusions from the Central Limit Theorem is provided below.

$$E(p') = E\left(\frac{x}{n}\right) = \left(\frac{1}{n}\right) E(x) = \left(\frac{1}{n}\right) np = p$$

(The expected value of X , $E(x)$, is simply the mean of the binomial distribution which we know to be np .)

$$\sigma_p^2 = \text{Var}(p') = \text{Var}\left(\frac{x}{n}\right) = \frac{1}{n^2} (\text{Var}(x)) = \frac{1}{n^2} (np(1-p)) = \frac{p(1-p)}{n}$$

The standard deviation of the sampling distribution for proportions is thus:

$$\sigma_{p'} = \sqrt{\frac{p(1-p)}{n}}$$

Table 7.4.2

Parameter	Population distribution	Sample	Sampling distribution of p 's
Mean	$\mu = np$	$p' = \frac{x}{n}$	p' and $E(p') = p$
Standard Deviation	$\sigma = \sqrt{npq}$		$\sigma_{p'} = \sqrt{\frac{p(1-p)}{n}}$

Table 7.4.2 summarizes these results and shows the relationship between the population, sample and sampling distribution. Notice the parallel between this Table and Table 7.4.1 for the case where the random variable is continuous and we were developing the sampling distribution for means.

Reviewing the formula for the standard deviation of the sampling distribution for proportions we see that as n increases the standard deviation decreases. This is the same observation we made for the standard deviation for the sampling distribution for means. Again, as the sample size increases, the point estimate for either μ or p is found to come from a distribution with a narrower and narrower distribution. We concluded that with a given level of probability, the range from which the point estimate comes is smaller as the sample size, n , increases. Figure 7.4.8 shows this result for the case of sample means. Simply substitute p' for \bar{x} and we can see the impact of the sample size on the estimate of the sample proportion.

This page titled [7.4: The Central Limit Theorem for Proportions](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.3: The Central Limit Theorem for Proportions](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.5: Finite Population Correction Factor

We saw that the sample size has an important effect on the variance and thus the standard deviation of the sampling distribution. Also of interest is the proportion of the total population that has been sampled. We have assumed that the population is extremely large and that we have sampled a small part of the population. As the population becomes smaller and we sample a larger number of observations the sample observations are not independent of each other. To correct for the impact of this, the Finite Correction Factor can be used to adjust the variance of the sampling distribution. It is appropriate when more than 5% of the population is being sampled and the population has a known population size. There are cases when the population is known, and therefore the correction factor must be applied. The issue arises for both the sampling distribution of the means and the sampling distribution of proportions. The Finite Population Correction Factor for the variance of the means shown in the standardizing formula is:

$$Z = \frac{\bar{x} - \mu}{\frac{\sigma}{\sqrt{n}} \cdot \sqrt{\frac{N-n}{N-1}}}$$

and for the variance of proportions is:

$$\sigma_{p'} = \sqrt{\frac{p(1-p)}{n}} \times \sqrt{\frac{N-n}{N-1}}$$

The following examples show how to apply the factor. Sampling variances get adjusted using the above formula.

Example 7.5.1

It is learned that the population of White German Shepherds in the USA is 4,000 dogs and the mean weight for German Shepherds is 75.45 pounds. It is also learned that the population standard deviation is 10.37 pounds. If the sample size is 100 dogs, then find the probability that a sample will have a mean that differs from the true probability mean by less than 2 pounds.

Answer

Solution 7.1

$$N = 4000, \quad n = 100, \quad \sigma = 10.37, \quad \mu = 75.45, \quad (\bar{x} - \mu) = \pm 2$$

$$Z = \frac{\bar{x} - \mu}{\frac{\sigma}{\sqrt{n}} \cdot \sqrt{\frac{N-n}{N-1}}} = \frac{\pm 2}{\frac{10.37}{\sqrt{100}} \cdot \sqrt{\frac{4000-100}{4000-1}}} = \pm 1.95$$

$$f(Z) = 0.4744 \cdot 2 = 0.9488$$

Note that "differs by less" references the area on both sides of the mean within 2 pounds right or left.

Example 7.5.2

When a customer places an order with Rudy's On-Line Office Supplies, a computerized accounting information system (AIS) automatically checks to see if the customer has exceeded his or her credit limit. Past records indicate that the probability of customers exceeding their credit limit is .06.

Suppose that on a given day, 3,000 orders are placed in total. If we randomly select 360 orders, what is the probability that between 10 and 20 customers will exceed their credit limit?

Answer

Solution 7.2

$$N = 3000, \quad n = 360, \quad p = 0.06$$

$$\sigma_{p'} = \sqrt{\frac{p(1-p)}{n}} \times \sqrt{\frac{N-n}{N-1}} = \sqrt{\frac{0.06(1-0.06)}{360}} \times \sqrt{\frac{3000-360}{3000-1}} = 0.0117$$

$$p_1 = \frac{10}{360} = 0.0278, \quad p_2 = \frac{20}{360} = 0.0556$$
$$Z = \frac{p' - p}{\sqrt{\frac{p(1-p)}{n}} \cdot \sqrt{\frac{N-n}{N-1}}} = \frac{0.0278 - 0.06}{0.011744} = -2.74$$

$\left[\frac{0.0278-0.06}{0.011744} < \frac{0.0556-0.06}{0.011744}\right]$

This page titled [7.5: Finite Population Correction Factor](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.4: Finite Population Correction Factor](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.6: Chapter Key Terms

Average

a number that describes the central tendency of the data; there are a number of specialized averages, including the arithmetic mean, weighted mean, median, mode, and geometric mean.

Central Limit Theorem

Given a random variable with known mean μ and known standard deviation, σ , we are sampling with size n , and we are interested in two new RVs: the sample mean, \bar{X} . If the size (n) of the sample is sufficiently large, then $\bar{X} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$. If the size (n) of the sample is sufficiently large, then the distribution of the sample means will approximate a normal distributions regardless of the shape of the population. The mean of the sample means will equal the population mean. The standard deviation of the distribution of the sample means, $\frac{\sigma}{\sqrt{n}}$, is called the standard error of the mean.

Finite Population Correction Factor

adjusts the variance of the sampling distribution if the population is known and more than 5% of the population is being sampled.

Mean

a number that measures the central tendency; a common name for mean is "average." The term "mean" is a shortened form of "arithmetic mean." By definition, the mean for a sample (denoted by \bar{x}) is $\bar{x} = \frac{\text{Sum of all values in the sample}}{\text{Number of values in the sample}}$, and the mean for a population (denoted by μ) is $\mu = \frac{\text{Sum of all values in the population}}{\text{Number of values in the population}}$.

Normal Distribution

a continuous random variable with pdf $f(x) = \frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\mu)^2}{2\sigma^2}}$, where μ is the mean of the distribution and σ is the standard deviation.; notation: $X \sim N(\mu, \sigma)$. If $\mu = 0$ and $\sigma = 1$, the random variable, Z , is called the **standard normal distribution**.

Sampling Distribution

Given simple random samples of size n from a given population with a measured characteristic such as mean, proportion, or standard deviation for each sample, the probability distribution of all the measured characteristics is called a sampling distribution.

Standard Error of the Mean

the standard deviation of the distribution of the sample means, or $\frac{\sigma}{\sqrt{n}}$.

Standard Error of the Proportion

the standard deviation of the sampling distribution of proportions

This page titled [7.6: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

7.7: Chapter Review

7.1 The Central Limit Theorem for Sample Means

In a population whose distribution may be known or unknown, if the size (n) of samples is sufficiently large, the distribution of the sample means will be approximately normal. The mean of the sample means will equal the population mean. The standard deviation of the distribution of the sample means, called the standard error of the mean, is equal to the population standard deviation divided by the square root of the sample size (n).

7.2 Using the Central Limit Theorem

The Central Limit Theorem can be used to illustrate the law of large numbers. The law of large numbers states that the larger the sample size you take from a population, the closer the sample mean \bar{x} gets to μ .

7.3 The Central Limit Theorem for Proportions

The Central Limit Theorem can also be used to illustrate that the sampling distribution of sample proportions is normally distributed with the expected value of p and a standard deviation of $\sigma_p = \sqrt{\frac{p(1-p)}{n}}$

This page titled [7.7: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **7.10: Chapter Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.8: Chapter Formula Review

7.1 The Central Limit Theorem for Sample Means

The Central Limit Theorem for Sample Means:

$$\bar{X} \sim N\left(\mu_{\bar{x}}, \frac{\sigma}{\sqrt{n}}\right)$$

$$Z = \frac{\bar{X} - \mu_{\bar{x}}}{\sigma_{\bar{x}}} = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$$

The Mean \bar{X} : $\mu_{\bar{x}}$

Central Limit Theorem for Sample Means z-score $z = \frac{\bar{x} - \mu_{\bar{x}}}{\left(\frac{\sigma}{\sqrt{n}}\right)}$

Standard Error of the Mean (Standard Deviation (\bar{X})): $\frac{\sigma}{\sqrt{n}}$

Finite Population Correction Factor for the sampling distribution of means: $Z = \frac{\bar{x} - \mu}{\frac{\sigma}{\sqrt{n}} \cdot \sqrt{\frac{N-n}{N-1}}}$

Finite Population Correction Factor for the sampling distribution of proportions: $\sigma_{p'} = \sqrt{\frac{p(1-p)}{n}} \times \sqrt{\frac{N-n}{N-1}}$

This page titled [7.8: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.5: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.9: Chapter Practice

Using the Central Limit Theorem

Use the following information to answer the next ten exercises: A manufacturer produces 25-pound lifting weights. The lowest actual weight is 24 pounds, and the highest is 26 pounds. Each weight is equally likely so the distribution of weights is uniform. A sample of 100 weights is taken.

1.

1. What is the distribution for the weights of one 25-pound lifting weight? What is the mean and standard deviation?
2. What is the distribution for the mean weight of 100 25-pound lifting weights?
3. Find the probability that the mean actual weight for the 100 weights is less than 24.9.

2.

Draw the graph from **Exercise 7.9.1**

3.

Find the probability that the mean actual weight for the 100 weights is greater than 25.2.

4.

Draw the graph from **Exercise 7.9.3**

5.

Find the 90th percentile for the mean weight for the 100 weights.

6.

Draw the graph from **Exercise 7.9.5**

7.

1. What is the distribution for the sum of the weights of 100 25-pound lifting weights?
2. Find $P(\sum x < 2,450)$.

8.

Draw the graph from **Exercise 7.9.7**

9.

Find the 90th percentile for the total weight of the 100 weights.

10.

Draw the graph from **Exercise 7.9.9**

Use the following information to answer the next five exercises: The length of time a particular smartphone's battery lasts follows an exponential distribution with a mean of ten months. A sample of 64 of these smartphones is taken.

11.

1. What is the standard deviation?
2. What is the parameter m ?

12.

What is the distribution for the length of time one battery lasts?

13.

What is the distribution for the mean length of time 64 batteries last?

14.

What is the distribution for the total length of time 64 batteries last?

15.

Find the probability that the sample mean is between seven and 11.

16.

Find the 80th percentile for the total length of time 64 batteries last.

17.

Find the *IQR* for the mean amount of time 64 batteries last.

18.

Find the middle 80% for the total amount of time 64 batteries last.

Use the following information to answer the next eight exercises: A uniform distribution has a minimum of six and a maximum of ten. A sample of 50 is taken.

19.

Find $P(\Sigma x > 420)$.

20.

Find the 90th percentile for the sums.

21.

Find the 15th percentile for the sums.

22.

Find the first quartile for the sums.

23.

Find the third quartile for the sums.

24.

Find the 80th percentile for the sums.

25.

A population has a mean of 25 and a standard deviation of 2. If it is sampled repeatedly with samples of size 49, what is the mean and standard deviation of the sample means?

26.

A population has a mean of 48 and a standard deviation of 5. If it is sampled repeatedly with samples of size 36, what is the mean and standard deviation of the sample means?

27.

A population has a mean of 90 and a standard deviation of 6. If it is sampled repeatedly with samples of size 64, what is the mean and standard deviation of the sample means?

28.

A population has a mean of 120 and a standard deviation of 2.4. If it is sampled repeatedly with samples of size 40, what is the mean and standard deviation of the sample means?

29.

A population has a mean of 17 and a standard deviation of 1.2. If it is sampled repeatedly with samples of size 50, what is the mean and standard deviation of the sample means?

30.

A population has a mean of 17 and a standard deviation of 0.2. If it is sampled repeatedly with samples of size 16, what is the expected value and standard deviation of the sample means?

31.

A population has a mean of 38 and a standard deviation of 3. If it is sampled repeatedly with samples of size 48, what is the expected value and standard deviation of the sample means?

32.

A population has a mean of 14 and a standard deviation of 5. If it is sampled repeatedly with samples of size 60, what is the expected value and standard deviation of the sample means?

The Central Limit Theorem for Proportions

33.

A question is asked of a class of 200 freshmen, and 23% of the students know the correct answer. If a sample of 50 students is taken repeatedly, what is the expected value of the mean of the sampling distribution of sample proportions?

34.

A question is asked of a class of 200 freshmen, and 23% of the students know the correct answer. If a sample of 50 students is taken repeatedly, what is the standard deviation of the mean of the sampling distribution of sample proportions?

35.

A game is played repeatedly. A player wins one-fifth of the time. If samples of 40 times the game is played are taken repeatedly, what is the expected value of the mean of the sampling distribution of sample proportions?

36.

A game is played repeatedly. A player wins one-fifth of the time. If samples of 40 times the game is played are taken repeatedly, what is the standard deviation of the mean of the sampling distribution of sample proportions?

37.

A virus attacks one in three of the people exposed to it. An entire large city is exposed. If samples of 70 people are taken, what is the expected value of the mean of the sampling distribution of sample proportions?

38.

A virus attacks one in three of the people exposed to it. An entire large city is exposed. If samples of 70 people are taken, what is the standard deviation of the mean of the sampling distribution of sample proportions?

39.

A company inspects products coming through its production process, and rejects detected products. One-tenth of the items are rejected. If samples of 50 items are taken, what is the expected value of the mean of the sampling distribution of sample proportions?

40.

A company inspects products coming through its production process, and rejects detected products. One-tenth of the items are rejected. If samples of 50 items are taken, what is the standard deviation of the mean of the sampling distribution of sample proportions?

Finite Population Correction Factor

41.

A fishing boat has 1,000 fish on board, with an average weight of 120 pounds and a standard deviation of 6.0 pounds. If sample sizes of 50 fish are checked, what is the probability the fish in a sample will have mean weight within 2.8 pounds the true mean of the population?

42.

An experimental garden has 500 sunflowers plants. The plants are being treated so they grow to unusual heights. The average height is 9.3 feet with a standard deviation of 0.5 foot. If sample sizes of 60 plants are taken, what is the probability the plants in a given sample will have an average height within 0.1 foot of the true mean of the population?

43.

A company has 800 employees. The average number of workdays between absence for illness is 123 with a standard deviation of 14 days. Samples of 50 employees are examined. What is the probability a sample has a mean of workdays with no absence for illness of at least 124 days?

44.

Cars pass an automatic speed check device that monitors 2,000 cars on a given day. This population of cars has an average speed of 67 miles per hour with a standard deviation of 2 miles per hour. If samples of 30 cars are taken, what is the probability a given sample will have an average speed within 0.50 mile per hour of the population mean?

45.

A town keeps weather records. From these records it has been determined that it rains on an average of 37% of the days each year. If 30 days are selected at random from one year, what is the probability that at least 5 and at most 11 days had rain?

46.

A maker of yardsticks has an ink problem that causes the markings to smear on 4% of the yardsticks. The daily production run is 2,000 yardsticks. What is the probability if a sample of 100 yardsticks is checked, there will be ink smeared on at most 4 yardsticks?

47.

A school has 300 students. Usually, there are an average of 21 students who are absent. If a sample of 30 students is taken on a certain day, what is the probability that at most 2 students in the sample will be absent?

48.

A college gives a placement test to 5,000 incoming students each year. On the average 1,213 place in one or more developmental courses. If a sample of 50 is taken from the 5,000, what is the probability at most 12 of those sampled will have to take at least one developmental course?

This page titled [7.9: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.8: Chapter Practice](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.10: Chapter Homework

The Central Limit Theorem for Sample Means

49

Previously, De Anza statistics students estimated that the amount of change daytime statistics students carry is exponentially distributed with a mean of \$0.88. Suppose that we randomly pick 25 daytime statistics students.

- In words, $X =$ _____
- $X \sim$ ____ (____, ____)
- In words, $\bar{X} =$ _____
- $\bar{X} \sim$ ____ (____, ____)
- Find the probability that an individual had between \$0.80 and \$1.00. Graph the situation, and shade in the area to be determined.
- Find the probability that the average of the 25 students was between \$0.80 and \$1.00. Graph the situation, and shade in the area to be determined.
- Explain why there is a difference in part e and part f.

Answer

- $X =$ amount of change students carry
- $X \sim E(0.88, 0.88)$
- $\bar{X} =$ average amount of change carried by a sample of 25 students.
- $\bar{X} \sim N(0.88, 0.176)$
- 0.0819
- 0.1882
- The distributions are different. Part 1 is exponential and part 2 is normal.

50.

Suppose that the distance of fly balls hit to the outfield (in baseball) is normally distributed with a mean of 250 feet and a standard deviation of 50 feet. We randomly sample 49 fly balls.

- If $\bar{X} =$ average distance in feet for 49 fly balls, then $\bar{X} \sim$ ____ (____, ____)
- What is the probability that the 49 balls traveled an average of less than 240 feet? Sketch the graph. Scale the horizontal axis for \bar{X} . Shade the region corresponding to the probability. Find the probability.
- Find the 80th percentile of the distribution of the average of 49 fly balls.

51.

According to the Internal Revenue Service, the average length of time for an individual to complete (keep records for, learn, prepare, copy, assemble, and send) IRS Form 1040 is 10.53 hours (without any attached schedules). The distribution is unknown. Let us assume that the standard deviation is two hours. Suppose we randomly sample 36 taxpayers.

- In words, $X =$ _____
- In words, $\bar{X} =$ _____
- $\bar{X} \sim$ ____ (____, ____)
- Would you be surprised if the 36 taxpayers finished their Form 1040s in an average of more than 12 hours? Explain why or why not in complete sentences.
- Would you be surprised if one taxpayer finished his or her Form 1040 in more than 12 hours? In a complete sentence, explain why.

52.

Suppose that a category of world-class runners are known to run a marathon (26 miles) in an average of 145 minutes with a standard deviation of 14 minutes. Consider 49 of the races. Let \bar{X} the average of the 49 races.

- $\bar{X} \sim$ ____ (____, ____)

2. Find the probability that the runner will average between 142 and 146 minutes in these 49 marathons.
3. Find the 80th percentile for the average of these 49 marathons.
4. Find the median of the average running times.

53.

The length of songs in a collector's iTunes album collection is uniformly distributed from two to 3.5 minutes. Suppose we randomly pick five albums from the collection. There are a total of 43 songs on the five albums.

1. In words, $X =$ _____
2. $X \sim$ _____
3. In words, $\bar{X} =$ _____
4. $\bar{X} \sim$ _____(_____, _____)
5. Find the first quartile for the average song length.
6. The *IQR* (interquartile range) for the average song length is from _____ - _____.

54.

In 1940 the average size of a U.S. farm was 174 acres. Let's say that the standard deviation was 55 acres. Suppose we randomly survey 38 farmers from 1940.

1. In words, $X =$ _____
2. In words, $\bar{X} =$ _____
3. $\bar{X} \sim$ _____(_____, _____)
4. The *IQR* for \bar{X} is from _____ acres to _____ acres.

55.

Determine which of the following are true and which are false. Then, in complete sentences, justify your answers.

1. When the sample size is large, the mean of \bar{X} is approximately equal to the mean of X .
2. When the sample size is large, \bar{X} is approximately normally distributed.
3. When the sample size is large, the standard deviation of \bar{X} is approximately the same as the standard deviation of X .

56.

The percent of fat calories that a person in America consumes each day is normally distributed with a mean of about 36 and a standard deviation of about ten. Suppose that 16 individuals are randomly chosen. Let \bar{X} = average percent of fat calories.

- a. $\bar{X} \sim$ _____(_____, _____)
- b. For the group of 16, find the probability that the average percent of fat calories consumed is more than five. Graph the situation and shade in the area to be determined.
- c. Find the first quartile for the average percent of fat calories.

57.

The distribution of income in some Third World countries is considered wedge shaped (many very poor people, very few middle income people, and even fewer wealthy people). Suppose we pick a country with a wedge shaped distribution. Let the average salary be \$2,000 per year with a standard deviation of \$8,000. We randomly survey 1,000 residents of that country.

1. In words, $X =$ _____
2. In words, $\bar{X} =$ _____
3. $\bar{X} \sim$ _____(_____, _____)
4. How is it possible for the standard deviation to be greater than the average?
5. Why is it more likely that the average of the 1,000 residents will be from \$2,000 to \$2,100 than from \$2,100 to \$2,200?

58.

Which of the following is NOT TRUE about the distribution for averages?

1. The mean, median, and mode are equal.
2. The area under the curve is one.

3. The curve never touches the x-axis.
4. The curve is skewed to the right.

59.

The cost of unleaded gasoline in the Bay Area once followed an unknown distribution with a mean of \$4.59 and a standard deviation of \$0.10. Sixteen gas stations from the Bay Area are randomly chosen. We are interested in the average cost of gasoline for the 16 gas stations. The distribution to use for the average cost of gasoline for the 16 gas stations is:

- a. $\bar{X} \sim N(4.59, 0.10)$
- b. $\bar{X} \sim N\left(4.59, \frac{0.10}{\sqrt{16}}\right)$
- c. $\bar{X} \sim N\left(4.59, \frac{16}{0.10}\right)$
- d. $\bar{X} \sim N\left(4.59, \frac{\sqrt{16}}{0.10}\right)$

Using the Central Limit Theorem

60.

A large population of 5,000 students take a practice test to prepare for a standardized test. The population mean is 140 questions correct, and the standard deviation is 80. What size samples should a researcher take to get a distribution of means of the samples with a standard deviation of 10?

61.

A large population has skewed data with a mean of 70 and a standard deviation of 6. Samples of size 100 are taken, and the distribution of the means of these samples is analyzed.

1. Will the distribution of the means be closer to a normal distribution than the distribution of the population?
2. Will the mean of the means of the samples remain close to 70?
3. Will the distribution of the means have a smaller standard deviation?
4. What is that standard deviation?

62.

A researcher is looking at data from a large population with a standard deviation that is much too large. In order to concentrate the information, the researcher decides to repeatedly sample the data and use the distribution of the means of the samples? The first effort used sample sized of 100. But the standard deviation was about double the value the researcher wanted. What is the smallest size samples the researcher can use to remedy the problem?

63.

A researcher looks at a large set of data, and concludes the population has a standard deviation of 40. Using sample sizes of 64, the researcher is able to focus the mean of the means of the sample to a narrower distribution where the standard deviation is 5. Then, the researcher realizes there was an error in the original calculations, and the initial standard deviation is really 20. Since the standard deviation of the means of the samples was obtained using the original standard deviation, this value is also impacted by the discovery of the error. What is the correct value of the standard deviation of the means of the samples?

64.

A population has a standard deviation of 50. It is sampled with samples of size 100. What is the variance of the means of the samples?

The Central Limit Theorem for Proportions

65.

A farmer picks pumpkins from a large field. The farmer makes samples of 260 pumpkins and inspects them. If one in fifty pumpkins are not fit to market and will be saved for seeds, what is the standard deviation of the mean of the sampling distribution of sample proportions?

66.

A store surveys customers to see if they are satisfied with the service they received. Samples of 25 surveys are taken. One in five people are unsatisfied. What is the variance of the mean of the sampling distribution of sample proportions for the number of unsatisfied customers? What is the variance for satisfied customers?

67.

A company gives an anonymous survey to its employees to see what percent of its employees are happy. The company is too large to check each response, so samples of 50 are taken, and the tendency is that three-fourths of the employees are happy. For the mean of the sampling distribution of sample proportions, answer the following questions, if the sample size is doubled.

1. How does this affect the mean?
2. How does this affect the standard deviation?
3. How does this affect the variance?

68.

A pollster asks a single question with only yes and no as answer possibilities. The poll is conducted nationwide, so samples of 100 responses are taken. There are four yes answers for each no answer overall. For the mean of the sampling distribution of sample proportions, find the following for yes answers.

- a. The expected value.
- b. The standard deviation.
- c. The variance.

69.

The mean of the sampling distribution of sample proportions has a value of p of 0.3, and sample size of 40.

1. Is there a difference in the expected value if p and q reverse roles?
2. Is there a difference in the calculation of the standard deviation with the same reversal?

Finite Population Correction Factor

70.

A company has 1,000 employees. The average number of workdays between absence for illness is 80 with a standard deviation of 11 days. Samples of 80 employees are examined. What is the probability a sample has a mean of workdays with no absence for illness of at least 78 days and at most 84 days?

71.

Trucks pass an automatic scale that monitors 2,000 trucks. This population of trucks has an average weight of 20 tons with a standard deviation of 2 tons. If a sample of 50 trucks is taken, what is the probability the sample will have an average weight within one-half ton of the population mean?

72.

A town keeps weather records. From these records it has been determined that it rains on an average of 12% of the days each year. If 30 days are selected at random from one year, what is the probability that at most 3 days had rain?

73.

A maker of greeting cards has an ink problem that causes the ink to smear on 7% of the cards. The daily production run is 500 cards. What is the probability that if a sample of 35 cards is checked, there will be ink smeared on at most 5 cards?

74.

A school has 500 students. Usually, there are an average of 20 students who are absent. If a sample of 30 students is taken on a certain day, what is the probability that at least 2 students in the sample will be absent?

This page titled [7.10: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.6: Chapter Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.11: Chapter References

7.1 The Central Limit Theorem for Sample Means

Baran, Daya. “20 Percent of Americans Have Never Used Email.” WebGuild, 2010. Available online at <http://www.webguild.org/20080519/20-...ver-used-email> (accessed May 17, 2013).

Data from The Flurry Blog, 2013. Available online at blog.flurry.com (accessed May 17, 2013).

Data from the United States Department of Agriculture.

This page titled [7.11: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [7.9: Chapter References](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

7.12: Chapter Solution (Practice + Homework)

1.

1. 3.

0.0003

25.07

7.

1. 9.

2,507.40

1. 13.

$N(10, \frac{10}{8})$

0.7799

17.

1.69

19.

0.0072

21.

391.54

23.

405.51

25.

Mean = 25, standard deviation = $\frac{2}{7}$

26.

Mean = 48, standard deviation = $\frac{5}{6}$

27.

Mean = 90, standard deviation = $\frac{3}{4}$

28.

Mean = 120, standard deviation = 0.38

29.

Mean = 17, standard deviation = 0.17

30.

Expected value = 17, standard deviation = 0.05

31.

Expected value = 38, standard deviation = 0.43

32.

Expected value = 14, standard deviation = 0.65

33.

0.23

34.

0.060

35.

1/5

36.

0.063

37.

1/3

38.

0.056

39.

1/10

40.

0.042

41.

0.999

42.

0.901

43.

0.301

44.

0.832

45.

0.483

46.

0.500

47.

0.502

48.

0.519

49.

1. 51.

1. 53.

1. 55.

1. 57.

1. 59.

b

64

61.

1. **62.**

400

2.5

64.

25

65.

0.0087

66.

0.0064, 0.0064

67.

1. **68.**

1. **69.**

a. **70.**

0.955

0.927

72.

0.648

73.

0.101

74.

0.273

This page titled [7.12: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **7.11: Chapter Solution (Practice + Homework)** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

8: Confidence Intervals

- 8.1: Introduction to Confidence Intervals
- 8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size
- 8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case
- 8.4: A Confidence Interval for A Population Proportion
- 8.5: Calculating the Sample Size n - Continuous and Binary Random Variables
- 8.6: Chapter Key Terms
- 8.7: Chapter Review
- 8.8: Chapter Formula Review
- 8.9: Chapter Practice
- 8.10: Chapter Homework
- 8.11: Chapter References
- 8.12: Chapter Solution (Practice + Homework)

This page titled [8: Confidence Intervals](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

8.1: Introduction to Confidence Intervals

Suppose you were trying to determine the mean rent of a two-bedroom apartment in your town. You might look in the classified section of the newspaper, write down several rents listed, and average them together. You would have obtained a point estimate of the true mean. If you are trying to determine the percentage of times you make a basket when shooting a basketball, you might count the number of shots you make and divide that by the number of shots you attempted. In this case, you would have obtained a point estimate for the true proportion the parameter p in the binomial probability density function.




Figure 8.1.1 Have you ever wondered what the average number of M&Ms in a bag at the grocery store is? You can use confidence intervals to answer this question. (credit: comedy_nose/flickr)

We use sample data to make generalizations about an unknown population. This part of statistics is called **inferential statistics**. **The sample data help us to make an estimate of a population parameter.** We realize that the point estimate is most likely not the exact value of the population parameter, but close to it. After calculating point estimates, we construct interval estimates, called confidence intervals. What statistics provides us beyond a simple average, or point estimate, is an estimate to which we can attach a probability of accuracy, what we will call a confidence level. We make inferences with a known level of probability.

In this chapter, you will learn to construct and interpret confidence intervals. You will also learn a new distribution, the Student's-t, and how it is used with these intervals. Throughout the chapter, it is important to keep in mind that the confidence interval is a random variable. It is the population parameter that is fixed.

If you worked in the marketing department of an entertainment company, you might be interested in the mean number of songs a consumer downloads a month from iTunes. If so, you could conduct a survey and calculate the sample mean, \bar{x} , and the sample standard deviation, s . You would use \bar{x} to estimate the population mean and s to estimate the population standard deviation. The sample mean, \bar{x} , is the point estimate for the population mean, μ . The sample standard deviation, s , is the point estimate for the population standard deviation, σ .

\bar{x} and s are each called a statistic.

A **confidence interval** is another type of estimate but, instead of being just one number, it is an interval of numbers. The interval of numbers is a range of values calculated from a given set of sample data. The confidence interval is likely to include the unknown population parameter.

Suppose, for the iTunes example, we do not know the population mean μ , but we do know that the population standard deviation is $\sigma = 1$ and our sample size is 100. Then, by the central limit theorem, the standard deviation of the sampling distribution of the sample means is

$$\frac{\sigma}{\sqrt{n}} = \frac{1}{\sqrt{100}} = 0.1.$$

The **empirical rule**, which applies to the normal distribution, says that in approximately 95% of the samples, the sample mean, \bar{x} , will be within two standard deviations of the population mean μ . For our iTunes example, two standard deviations is $(2)(0.1) = 0.2$. The sample mean \bar{x} is likely to be within 0.2 units of μ .

Because \bar{x} is within 0.2 units of μ , which is unknown, then μ is likely to be within 0.2 units of \bar{x} with 95% probability. The population mean μ is contained in an interval whose lower number is calculated by taking the sample mean and subtracting two

standard deviations $(2)(0.1)$ and whose upper number is calculated by taking the sample mean and adding two standard deviations. In other words, μ is between $\bar{x} - 0.2$ and $\bar{x} + 0.2$ in 95% of all the samples.

For the iTunes example, suppose that a sample produced a sample mean $\bar{x} = 2$. Then with 95% probability the unknown population mean μ is between

$$\bar{x} - 0.2 = 2 - 0.2 = 1.8 \text{ and } \bar{x} + 0.2 = 2 + 0.2 = 2.2$$

We say that we are **95% confident** that the unknown population mean number of songs downloaded from iTunes per month is between 1.8 and 2.2. **The 95% confidence interval is (1.8, 2.2)**. Please note that we talked in terms of 95% confidence using the empirical rule. The empirical rule for two standard deviations is only approximately 95% of the probability under the normal distribution. To be precise, two standard deviations under a normal distribution is actually 95.44% of the probability. To calculate the exact 95% confidence level we would use 1.96 standard deviations.

The 95% confidence interval implies two possibilities. Either the interval (1.8, 2.2) contains the true mean μ , or our sample produced an \bar{x} that is not within 0.2 units of the true mean μ . The second possibility happens for only 5% of all the samples (95% minus 100% = 5%).

Remember that a confidence interval is created for an unknown population parameter like the population mean, μ .

For the confidence interval for a mean the formula would be:

$$\mu = \bar{X} \pm Z_{\alpha}\sigma/\sqrt{n}$$

Or written another way as:

$$\bar{X} - Z_{\alpha}\sigma/\sqrt{n} \leq \mu \leq \bar{X} + Z_{\alpha}\sigma/\sqrt{n}$$

Where \bar{x} is the sample mean. Z_{α} is determined by the level of confidence desired by the analyst, and σ/\sqrt{n} is the standard deviation of the sampling distribution for means given to us by the Central Limit Theorem.

This page titled [8.1: Introduction to Confidence Intervals](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **8.0: Introduction to Confidence Intervals** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size

A confidence interval for a population mean with a known population standard deviation is based on the conclusion of the Central Limit Theorem that the sampling distribution of the sample means follow an approximately normal distribution.

Calculating the Confidence Interval

Consider the standardizing formula for the sampling distribution developed in the discussion of the Central Limit Theorem:

$$Z_1 = \frac{\bar{X} - \mu_{\bar{X}}}{\sigma_{\bar{X}}} = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}}$$

Notice that μ is substituted for $\mu_{\bar{X}}$ because we know that the expected value of $\mu_{\bar{X}}$ is μ from the Central Limit theorem and $\sigma_{\bar{X}}$ is replaced with σ/\sqrt{n} , also from the Central Limit Theorem.

In this formula we know \bar{X} , $\sigma_{\bar{X}}$ and n , the sample size. (In actuality we do not know the population standard deviation, but we do have a point estimate for it, s , from the sample we took. More on this later.) What we do not know is μ or Z_1 . We can solve for either one of these in terms of the other. Solving for μ in terms of Z_1 gives:

$$\mu = \bar{X} \pm Z_1 \sigma / \sqrt{n}$$

Remembering that the Central Limit Theorem tells us that the distribution of the \bar{X} 's, the sampling distribution for means, is normal, and that the normal distribution is symmetrical, we can rearrange terms thus:

$$\bar{X} - Z_\alpha(\sigma/\sqrt{n}) \leq \mu \leq \bar{X} + Z_\alpha(\sigma/\sqrt{n})$$

This is the formula for a confidence interval for the mean of a population.

Notice that Z_α has been substituted for Z_1 in this equation. This is where a choice must be made by the statistician. The analyst must decide the level of confidence they wish to impose on the confidence interval. α is the probability that the interval will not contain the true population mean. The confidence level is defined as $(1 - \alpha)$. Z_α is the number of standard deviations \bar{X} lies from the mean with a certain probability. If we chose $Z_\alpha = 1.96$ we are asking for the 95% confidence interval because we are setting the probability that the true mean lies within the range at 0.95. If we set Z_α at 1.64 we are asking for the 90% confidence interval because we have set the probability at 0.90. These numbers can be verified by consulting the Standard Normal table. Divide either 0.95 or 0.90 in half and find that probability inside the body of the table. Then read on the top and left margins the number of standard deviations it takes to get this level of probability.

In reality, we can set whatever level of confidence we desire simply by changing the Z_α value in the formula. It is the analyst's choice. Common convention in Economics and most social sciences sets confidence intervals at either 90, 95, or 99 percent levels. Levels less than 90% are considered of little value. The level of confidence of a particular interval estimate is called by $(1 - \alpha)$.

A good way to see the development of a confidence interval is to graphically depict the solution to a problem requesting a confidence interval. This is presented in Figure 8.2.2 for the example in the introduction concerning the number of downloads from iTunes. That case was for a 95% confidence interval, but other levels of confidence could have just as easily been chosen depending on the need of the analyst. However, the level of confidence MUST be pre-set and not subject to revision as a result of the calculations.




Figure 8.2.2

For this example, let's say we know that the actual population mean number of iTunes downloads is 2.1. The true population mean falls within the range of the 95% confidence interval. There is absolutely nothing to guarantee that this will happen. **Further, if the true mean falls outside of the interval we will never know it. We must always remember that we will never ever know the**

true mean. Statistics simply allows us, with a given level of probability (confidence), to say that the true mean is within the range calculated. This is what was called in the introduction, the "level of ignorance admitted".

Changing the Confidence Level or Sample Size

Here again is the formula for a confidence interval for an unknown population mean assuming we know the population standard deviation:

$$\bar{X} - Z_{\alpha}(\sigma/\sqrt{n}) \leq \mu \leq \bar{X} + Z_{\alpha}(\sigma/\sqrt{n})$$

It is clear that the confidence interval is driven by two things, the chosen level of confidence, Z_{α} , and the standard deviation of the sampling distribution. The Standard deviation of the sampling distribution is further affected by two things, the standard deviation of the population and the sample size we chose for our data. Here we wish to examine the effects of each of the choices we have made on the calculated confidence interval, the confidence level and the sample size.

For a moment we should ask just what we desire in a confidence interval. Our goal was to estimate the population mean from a sample. We have forsaken the hope that we will ever find the true population mean, and population standard deviation for that matter, for any case except where we have an extremely small population and the cost of gathering the data of interest is very small. In all other cases we must rely on samples. With the Central Limit Theorem we have the tools to provide a meaningful confidence interval with a given level of confidence, meaning a known probability of being wrong. By meaningful confidence interval we mean one that is useful. Imagine that you are asked for a confidence interval for the ages of your classmates. You have taken a sample and find a mean of 19.8 years. You wish to be very confident so you report an interval between 9.8 years and 29.8 years. This interval would certainly contain the true population mean and have a very high confidence level. However, it hardly qualifies as meaningful. The very best confidence interval is narrow while having high confidence. There is a natural tension between these two goals. The higher the level of confidence the wider the confidence interval as the case of the students' ages above. We can see this tension in the equation for the confidence interval.

$$\mu = \bar{x} \pm Z_{\alpha} \left(\frac{\sigma}{\sqrt{n}} \right)$$

The confidence interval will increase in width as Z_{α} increases, Z_{α} increases as the level of confidence increases. There is a tradeoff between the level of confidence and the width of the interval. Now let's look at the formula again and we see that the sample size also plays an important role in the width of the confidence interval. The sample sized, nn , shows up in the denominator of the standard deviation of the sampling distribution. As the sample size increases, the standard deviation of the sampling distribution decreases and thus the width of the confidence interval, while holding constant the level of confidence. This relationship was demonstrated in Figure 8.2.8. Again we see the importance of having large samples for our analysis although we then face a second constraint, the cost of gathering data.

Calculating the Confidence Interval: An Alternative Approach

Another way to approach confidence intervals is through the use of something called the **Error Bound**. The Error Bound gets its name from the recognition that it provides the boundary of the interval derived from the standard error of the sampling distribution. In the equations above it is seen that the interval is simply the estimated mean, sample mean, plus or minus something. That something is the Error Bound and is driven by the probability we desire to maintain in our estimate, Z_{α} , times the standard deviation of the sampling distribution. The Error Bound for a mean is given the name, **Error Bound Mean**, or *EBM*.

To construct a confidence interval for a single unknown population mean μ , **where the population standard deviation is known**, we need \bar{x} as an estimate for μ and we need the margin of error. Here, the margin of error (*EBM*) is called the error bound for a population mean (abbreviated **EBM**). The sample mean \bar{x} is the **point estimate** of the unknown population mean μ .

The confidence interval estimate will have the form:

(point estimate - error bound, point estimate + error bound) or, in symbols, $(\bar{x} - EBM, \bar{x} + EBM)$

The mathematical formula for this confidence interval is:

$$\bar{X} - Z_{\alpha}(\sigma/\sqrt{n}) \leq \mu \leq \bar{X} + Z_{\alpha}(\sigma/\sqrt{n}) \quad (8.2.1)$$

The margin of error (EBM) depends on the **confidence level** (abbreviated **CL**). The confidence level is often considered the probability that the calculated confidence interval estimate will contain the true population parameter. However, it is more accurate

to state that the confidence level is the percent of confidence intervals that contain the true population parameter when repeated samples are taken. Most often, it is the choice of the person constructing the confidence interval to choose a confidence level of 90% or higher because that person wants to be reasonably certain of his or her conclusions.

There is another probability called alpha (α). α is related to the confidence level, CL . α is the probability that the interval does not contain the unknown population parameter.

Mathematically, $1 - \alpha = CL$.

A confidence interval for a population mean with a **known** standard deviation is based on the fact that the sampling distribution of the sample means follow an approximately normal distribution. Suppose that our sample has a mean of $\bar{x} = 10$, and we have constructed the 90% confidence interval (5, 15) where $EBM = 5$.

To get a 90% confidence interval, we must include the central 90% of the probability of the normal distribution. If we include the central 90%, we leave out a total of $\alpha = 10$ in both tails, or 5% in each tail, of the normal distribution.




Figure 8.2.3

To capture the central 90%, we must go out 1.645 standard deviations on either side of the calculated sample mean. The value 1.645 is the z-score from a standard normal probability distribution that puts an area of 0.90 in the center, an area of 0.05 in the far left tail, and an area of 0.05 in the far right tail.

It is important that the standard deviation used must be appropriate for the parameter we are estimating, so in this section we need to use the standard deviation that applies to the sampling distribution for means which we studied with the Central Limit Theorem and is, $\frac{\sigma}{\sqrt{n}}$.

Calculating the Confidence Interval Using EMB

To construct a confidence interval estimate for an unknown population mean, we need data from a random sample. The steps to construct and interpret the confidence interval are:

- Calculate the sample mean \bar{x} from the sample data. Remember, in this section we know the population standard deviation σ .
- Find the z-score from the standard normal table that corresponds to the confidence level desired.
- Calculate the error bound EBM .
- Construct the confidence interval.
- Write a sentence that interprets the estimate in the context of the situation in the problem.

We will first examine each step in more detail, and then illustrate the process with some examples.

Finding the z-score for the Stated Confidence Level

When we know the population standard deviation σ , we use a standard normal distribution to calculate the error bound EBM and construct the confidence interval. We need to find the value of z that puts an area equal to the confidence level (in decimal form) in the middle of the standard normal distribution $Z \sim N(0, 1)$.

The confidence level, CL , is the area in the middle of the standard normal distribution. $CL = 1 - \alpha$, so α is the area that is split equally between the two tails. Each of the tails contains an area equal to $\frac{\alpha}{2}$.

The z-score that has an area to the right of $\frac{\alpha}{2}$ is denoted by $Z_{\frac{\alpha}{2}}$.

For example, when $CL = 0.95$, $\alpha = 0.05$ and $\frac{\alpha}{2} = 0.025$; we write $Z_{\frac{\alpha}{2}} = Z_{0.025}$.

The area to the right of $Z_{0.025}$ is 0.025 and the area to the left of $Z_{0.025}$ is $1 - 0.025 = 0.975$

$Z_{\frac{\alpha}{2}} = Z_{0.025} = 1.96$, using a standard normal probability table. We will see later that we can use a different probability table, the Student's t-distribution, for finding the number of standard deviations of commonly used levels of confidence.

Calculating the Error Bound (EBM)

The error bound formula for an unknown population mean μ when the population standard deviation σ is known is

- $EBM = \left(Z_{\frac{\alpha}{2}} \right) \left(\frac{\sigma}{\sqrt{n}} \right)$

Constructing the Confidence Interval

- The confidence interval estimate has the format $(\bar{x} - EBM, \bar{x} + EBM)$ or the formula:

$$\bar{X} - Z_{\alpha}(\sigma/\sqrt{n}) \leq \mu \leq \bar{X} + Z_{\alpha}(\sigma/\sqrt{n})$$

The graph gives a picture of the entire situation.

$$CL + \frac{\alpha}{2} + \frac{\alpha}{2} = CL + \alpha = 1 .$$




Figure 8.2.4

Example 8.2.1

Suppose we are interested in the mean scores on an exam. A random sample of 36 scores is taken and gives a sample mean (sample mean score) of 68 ($\bar{X} = 68$). In this example we have the unusual knowledge that the population standard deviation is 3 points. Do not count on knowing the population parameters outside of textbook examples. Find a confidence interval estimate for the population mean exam score (the mean score on all exams).

Find a 90% confidence interval for the true (population) mean of statistics exam scores.

Answer

Solution 8.1

- The solution is shown step-by-step.

To find the confidence interval, you need the sample mean, \bar{x} , and the EBM .

- $\bar{x} = 68$
- $EBM = \left(Z_{\frac{\alpha}{2}} \right) \left(\frac{\sigma}{\sqrt{n}} \right)$
- $\sigma = 3$; $n = 36$; The confidence level is 90% ($CL = 0.90$)

$$CL = 0.90 \text{ so } \alpha = 1 - CL = 1 - 0.90 = 0.10$$

$$\frac{\alpha}{2} = 0.05, Z_{\frac{\alpha}{2}} = z_{0.05}$$

The area to the right of $Z_{0.05}$ is 0.05 and the area to the left of $Z_{0.05}$ is $1 - 0.05 = 0.95$

$$Z_{\frac{\alpha}{2}} = Z_{0.05} = 1.645$$

This can be found using a computer, or using a probability table for the standard normal distribution. Because the common levels of confidence in the social sciences are 90%, 95% and 99% it will not be long until you become familiar with the numbers, 1.645, 1.96, and 2.56

$$EBM = (1.645) \left(\frac{3}{\sqrt{36}} \right) = 0.8225$$

$$\bar{x} - EBM = 68 - 0.8225 = 67.1775$$

$$\bar{x} + EBM = 68 + 0.8225 = 68.8225$$

The 90% confidence interval is **(67.1775, 68.8225)**.

Interpretation

We estimate with 90% confidence that the true population mean exam score for all statistics students is between 67.18 and 68.82.

Example 8.2.2

Suppose we change the original problem in [Example 8.2.1](#) by using a 95% confidence level. Find a 95% confidence interval for the true (population) mean statistics exam score.

Answer

Solution 8.2




Figure 8.2.5

$$\mu = \bar{x} \pm Z_{\alpha} \left(\frac{\sigma}{\sqrt{n}} \right)$$

$$\mu = 68 \pm 1.96 \left(\frac{3}{\sqrt{36}} \right)$$

$$67.02 \leq \mu \leq 68.98$$

$\sigma = 3$; $n = 36$; The confidence level is 95% ($CL = 0.95$).

$CL = 0.95$ so $\alpha = 1 - CL = 1 - 0.95 = 0.05$

$$Z_{\frac{\alpha}{2}} = Z_{0.025} = 1.96$$

Notice that the *EBM* is larger for a 95% confidence level in the original problem.

Comparing the results

The 90% confidence interval is (67.18, 68.82). The 95% confidence interval is (67.02, 68.98). The 95% confidence interval is wider. If you look at the graphs, because the area 0.95 is larger than the area 0.90, it makes sense that the 95% confidence interval is wider. To be more confident that the confidence interval actually does contain the true value of the population mean for all statistics exam scores, the confidence interval necessarily needs to be wider. This demonstrates a very important principle of confidence intervals. There is a trade off between the level of confidence and the width of the interval. Our desire is to have a narrow confidence interval, huge wide intervals provide little information that is useful. But we would also like to have a high level of confidence in our interval. This demonstrates that we cannot have both.




Figure 8.2.6

Summary: Effect of Changing the Confidence Level

- Increasing the confidence level makes the confidence interval wider.
- Decreasing the confidence level makes the confidence interval narrower.

And again here is the formula for a confidence interval for an unknown mean assuming we have the population standard deviation:

$$\bar{X} - Z_{\alpha}(\sigma/\sqrt{n}) \leq \mu \leq \bar{X} + Z_{\alpha}(\sigma/\sqrt{n})$$

The standard deviation of the sampling distribution was provided by the Central Limit Theorem as σ/\sqrt{n} . While we infrequently get to choose the sample size it plays an important role in the confidence interval. Because the sample size is in the denominator of the equation, as n increases it causes the standard deviation of the sampling distribution to decrease and thus the width of the confidence interval to decrease. We have met this before as we reviewed the effects of sample size on the Central Limit Theorem. There we saw that as n increases the sampling distribution narrows until in the limit it collapses on the true population mean.

Example 8.2.3

Suppose we change the original problem in [Example 8.2.1](#) to see what happens to the confidence interval if the sample size is changed.

Leave everything the same except the sample size. Use the original 90% confidence level. What happens to the confidence interval if we increase the sample size and use $n = 100$ instead of $n = 36$? What happens if we decrease the sample size to $n = 25$ instead of $n = 36$?

Answer

Solution 8.3

Solution A

$$\mu = \bar{x} \pm Z_{\alpha} \left(\frac{\sigma}{\sqrt{n}} \right)$$

$$\mu = 68 \pm 1.645 \left(\frac{3}{\sqrt{100}} \right)$$

$$67.5065 \leq \mu \leq 68.4935$$

If we **increase** the sample size n to 100, we **decrease** the width of the confidence interval relative to the original sample size of 36 observations.

Answer

Solution 8.3

Solution B

$$\mu = \bar{x} \pm Z_{\alpha} \left(\frac{\sigma}{\sqrt{n}} \right)$$

$$\mu = 68 \pm 1.645 \left(\frac{3}{\sqrt{25}} \right)$$

$$67.013 \leq \mu \leq 68.987$$

If we **decrease** the sample size n to 25, we **increase** the width of the confidence interval by comparison to the original sample size of 36 observations.

Summary: Effect of Changing the Sample Size

- Increasing the sample size makes the confidence interval narrower.
- Decreasing the sample size makes the confidence interval wider.

We have already seen this effect when we reviewed the effects of changing the size of the sample, n , on the Central Limit Theorem. See Figure 8.2.7 to see this effect. Before we saw that as the sample size increased the standard deviation of the sampling distribution decreases. This was why we choose the sample mean from a large sample as compared to a small sample, all other things held constant.

Thus far we assumed that we knew the population standard deviation. This will virtually never be the case. We will have the sample standard deviation, s , however. This is a point estimate for the population standard deviation and can be substituted into the formula for confidence intervals for a mean under certain circumstances. We just saw the effect the sample size has on the width of confidence interval and the impact on the sampling distribution for our discussion of the Central Limit Theorem. We can invoke this to substitute the point estimate for the standard deviation if the sample size is large "enough". Simulation studies indicate that 30 observations or more will be sufficient to eliminate any meaningful bias in the estimated confidence interval.

Example 8.2.4

Spring break can be a very expensive holiday. A sample of 80 students is surveyed, and the average amount spent by students on travel and beverages is \$593.84. The sample standard deviation is approximately \$369.34.

Construct a 92% confidence interval for the population mean amount of money spent by spring breakers.

Answer

Solution 8.4

We begin with the confidence interval for a mean. We use the formula for a mean because the random variable is dollars spent and this is a continuous random variable. The point estimate for the population standard deviation, s , has been substituted for the true population standard deviation because with 80 observations there is no concern for bias in the estimate of the confidence interval.

$$\mu = \bar{x} \pm \left[Z_{(\alpha/2)} \frac{s}{\sqrt{n}} \right]$$

Substituting the values into the formula, we have:

$$\mu = 593.84 \pm \left[1.75 \frac{369.34}{\sqrt{80}} \right]$$

$Z_{(\alpha/2)}$ is found on the standard normal table by looking up 0.46 in the body of the table and finding the number of standard deviations on the side and top of the table; 1.75. The solution for the interval is thus:

$$\mu = 593.84 \pm 72.2636 = (521.57, 666.10)$$

$$\$521.58 \leq \mu \leq \$666.10$$




Figure 8.2.7

Formula Review

The general form for a confidence interval for a single population mean, known standard deviation, normal distribution is given by $\bar{X} - Z_{\alpha}(\sigma/\sqrt{n}) \leq \mu \leq \bar{X} + Z_{\alpha}(\sigma/\sqrt{n})$. This formula is used when the population standard deviation is known.

CL = confidence level, or the proportion of confidence intervals created that are expected to contain the true population parameter

$\alpha = 1 - CL$ = the proportion of confidence intervals that will not contain the population parameter

$z_{\frac{\alpha}{2}}$ = the z-score with the property that the area to the right of the z-score is $\frac{\alpha}{2}$. This is the z-score used in the calculation of "EBM" where $\alpha = 1 - CL$.

This page titled [8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.1: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size](#) by [OpenStax](#) is licensed [CC BY 4.0](#).
Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case

In practice, we rarely know the population **standard deviation**. In the past, when the sample size was large, this did not present a problem to statisticians. They used the sample standard deviation s as an estimate for σ and proceeded as before to calculate a **confidence interval** with close enough results. This is what we did in Example 8.3.4 above. The point estimate for the standard deviation, s , was substituted in the formula for the confidence interval for the population standard deviation. In this case there 80 observation well above the suggested 30 observations to eliminate any bias from a small sample. However, statisticians ran into problems when the sample size was small. A small sample size caused inaccuracies in the confidence interval.

William S. Goset (1876–1937) of the Guinness brewery in Dublin, Ireland ran into this problem. His experiments with hops and barley produced very few samples. Just replacing σ with s did not produce accurate results when he tried to calculate a confidence interval. He realized that he could not use a normal distribution for the calculation; he found that the actual distribution depends on the sample size. This problem led him to "discover" what is called the **Student's t-distribution**. The name comes from the fact that Gosset wrote under the pen name "A Student."

Up until the mid-1970s, some statisticians used the **normal distribution** approximation for large sample sizes and used the Student's t-distribution only for sample sizes of at most 30 observations.

If you draw a simple random sample of size n from a population with mean μ and unknown population standard deviation σ and calculate the t-score

$$t = \frac{\bar{x} - \mu}{\left(\frac{s}{\sqrt{n}}\right)} \quad (8.3.1)$$

then the t-scores follow a **Student's t-distribution with $n-1$ degrees of freedom**. The t-score has the same interpretation as the z-score. It measures how far in standard deviation units \bar{x} is from its mean μ . For each sample size n , there is a different Student's t-distribution.

The **degrees of freedom, $n-1$** , come from the calculation of the sample standard deviation s . Remember when we first calculated a sample standard deviation we divided the sum of the squared deviations by $n-1$, but we used n deviations (\bar{x} values) to calculate s . Because the sum of the deviations is zero, we can find the last deviation once we know the other $n-1$ deviations. The other $n-1$ deviations can change or vary freely. **We call the number $n-1$ the degrees of freedom (df)** in recognition that one is lost in the calculations. The effect of losing a degree of freedom is that the t-value increases and the confidence interval increases in width.

Properties of the Student's t-Distribution

- The graph for the Student's t-distribution is similar to the standard normal curve and at infinite degrees of freedom it is the normal distribution. You can confirm this by reading the bottom line at infinite degrees of freedom for a familiar level of confidence, e.g. at column 0.05, 95% level of confidence, we find the t-value of 1.96 at infinite degrees of freedom.
- The mean for the Student's t-distribution is zero and the distribution is symmetric about zero, again like the standard normal distribution.
- The Student's t-distribution has more probability in its tails than the standard normal distribution because the spread of the t-distribution is greater than the spread of the standard normal. So the graph of the Student's t-distribution will be thicker in the tails and shorter in the center than the graph of the standard normal distribution.
- The exact shape of the Student's t-distribution depends on the degrees of freedom. As the degrees of freedom increases, the graph of Student's t-distribution becomes more like the graph of the standard normal distribution.
- The underlying population of individual observations is assumed to be normally distributed with unknown population mean μ and unknown population standard deviation σ . This assumption comes from the Central Limit theorem because the individual observations in this case are the \bar{x} s of the sampling distribution. The size of the underlying population is generally not relevant unless it is very small. If it is normal then the assumption is met and doesn't need discussion.

A probability table for the Student's t-distribution is used to calculate t-values at various commonly-used levels of confidence. The table gives t-scores that correspond to the confidence level (column) and degrees of freedom (row). When using a t-table, note that some tables are formatted to show the confidence level in the column headings, while the column headings in some tables may

show only corresponding area in one or both tails. Notice that at the bottom the table will show the t -value for infinite degrees of freedom. Mathematically, as the degrees of freedom increase, the t distribution approaches the standard normal distribution. You can find familiar Z -values by looking in the relevant alpha column and reading value in the last row.

A Student's t table (Table 8.3.6) gives t -scores given the degrees of freedom and the right-tailed probability.

The Student's t distribution has one of the most desirable properties of the normal: it is symmetrical. What the Student's t distribution does is spread out the horizontal axis so it takes a larger number of standard deviations to capture the same amount of probability. In reality there are an infinite number of Student's t distributions, one for each adjustment to the sample size. As the sample size increases, the Student's t distribution become more and more like the normal distribution. When the sample size reaches 30 the normal distribution is usually substituted for the Student's t because they are so much alike. This relationship between the Student's t distribution and the normal distribution is shown in Figure 8.3.8.




Figure 8.3.1

This is another example of one distribution limiting another one, in this case the normal distribution is the limiting distribution of the Student's t when the degrees of freedom in the Student's t approaches infinity. This conclusion comes directly from the derivation of the Student's t distribution by Mr. Gosset. He recognized the problem as having few observations and no estimate of the population standard deviation. He was substituting the sample standard deviation and getting volatile results. He therefore created the Student's t distribution as a ratio of the normal distribution and Chi squared distribution. The Chi squared distribution is itself a ratio of two variances, in this case the sample variance and the unknown population variance. The Student's t distribution thus is tied to the normal distribution, but has degrees of freedom that come from those of the Chi squared distribution. The algebraic solution demonstrates this result.

Development of Student's t -distribution:

$$1. t = \frac{z}{\sqrt{\frac{\chi^2}{v}}}$$

Where Z is the standard normal distribution and χ^2 is the chi-squared distribution with v degrees of freedom.

$$2. t = \frac{\frac{(\bar{x} - \mu)}{\sigma}}{\sqrt{\frac{\frac{s^2}{(n-1)}}{\frac{\sigma^2}{(n-1)}}}}$$

by substitution, and thus Student's t with $v = n - 1$ degrees of freedom is:

$$3. t = \frac{\bar{x} - \mu}{\frac{s}{\sqrt{n}}}$$

Restating the formula for a confidence interval for the mean for cases when the sample size is smaller than 30 and we do not know the population standard deviation, σ :

$$\bar{x} - t_{\nu, \alpha} \left(\frac{s}{\sqrt{n}} \right) \leq \mu \leq \bar{x} + t_{\nu, \alpha} \left(\frac{s}{\sqrt{n}} \right)$$

Here the point estimate of the population standard deviation, s has been substituted for the population standard deviation, σ , and $t_{\nu, \alpha}$ has been substituted for Z_{α} . The Greek letter ν (pronounced nu) is placed in the general formula in recognition that there are many Student t_{ν} distributions, one for each sample size. ν is the symbol for the degrees of freedom of the distribution and depends on the size of the sample. Often df is used to abbreviate degrees of freedom. **For this type of problem**, the degrees of freedom is

$\nu = n - 1$, where n is the sample size. To look up a probability in the Student's t table we have to know the degrees of freedom in the problem.

Example 8.3.1

The average earnings per share (EPS) for 10 industrial stocks randomly selected from those listed on the Dow-Jones Industrial Average was found to be $\bar{X} = 1.85$ with a standard deviation of $s = 0.395$. Calculate a 99% confidence interval for the average EPS of all the industrials listed on the *DJIA*.

$$\bar{x} - t_{\nu, \alpha} \left(\frac{s}{\sqrt{n}} \right) \leq \mu \leq \bar{x} + t_{\nu, \alpha} \left(\frac{s}{\sqrt{n}} \right)$$

Answer

To help visualize the process of calculating a confident interval we draw the appropriate distribution for the problem. In this case this is the Student's t because we do not know the population standard deviation and the sample is small, less than 30.




Figure 8.3.2

To find the appropriate t -value requires two pieces of information, the level of confidence desired and the degrees of freedom. The question asked for a 99% confidence level. On the graph this is shown where $(1 - \alpha)$, the level of confidence, is in the unshaded area. The tails, thus, have .005 probability each, $\alpha/2$. The degrees of freedom for this type of problem is $n - 1 = 9$. From the Student's t table, at the row marked 9 and column marked .005, is the number of standard deviations to capture 99% of the probability, 3.2498. These are then placed on the graph remembering that the Student's t is symmetrical and so the t -value is both plus or minus on each side of the mean.

Inserting these values into the formula gives the result. These values can be placed on the graph to see the relationship between the distribution of the sample means, \bar{X} 's and the Student's t distribution.

$$\mu = \bar{X} \pm t_{\alpha/2, df=n-1} \frac{s}{\sqrt{n}} = 1.851 \pm 3.2498 \frac{0.395}{\sqrt{10}} = 1.8551 \pm 0.406$$

$$1.445 \leq \mu \leq 2.257$$

We state the formal conclusion as :

With 99% confidence level, the average *EPS* of all the industries listed at *DJIA* is from \$1.44 to \$2.26.

Exercise 8.3.2

You do a study of hypnotherapy to determine how effective it is in increasing the number of hours of sleep subjects get each night. You measure hours of sleep for 12 subjects with the following results. Construct a 95% confidence interval for the mean number of hours slept for the population (assumed normal) from which you took the data.

8.2; 9.1; 7.7; 8.6; 6.9; 11.2; 10.1; 9.9; 8.9; 9.2; 7.5; 10.5

This page titled [8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.2: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case](#) by [OpenStax](#) is licensed [CC BY 4.0](#).
Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.4: A Confidence Interval for A Population Proportion

During an election year, we see articles in the newspaper that state **confidence intervals** in terms of proportions or percentages. For example, a poll for a particular candidate running for president might show that the candidate has 40% of the vote within three percentage points (if the sample is large enough). Often, election polls are calculated with 95% confidence, so, the pollsters would be 95% confident that the true proportion of voters who favored the candidate would be between 0.37 and 0.43.

Investors in the stock market are interested in the true proportion of stocks that go up and down each week. Businesses that sell personal computers are interested in the proportion of households in the United States that own personal computers. Confidence intervals can be calculated for the true proportion of stocks that go up or down each week and for the true proportion of households in the United States that own personal computers.

The procedure to find the confidence interval for a population proportion is similar to that for the population mean, but the formulas are a bit different although conceptually identical. While the formulas are different, they are based upon the same mathematical foundation given to us by the Central Limit Theorem. Because of this we will see the same basic format using the same three pieces of information: the sample value of the parameter in question, the standard deviation of the relevant sampling distribution, and the number of standard deviations we need to have the confidence in our estimate that we desire.

How do you know you are dealing with a proportion problem? First, the underlying **distribution has a binary random variable and therefore is a binomial distribution**. (There is no mention of a mean or average.) If X is a binomial random variable, then $X \sim B(n, p)$ where n is the number of trials and p is the probability of a success. To form a sample proportion, take X , the random variable for the number of successes and divide it by n , the number of trials (or the sample size). The random variable P' (read "P prime") is the sample proportion,

$$P' = \frac{X}{n}$$

(Sometimes the random variable is denoted as \hat{P} , read "P hat".)

- P' = the **estimated proportion** of successes or sample proportion of successes (P' is a **point estimate** for p , the true population proportion, and thus q is the probability of a failure in any one trial.)
- x = the **number** of successes in the sample
- n = the size of the sample

The formula for the confidence interval for a population proportion follows the same format as that for an estimate of a population mean. Remembering the sampling distribution for the proportion from [Chapter 7](#), the standard deviation was found to be:

$$\sigma_{p'} = \sqrt{\frac{p(1-p)}{n}}$$

The confidence interval for a population proportion, therefore, becomes:

$$p = p' \pm \left[Z_{\left(\frac{\alpha}{2}\right)} \sqrt{\frac{p'(1-p')}{n}} \right]$$

$Z_{\left(\frac{\alpha}{2}\right)}$ is set according to our desired degree of confidence and $\sqrt{\frac{p'(1-p')}{n}}$ is the standard deviation of the sampling distribution.

The **sample proportions p' and q' are estimates of the unknown population proportions p and q** . The estimated proportions p' and q' are used because p and q are not known.

Remember that as p moves further from 0.5 the binomial distribution becomes less symmetrical. Because we are estimating the binomial with the symmetrical normal distribution the further away from symmetrical the binomial becomes the less confidence we have in the estimate.

This conclusion can be demonstrated through the following analysis. Proportions are based upon the binomial probability distribution. The possible outcomes are binary, either "success" or "failure". This gives rise to a proportion, meaning the percentage of the outcomes that are "successes". It was shown that the binomial distribution could be fully understood if we knew only the probability of a success in any one trial, called p . The mean and the standard deviation of the binomial were found to be:

$$\mu = np$$

$$\sigma = \sqrt{npq}$$

It was also shown that the binomial could be estimated by the normal distribution if BOTH np AND nq were greater than 5. From the discussion above, it was found that the standardizing formula for the binomial distribution is:

$$Z = \frac{p' - p}{\sqrt{\left(\frac{pq}{n}\right)}}$$

which is nothing more than a restatement of the general standardizing formula with appropriate substitutions for μ and σ from the binomial. We can use the standard normal distribution, the reason Z is in the equation, because the normal distribution is the limiting distribution of the binomial. This is another example of the Central Limit Theorem. We have already seen that the sampling distribution of means is normally distributed. Recall the extended discussion in [Chapter 7](#) concerning the sampling distribution of proportions and the conclusions of the Central Limit Theorem.

We can now manipulate this formula in just the same way we did for finding the confidence intervals for a mean, but to find the confidence interval for the binomial population parameter, p .

$$p' - Z_{\alpha} \sqrt{\frac{p'q'}{n}} \leq p \leq p' + Z_{\alpha} \sqrt{\frac{p'q'}{n}}$$

Where $p' = x/n$, the point estimate of p taken from the sample. Notice that p' has replaced p in the formula. This is because we do not know p , indeed, this is just what we are trying to estimate.

Unfortunately, there is no correction factor for cases where the sample size is small so np' and nq' must always be greater than 5 to develop an interval estimate for p .

Example 8.4.1

Suppose that a market research firm is hired to estimate the percent of adults living in a large city who have cell phones. Five hundred randomly selected adult residents in this city are surveyed to determine whether they have cell phones. Of the 500 people sampled, 421 responded yes - they own cell phones. Using a 95% confidence level, compute a confidence interval estimate for the true proportion of adult residents of this city who have cell phones.

Answer

The solution step-by-step.

Let X = the number of people in the sample who have cell phones. X is binomial: the random variable is binary, people either have a cell phone or they do not.

To calculate the confidence interval, we must find p' , q' .

$$n = 500$$

$$x = \text{the number of successes in the sample} = 421$$

$$p' = \frac{x}{n} = \frac{421}{500} = 0.842$$

$p' = 0.842$ is the sample proportion; this is the point estimate of the population proportion.

$$q' = 1 - p' = 1 - 0.842 = 0.158$$

Since the requested confidence level is $CL = 0.95$, then $\alpha = 1 - CL = 1 - 0.95 = 0.05$ ($\frac{\alpha}{2}$) = 0.025 .

$$\text{Then } z_{\frac{\alpha}{2}} = z_{0.025} = 1.96$$

This can be found using the Standard Normal probability table in [Table 8.4.6](#) This can also be found in the students t table at the 0.025 column and infinity degrees of freedom because at infinite degrees of freedom the students t distribution becomes the standard normal distribution, Z .

The confidence interval for the true binomial population proportion is

$$p' - Z_{\alpha} \sqrt{\frac{p'q'}{n}} \leq p \leq p' + Z_{\alpha} \sqrt{\frac{p'q'}{n}}$$

Substituting in the values from above we find the confidence interval is : $0.810 \leq p \leq 0.874$

Interpretation

We estimate with 95% confidence that between 81% and 87.4% of all adult residents of this city have cell phones.

Explanation of 95% Confidence Level

Ninety-five percent of the confidence intervals constructed in this way would contain the true value for the population proportion of all adult residents of this city who have cell phones.

Exercise 8.4.1

Suppose 250 randomly selected people are surveyed to determine if they own a tablet. Of the 250 surveyed, 98 reported owning a tablet. Using a 95% confidence level, compute a confidence interval estimate for the true proportion of people who own tablets.

Example 8.4.2

The Dundee Dog Training School has a larger than average proportion of clients who compete in competitive professional events. A confidence interval for the population proportion of dogs that compete in professional events from 150 different training schools is constructed. The lower limit is determined to be 0.08 and the upper limit is determined to be 0.16. Determine the level of confidence used to construct the interval of the population proportion of dogs that compete in professional events.

Answer

We begin with the formula for a confidence interval for a proportion because the random variable is binary; either the client competes in professional competitive dog events or they don't.

$$p = p' \pm \left[Z_{\left(\frac{\alpha}{2}\right)} \sqrt{\frac{p'(1-p')}{n}} \right]$$

Next we find the sample proportion:

$$p' = \frac{0.08 + 0.16}{2} = 0.12$$

The \pm that makes up the confidence interval is thus 0.04; $0.12 + 0.04 = 0.16$ and $0.12 - 0.04 = 0.08$ the boundaries of the confidence interval. Finally, we solve for Z .

$$\left[Z \cdot \sqrt{\frac{0.12(1-0.12)}{150}} \right] = 0.04, \text{ therefore } z = 1.51$$

And then look up the probability for 1.51 standard deviations on the standard normal table.

$$p(Z = 1.51) = 0.4345, p(Z) \cdot 2 = 0.8690 \text{ or } 86.90\%$$

Example 8.4.3

A financial officer for a company wants to estimate the percent of accounts receivable that are more than 30 days overdue. He surveys 500 accounts and finds that 300 are more than 30 days overdue. Compute a 90% confidence interval for the true percent of accounts receivable that are more than 30 days overdue, and interpret the confidence interval.

Answer

The solution is step-by-step:

Ninety percent of all confidence intervals constructed in this way contain the true value for the population percent of accounts receivable that are overdue 30 days.

Explanation of 90% Confidence Level

$$x = 300 \text{ and } n = 500$$

$$p' = \frac{x}{n} = \frac{300}{500} = 0.600$$

$$q' = 1 - p' = 1 - 0.600 = 0.400$$

Since confidence level = 0.90 then $\alpha = 1 - \text{confidence level} = (1 - 0.90) = 0.10$ ($\frac{\alpha}{2}$) = 0.05

$$Z_{\frac{\alpha}{2}} = Z_{0.05} = 1.645$$

This Z-value can be found using a standard normal probability table. The student's t-table can also be used by entering the table at the 0.05 column and reading at the line for infinite degrees of freedom. The t-distribution is the normal distribution at infinite degrees of freedom. This is a handy trick to remember in finding Z-values for commonly used levels of confidence. We use this formula for a confidence interval for a proportion:

$$p' - Z_{\alpha} \sqrt{\frac{p'q'}{n}} \leq p \leq p' + Z_{\alpha} \sqrt{\frac{p'q'}{n}}$$

Substituting in the values from above we find the confidence interval for the true binomial population proportion is $0.564 \leq p \leq 0.636$

Interpretation

We estimate with 90% confidence that the true percent of all accounts receivable overdue 30 days is between 56.4% and 63.6%. Alternate Wording: We estimate with 90% confidence that between 56.4% and 63.6% of ALL accounts are overdue 30 days.

Exercise 8.4.2

A student polls his school to see if students in the school district are for or against the new legislation regarding school uniforms. She surveys 600 students and finds that 480 are against the new legislation.

- Compute a 90% confidence interval for the true percent of students who are against the new legislation, and interpret the confidence interval.
- In a sample of 300 students, 68% said they own an iPod and a smart phone. Compute a 97% confidence interval for the true percent of students who own an iPod and a smartphone.

This page titled [8.4: A Confidence Interval for A Population Proportion](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.3: A Confidence Interval for A Population Proportion](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.5: Calculating the Sample Size n - Continuous and Binary Random Variables

Continuous Random Variables

Usually we have no control over the sample size of a data set. However, if we are able to set the sample size, as in cases where we are taking a survey, it is very helpful to know just how large it should be to provide the most information. Sampling can be very costly in both time and product. Simple telephone surveys will cost approximately \$30.00 each, for example, and some sampling requires the destruction of the product.

If we go back to our standardizing formula for the sampling distribution for means, we can see that it is possible to solve it for n . If we do this we have $(\bar{X} - \mu)$ in the denominator.

$$n = \frac{Z_{\alpha}^2 \sigma^2}{(\bar{X} - \mu)^2} = \frac{Z_{\alpha}^2 \sigma^2}{e^2}$$

Because we have not taken a sample yet we do not know any of the variables in the formula except that we can set Z_{α} to the level of confidence we desire just as we did when determining confidence intervals. If we set a predetermined acceptable error, or tolerance, for the difference between \bar{X} and μ , called e in the formula, we are much further in solving for the sample size n . We still do not know the population standard deviation, σ . In practice, a pre-survey is usually done which allows for fine tuning the questionnaire and will give a sample standard deviation that can be used. In other cases, previous information from other surveys may be used for σ in the formula. While crude, this method of determining the sample size may help in reducing cost significantly. It will be the actual data gathered that determines the inferences about the population, so caution in the sample size is appropriate calling for high levels of confidence and small sampling errors.

Binary Random Variables

What was done in cases when looking for the mean of a distribution can also be done when sampling to determine the population parameter p for proportions. Manipulation of the standardizing formula for proportions gives:

$$n = \frac{Z_{\alpha}^2 pq}{e^2}$$

where $e = (p' - p)$, and is the acceptable sampling error, or tolerance, for this application. This will be measured in percentage points.

In this case the very object of our search is in the formula, p , and of course q because $q = 1 - p$. This result occurs because the binomial distribution is a one parameter distribution. If we know p then we know the mean and the standard deviation. Therefore, p shows up in the standard deviation of the sampling distribution which is where we got this formula. If, in an abundance of caution, we substitute 0.5 for p we will draw the largest required sample size that will provide the level of confidence specified by Z_{α} and the tolerance we have selected. This is true because of all combinations of two fractions that add to one, the largest multiple is when each is 0.5. Without any other information concerning the population parameter p , this is the common practice. This may result in oversampling, but certainly not under sampling, thus, this is a cautious approach.

There is an interesting trade-off between the level of confidence and the sample size that shows up here when considering the cost of sampling. Table 8.5.1 shows the appropriate sample size at different levels of confidence and different level of the acceptable error, or tolerance.

Table 8.5.1

Required sample size (90%)	Required sample size (95%)	Tolerance level
1691	2401	2%
752	1067	3%
271	384	5%
68	96	10%

This table is designed to show the maximum sample size required at different levels of confidence given an assumed $p = 0.5$ and $q = 0.5$ as discussed above.

The acceptable error, called tolerance in the table, is measured in plus or minus values from the actual proportion. For example, an acceptable error of 5% means that if the sample proportion was found to be 26 percent, the conclusion would be that the actual population proportion is between 21 and 31 percent with a 90 percent level of confidence if a sample of 271 had been taken. Likewise, if the acceptable error was set at 2%, then the population proportion would be between 24 and 28 percent with a 90 percent level of confidence, but would require that the sample size be increased from 271 to 1,691. If we wished a higher level of confidence, we would require a larger sample size. Moving from a 90 percent level of confidence to a 95 percent level at a plus or minus 5% tolerance requires changing the sample size from 271 to 384. A very common sample size often seen reported in political surveys is 384. With the survey results it is frequently stated that the results are good to a plus or minus 5% level of “accuracy”.

Example 8.5.9

Suppose a mobile phone company wants to determine the current percentage of customers aged 50+ who use text messaging on their cell phones. How many customers aged 50+ should the company survey in order to be 90% confident that the estimated (sample) proportion is within three percentage points of the true population proportion of customers aged 50+ who use text messaging on their cell phones.

Answer

Solution 8.9

From the problem, we know that the acceptable error, e , is **0.03** ($3\%=0.03$) and $z_{\frac{\alpha}{2}} Z_{0.05} = 1.645$ because the confidence level is 90%. The acceptable error, e , is the difference between the actual population proportion p , and the sample proportion we expect to get from the sample.

However, in order to find n , we need to know the estimated (sample) proportion p' . Remember that $q' = 1 - p'$. But, we do not know p' yet. Since we multiply p' and q' together, we make them both equal to 0.5 because $p'q' = (0.5)(0.5) = 0.25$ results in the largest possible product. (Try other products: $(0.6)(0.4) = 0.24$; $(0.3)(0.7) = 0.21$; $(0.2)(0.8) = 0.16$ and so on). The largest possible product gives us the largest n . This gives us a large enough sample so that we can be 90% confident that we are within three percentage points of the true population proportion. To calculate the sample size n , use the formula and make the substitutions.

$$n = \frac{z^2 p' q'}{e^2} \text{ gives } n = \frac{1.645^2 (0.5)(0.5)}{0.03^2} = 751.7$$

Round the answer to the next higher value. The sample size should be 752 cell phone customers aged 50+ in order to be 90% confident that the estimated (sample) proportion is within three percentage points of the true population proportion of all customers aged 50+ who use text messaging on their cell phones.

Exercise 8.5.9

Suppose an internet marketing company wants to determine the current percentage of customers who click on ads on their smartphones. How many customers should the company survey in order to be 90% confident that the estimated proportion is within five percentage points of the true population proportion of customers who click on ads on their smartphones?

This page titled [8.5: Calculating the Sample Size n- Continuous and Binary Random Variables](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.4: Calculating the Sample Size n- Continuous and Binary Random Variables](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.6: Chapter Key Terms

Binomial Distribution

a discrete random variable (RV) which arises from Bernoulli trials; there are a fixed number, n , of independent trials. “Independent” means that the result of any trial (for example, trial 1) does not affect the results of the following trials, and all trials are conducted under the same conditions. Under these circumstances the binomial RV X is defined as the number of successes in n trials. The notation is: $X \sim B(\mathbf{n}, \mathbf{p})$. The mean is $\mu = np$ and the standard deviation is $\sigma = \sqrt{npq}$. The probability of exactly x successes in n trials is $P(X = x) = \binom{n}{x} p^x q^{n-x}$.

Confidence Interval (CI)

an interval estimate for an unknown population parameter. This depends on:

- the desired confidence level,
- information that is known about the distribution (for example, known standard deviation),
- the sample and its size.

Confidence Level (CL)

the percent expression for the probability that the confidence interval contains the true population parameter; for example, if the CL = 90%, then in 90 out of 100 samples the interval estimate will enclose the true population parameter.

Degrees of Freedom (df)

the number of objects in a sample that are free to vary

Error Bound for a Population Mean (EBM)

the margin of error; depends on the confidence level, sample size, and known or estimated population standard deviation.

Error Bound for a Population Proportion (EBP)

the margin of error; depends on the confidence level, the sample size, and the estimated (from the sample) proportion of successes.

Inferential Statistics

also called statistical inference or inductive statistics; this facet of statistics deals with estimating a population parameter based on a sample statistic. For example, if four out of the 100 calculators sampled are defective we might infer that four percent of the production is defective.

Normal Distribution

a continuous random variable (RV) with pdf $f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}$, where μ is the mean of the distribution and σ is the standard deviation, notation: $X \sim N(\mu, \sigma)$. If $\mu = 0$ and $\sigma = 1$, the RV is called **the standard normal distribution**.

Parameter

a numerical characteristic of a population

Point Estimate

a single number computed from a sample and used to estimate a population parameter

Standard Deviation

a number that is equal to the square root of the variance and measures how far data values are from their mean; notation: s for sample standard deviation and σ for population standard deviation

Student's t-Distribution

investigated and reported by William S. Gossett in 1908 and published under the pseudonym Student; the major characteristics of this random variable (RV) are:

- It is continuous and assumes any real values.

- The pdf is symmetrical about its mean of zero.
- It approaches the standard normal distribution as n get larger.
- There is a "family" of t-distributions: each representative of the family is completely defined by the number of degrees of freedom, which depends upon the application for which the t is being used.

This page titled [8.6: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.7: Chapter Key Terms](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.7: Chapter Review

8.2 A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case

In many cases, the researcher does not know the population standard deviation, σ , of the measure being studied. In these cases, it is common to use the sample standard deviation, s , as an estimate of σ . The normal distribution creates accurate confidence intervals when σ is known, but it is not as accurate when s is used as an estimate. In this case, the Student's t -distribution is much better. Define a t -score using the following formula:

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

The t -score follows the Student's t -distribution with $n-1$ degrees of freedom. The confidence interval under this distribution is calculated with $\bar{x} \pm \left(t_{\frac{\alpha}{2}} \right) \frac{s}{\sqrt{n}}$ where $t_{\frac{\alpha}{2}}$ is the t -score with area to the right equal to $\frac{\alpha}{2}$, s is the sample standard deviation, and n is the sample size. Use a table, calculator, or computer to find $t_{\frac{\alpha}{2}}$ for a given α .

8.3 A Confidence Interval for A Population Proportion

Some statistical measures, like many survey questions, measure qualitative rather than quantitative data. In this case, the population parameter being estimated is a proportion. It is possible to create a confidence interval for the true population proportion following procedures similar to those used in creating confidence intervals for population means. The formulas are slightly different, but they follow the same reasoning.

Let p' represent the sample proportion, x/n , where x represents the number of successes and n represents the sample size. Let $q' = 1 - p'$. Then the confidence interval for a population proportion is given by the following formula:

$$p' - Z_{\alpha} \sqrt{\frac{p'q'}{n}} \leq p \leq p' + Z_{\alpha} \sqrt{\frac{p'q'}{n}}$$

8.4 Calculating the Sample Size n : Continuous and Binary Random Variables

Sometimes researchers know in advance that they want to estimate a population mean within a specific margin of error for a given level of confidence. In that case, solve the relevant confidence interval formula for n to discover the size of the sample that is needed to achieve this goal:

$$n = \frac{Z_{\alpha}^2 \sigma^2}{(\bar{x} - \mu)^2}$$

If the random variable is binary then the formula for the appropriate sample size to maintain a particular level of confidence with a specific tolerance level is given by

$$n = \frac{Z_{\alpha}^2 pq}{e^2}$$

This page titled [8.7: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **8.10: Chapter Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.8: Chapter Formula Review

A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case

s = the standard deviation of sample values.

$t = \frac{\bar{x} - \mu}{\frac{s}{\sqrt{n}}}$ is the formula for the t-score which measures how far away a measure is from the population mean in the Student's t-distribution

$df = n - 1$; the degrees of freedom for a Student's t-distribution where n represents the size of the sample

$T \sim t_{df}$ the random variable, T , has a Student's t-distribution with df degrees of freedom

The general form for a confidence interval for a single mean, population standard deviation unknown, and sample size less than 30 Student's t is given by: $\bar{x} - t_{v,\alpha} \left(\frac{s}{\sqrt{n}} \right) \leq \mu \leq \bar{x} + t_{v,\alpha} \left(\frac{s}{\sqrt{n}} \right)$

A Confidence Interval for A Population Proportion

$p' = \frac{x}{n}$ where x represents the number of successes in a sample and n represents the sample size. The variable p' is the sample proportion and serves as the point estimate for the true population proportion.

$$q' = 1 - p'$$

The variable p' has a binomial distribution that can be approximated with the normal distribution shown here. The confidence interval for the true population proportion is given by the formula:

$$p' - Z_{\alpha} \sqrt{\frac{p'q'}{n}} \leq p \leq p' + Z_{\alpha} \sqrt{\frac{p'q'}{n}}$$

$n = \frac{Z_{\alpha}^2 p'q'}{e^2}$ provides the number of observations needed to sample to estimate the population proportion, p , with confidence $1 - \alpha$ and margin of error e . Where e = the acceptable difference between the actual population proportion and the sample proportion.

Calculating the Sample Size n : Continuous and Binary Random Variables

$n = \frac{Z^2 \sigma^2}{(\bar{x} - \mu)^2}$ = the formula used to determine the sample size (n) needed to achieve a desired margin of error at a given level of confidence for a continuous random variable

$n = \frac{Z_{\alpha}^2 pq}{e^2}$ = the formula used to determine the sample size if the random variable is binary

This page titled [8.8: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.5: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.9: Chapter Practice

8.2 A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case

Use the following information to answer the next five exercises. A hospital is trying to cut down on emergency room wait times. It is interested in the amount of time patients must wait before being called back to be examined. An investigation committee randomly surveyed 70 patients. The sample mean was 1.5 hours with a sample standard deviation of 0.5 hours.

1.

Identify the following:

1. Use the following information to answer the next six exercises: One hundred eight Americans were surveyed to determine the number of hours they spend watching television each month. It was revealed that they watched an average of 151 hours each month with a standard deviation of 32 hours. Assume that the underlying population distribution is normal.⁶

Identify the following:

1. Use the following information to answer the next 13 exercises: The data in Table 8.9.2 are the result of a random survey of 39 national flags (with replacement between picks) from various countries. We are interested in finding a confidence interval for the true mean number of colors on a national flag. Let X = the number of colors on a national flag.

X	Freq.
1	1
2	7
3	18
4	7
5	6

12.

Calculate the following:

1. Construct a 95% confidence interval for the true mean number of colors on national flags.¹⁷

How much area is in both tails (combined)?

18.

How much area is in each tail?

19.

Calculate the following:

1. Use the following information to answer the next two exercises: Marketing companies are interested in knowing the population percent of women who make the majority of household purchasing decisions.²⁵

When designing a study to determine this population proportion, what is the minimum number you would need to survey to be 90% confident that the population proportion is estimated to within 0.05?

26.

If it were later determined that it was important to be more than 90% confident and a new survey were commissioned, how would it affect the minimum number you need to survey? Why?

27.

Identify the following:

1. Use the following information to answer the next five exercises: Of 1,050 randomly selected adults, 360 identified themselves as manual laborers, 280 identified themselves as non-manual wage earners, 250 identified themselves as mid-level managers, and 160 identified themselves as executives. In the survey, 82% of manual laborers

preferred trucks, 62% of non-manual wage earners preferred trucks, 54% of mid-level managers preferred trucks, and 26% of executives preferred trucks.³²

We are interested in finding the 95% confidence interval for the percent of executives who prefer trucks. Define random variables X and p' in words.

33.

Which distribution should you use for this problem?

34.

Construct a 95% confidence interval. State the confidence interval, sketch the graph, and calculate the error bound.

35.

Suppose we want to lower the sampling error. What is one way to accomplish that?

36.

The sampling error given in the survey is $\pm 2\%$. Explain what the $\pm 2\%$ means.

37.

Define the random variable X in words.

38.

Define the random variable p' in words.

39.

Which distribution should you use for this problem?

40.

Construct a 90% confidence interval, and state the confidence interval and the error bound.

41.

What would happen to the confidence interval if the level of confidence were 95%?

Use the following information to answer the next 16 exercises: The Ice Chalet offers dozens of different beginning ice-skating classes. All of the class names are put into a bucket. The 5 P.M., Monday night, ages 8 to 12, beginning ice-skating class was picked. In that class were 64 girls and 16 boys. Suppose that we are interested in the true proportion of girls, ages 8 to 12, in all beginning ice-skating classes at the Ice Chalet. Assume that the children in the selected class are a random sample of the population.

42.

What is being counted?

43.

In words, define the random variable X .

44.

Calculate the following:

1. *Use the following information to answer the next five exercises:* The standard deviation of the weights of elephants is known to be approximately 15 pounds. We wish to construct a 95% confidence interval for the mean weight of newborn elephant calves. Fifty newborn elephants are weighed. The sample mean is 244 pounds. The sample standard deviation is 11 pounds.⁵⁸

Identify the following:

1. *Use the following information to answer the next seven exercises:* The U.S. Census Bureau conducts a study to determine the time needed to complete the short form. The Bureau surveys 200 people. The

sample mean is 8.2 minutes. There is a known standard deviation of 2.2 minutes. The population distribution is assumed to be normal.**63.**

Identify the following:

1. *Use the following information to answer the next ten exercises:* A sample of 20 heads of lettuce was selected. Assume that the population distribution of head weight is normal. The weight of each head of lettuce was then recorded. The mean weight was 2.2 pounds with a standard deviation of 0.1 pounds. The population standard deviation is known to be 0.2 pounds.**70.**

Identify the following:

1. *Use the following information to answer the next 14 exercises:* The mean age for all Foothill College students for a recent Fall term was 33.2. The population standard deviation has been pretty consistent at 15. Suppose that twenty-five Winter students were randomly selected. The mean age for the sample was 30.4. We are interested in the true mean age for Winter Foothill College students. Let X = the age of a Winter Foothill College student.**80.**

$$\bar{x} = \underline{\hspace{2cm}}$$

81.

$$n = \underline{\hspace{2cm}}$$

82.

$$\underline{\hspace{2cm}} = 15$$

83.

In words, define the random variable \bar{X} .

84.

What is \bar{x} estimating?

85.

Is σ_x known?

86.

As a result of your answer to **Exercise 8.9.83**, state the exact distribution to use when calculating the confidence interval.

87.

How much area is in both tails (combined)? $\alpha = \underline{\hspace{2cm}}$

88.

How much area is in each tail? $\frac{\alpha}{2} = \underline{\hspace{2cm}}$

89.

Identify the following specifications:

1. lower limit
2. upper limit
3. error bound

90.

The 95% confidence interval is: $\underline{\hspace{4cm}}$.

91.

Fill in the blanks on the graph with the areas, upper and lower limits of the confidence interval, and the sample mean.




Figure 8.9.12

92.

In one complete sentence, explain what the interval means.

93.

Using the same mean, standard deviation, and level of confidence, suppose that n were 69 instead of 25. Would the error bound become larger or smaller? How do you know?

94.

Using the same mean, standard deviation, and sample size, how would the error bound change if the confidence level were reduced to 90%? Why?

95.

Find the value of the sample size needed to if the confidence interval is 90% that the sample proportion and the population proportion are within 4% of each other. The sample proportion is 0.60. Note: Round all fractions up for n .

96.

Find the value of the sample size needed to if the confidence interval is 95% that the sample proportion and the population proportion are within 2% of each other. The sample proportion is 0.650. Note: Round all fractions up for n .

97.

Find the value of the sample size needed to if the confidence interval is 96% that the sample proportion and the population proportion are within 5% of each other. The sample proportion is 0.70. Note: Round all fractions up for n .

98.

Find the value of the sample size needed to if the confidence interval is 90% that the sample proportion and the population proportion are within 1% of each other. The sample proportion is 0.50. Note: Round all fractions up for n .

99.

Find the value of the sample size needed to if the confidence interval is 94% that the sample proportion and the population proportion are within 2% of each other. The sample proportion is 0.65. Note: Round all fractions up for n .

100.

Find the value of the sample size needed to if the confidence interval is 95% that the sample proportion and the population proportion are within 4% of each other. The sample proportion is 0.45. Note: Round all fractions up for n .

101.

Find the value of the sample size needed to if the confidence interval is 90% that the sample proportion and the population proportion are within 2% of each other. The sample proportion is 0.3.

Note: Round all fractions up for n .

This page titled [8.9: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **8.8: Chapter Practice** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.10: Chapter Homework

8.2 A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case 102.

In six packages of “The Flintstones® Real Fruit Snacks” there were five Bam-Bam snack pieces. The total number of snack pieces in the six bags was 68. We wish to calculate a 96% confidence interval for the population proportion of Bam-Bam snack pieces.

1. The FEC has reported financial information for 556 Leadership PACs that operating during the 2011–2012 election cycle. The following table shows the total receipts during this cycle for a random selection of 30 Leadership PACs.

Table 8.10.3

\$46,500.00	\$0	\$40,966.50	\$105,887.20	\$5,175.00
\$29,050.00	\$19,500.00	\$181,557.20	\$31,500.00	\$149,970.80
\$2,555,363.20	\$12,025.00	\$409,000.00	\$60,521.70	\$18,000.00
\$61,810.20	\$76,530.80	\$119,459.20	\$0	\$63,520.00
\$6,500.00	\$502,578.00	\$705,061.10	\$708,258.90	\$135,810.00
\$2,000.00	\$2,000.00	\$0	\$1,287,933.80	\$219,148.30

$$s = \$521,130.41$$

Use this sample data to construct a 95% confidence interval for the mean amount of money raised by all Leadership PACs during the 2011–2012 election cycle. Use the Student's t-distribution.

108.

Forbes magazine published data on the best small firms in 2012. These were firms that had been publicly traded for at least a year, have a stock price of at least \$5 per share, and have reported annual revenue between \$5 million and \$1 billion. The Table 8.10.4 shows the ages of the corporate CEOs for a random sample of these firms.

48	58	51	61	56
59	74	63	53	50
59	60	60	57	46
55	63	57	47	55
57	43	61	62	49
67	67	55	55	49

Table 8.4

Use this sample data to construct a 90% confidence interval for the mean age of CEO's for these top small firms. Use the Student's t-distribution.

109.

Unoccupied seats on flights cause airlines to lose revenue. Suppose a large airline wants to estimate its mean number of unoccupied seats per flight over the past year. To accomplish this, the records of 225 flights are randomly selected and the number of unoccupied seats is noted for each of the sampled flights. The sample mean is 11.6 seats and the sample standard deviation is 4.1 seats.

1. Use the following information to answer the next two exercises: A quality control specialist for a restaurant chain takes a random sample of size 12 to check the amount of soda served in the 16 oz. serving size. The sample mean is 13.30 with a sample standard deviation of 1.55. Assume the underlying population is normally distributed. **113.**

Find the 95% Confidence Interval for the true population mean for the amount of soda served.

- a. (12.42, 14.18)

- b. (12.32, 14.29)
- c. (12.50, 14.10)
- d. Impossible to determine

8.3 A Confidence Interval for A Population Proportion

114.

Insurance companies are interested in knowing the population percent of drivers who always buckle up before riding in a car.

1. When designing a study to determine this population proportion, what is the minimum number you would need to survey to be 95% confident that the population proportion is estimated to within 0.03?
2. If it were later determined that it was important to be more than 95% confident and a new survey was commissioned, how would that affect the minimum number you would need to survey? Why?

115.

Suppose that the insurance companies did do a survey. They randomly surveyed 400 drivers and found that 320 claimed they always buckle up. We are interested in the population proportion of drivers who claim they always buckle up.

1.
 - $x =$ _____
 - $n =$ _____
 - $p' =$ _____
2. Define the random variables X and P' , in words.
3. Which distribution should you use for this problem? Explain your choice.
4. Construct a 95% confidence interval for the population proportion who claim they always buckle up.
 - State the confidence interval.
 - Sketch the graph.
5. If this survey were done by telephone, list three difficulties the companies might have in obtaining random results.

116.

According to a recent survey of 1,200 people, 61% feel that the president is doing an acceptable job. We are interested in the population proportion of people who feel the president is doing an acceptable job.

1. Define the random variables X and P' in words.
2. Which distribution should you use for this problem? Explain your choice.
3. Construct a 90% confidence interval for the population proportion of people who feel the president is doing an acceptable job.
 - State the confidence interval.
 - Sketch the graph.

117.

An article regarding interracial dating and marriage recently appeared in the Washington Post. Of the 1,709 randomly selected adults, 315 identified themselves as Latinos, 323 identified themselves as blacks, 254 identified themselves as Asians, and 779 identified themselves as whites. In this survey, 86% of blacks said that they would welcome a white person into their families. Among Asians, 77% would welcome a white person into their families, 71% would welcome a Latino, and 66% would welcome a black person.

1. We are interested in finding the 95% confidence interval for the percent of all black adults who would welcome a white person into their families. Define the random variables X and P' , in words.
2. Which distribution should you use for this problem? Explain your choice.
3. Construct a 95% confidence interval.
 - State the confidence interval.
 - Sketch the graph.

118.

Refer to the information in Table 8.10.5 shows the total receipts from individuals for a random selection of 40 House candidates rounded to the nearest \$100. The standard deviation for this data to the nearest hundred is $\sigma = \$909,200$.

Table 8.10.5

\$3,600	\$1,243,900	\$10,900	\$385,200	\$581,500
\$7,400	\$2,900	\$400	\$3,714,500	\$632,500
\$391,000	\$467,400	\$56,800	\$5,800	\$405,200
\$733,200	\$8,000	\$468,700	\$75,200	\$41,000
\$13,300	\$9,500	\$953,800	\$1,113,500	\$1,109,300
\$353,900	\$986,100	\$88,600	\$378,200	\$13,200
\$3,800	\$745,100	\$5,800	\$3,072,100	\$1,626,700
\$512,900	\$2,309,200	\$6,600	\$202,400	\$15,800

1. Find the point estimate for the population mean.
2. Using 95% confidence, calculate the error bound.
3. Create a 95% confidence interval for the mean total individual contributions.
4. Interpret the confidence interval in the context of the problem.

137.

The American Community Survey (ACS), part of the United States Census Bureau, conducts a yearly census similar to the one taken every ten years, but with a smaller percentage of participants. The most recent survey estimates with 90% confidence that the mean household income in the U.S. falls between \$69,720 and \$69,922. Find the point estimate for mean U.S. household income and the error bound for mean U.S. household income.

138.

The average height of young adult males has a normal distribution with standard deviation of 2.5 inches. You want to estimate the mean height of students at your college or university to within one inch with 93% confidence. How many male students must you measure?

139.

If the confidence interval is change to a higher probability, would this cause a lower, or a higher, minimum sample size?

140.

If the tolerance is reduced by half, how would this affect the minimum sample size?

141.

If the value of p is reduced, would this necessarily reduce the sample size needed?

142.

Is it acceptable to use a higher sample size than the one calculated by $\frac{z^2 pq}{e^2}$?

143.

A company has been running an assembly line with 97.42% of the products made being acceptable. Then, a critical piece broke down. After the repairs the decision was made to see if the number of defective products made was still close enough to the long standing production quality. Samples of 500 pieces were selected at random, and the defective rate was found to be 0.025%.

1. Is this sample size adequate to claim the company is checking within the 90% confidence interval?
2. The 95% confidence interval?

This page titled [8.10: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **8.6: Chapter Homework** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.11: Chapter References

A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size

- “American Fact Finder.” U.S. Census Bureau. Available online at <http://factfinder2.census.gov/faces/...html?refresh=t> (accessed July 2, 2013).
- “Disclosure Data Catalog: Candidate Summary Report 2012.” U.S. Federal Election Commission. Available online at www.fec.gov/data/index.jsp (accessed July 2, 2013).
- “Headcount Enrollment Trends by Student Demographics Ten-Year Fall Trends to Most Recently Completed Fall.” Foothill De Anza Community College District. Available online at research.fhda.edu/factbook/FH...phicTrends.htm (accessed September 30, 2013).
- Kuczumarski, Robert J., Cynthia L. Ogden, Shumei S. Guo, Laurence M. Grummer-Strawn, Katherine M. Flegal, Zugu Mei, Rong Wei, Lester R. Curtin, Alex F. Roche, Clifford L. Johnson. “2000 CDC Growth Charts for the United States: Methods and Development.” Centers for Disease Control and Prevention. Available online at <http://www.cdc.gov/growthcharts/2000...thchart-us.pdf> (accessed July 2, 2013).
- La, Lynn, Kent German. “Cell Phone Radiation Levels.” c|net part of CBX Interactive Inc. Available online at <http://reviews.cnet.com/cell-phone-radiation-levels/> (accessed July 2, 2013).
- “Mean Income in the Past 12 Months (in 2011 Inflation-Adjusted Dollars): 2011 American Community Survey 1-Year Estimates.” American Fact Finder, U.S. Census Bureau. Available online at <http://factfinder2.census.gov/faces/...prodType=table> (accessed July 2, 2013).
- “Metadata Description of Candidate Summary File.” U.S. Federal Election Commission. Available online at www.fec.gov/finance/disclosure...esummary.shtml (accessed July 2, 2013).
- “National Health and Nutrition Examination Survey.” Centers for Disease Control and Prevention. Available online at <http://www.cdc.gov/nchs/nhanes.htm> (accessed July 2, 2013).

A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case

- “America’s Best Small Companies.” Forbes, 2013. Available online at <http://www.forbes.com/best-small-companies/list/> (accessed July 2, 2013).
- Data from *Microsoft Bookshelf*.
- Data from <http://www.businessweek.com/>.
- Data from <http://www.forbes.com/>.
- “Disclosure Data Catalog: Leadership PAC and Sponsors Report, 2012.” Federal Election Commission. Available online at www.fec.gov/data/index.jsp (accessed July 2, 2013).
- “Human Toxome Project: Mapping the Pollution in People.” Environmental Working Group. Available online at www.ewg.org/sites/humantoxome...tero%2Fnewborn (accessed July 2, 2013).
- “Metadata Description of Leadership PAC List.” Federal Election Commission. Available online at www.fec.gov/finance/disclosure...pPacList.shtml (accessed July 2, 2013).

A Confidence Interval for A Population Proportion

- Jensen, Tom. “Democrats, Republicans Divided on Opinion of Music Icons.” Public Policy Polling. Available online at www.publicpolicypolling.com/Day2MusicPoll.pdf (accessed July 2, 2013).
- Madden, Mary, Amanda Lenhart, Sandra Coresi, Urs Gasser, Maeve Duggan, Aaron Smith, and Meredith Beaton. “Teens, Social Media, and Privacy.” PewInternet, 2013. Available online at www.pewinternet.org/Reports/2...d-Privacy.aspx (accessed July 2, 2013).
- Prince Survey Research Associates International. “2013 Teen and Privacy Management Survey.” Pew Research Center: Internet and American Life Project. Available online at www.pewinternet.org/~media/...al%20Media.pdf (accessed July 2, 2013).
- Saad, Lydia. “Three in Four U.S. Workers Plan to Work Past Retirement Age: Slightly more say they will do this by choice rather than necessity.” Gallup® Economy, 2013. Available online at <http://www.gallup.com/poll/162758/th...ement-age.aspx> (accessed July 2, 2013).
- The Field Poll. Available online at field.com/fieldpollonline/subscribers/ (accessed July 2, 2013).
- Zogby. “New SUNYIT/Zogby Analytics Poll: Few Americans Worry about Emergency Situations Occurring in Their Community; Only one in three have an Emergency Plan; 70% Support Infrastructure ‘Investment’ for National Security.” Zogby Analytics, 2013. Available online at <http://www.zogbyanalytics.com/news/2...analytics-poll> (accessed July 2, 2013).

- “52% Say Big-Time College Athletics Corrupt Education Process.” Rasmussen Reports, 2013. Available online at http://www.rasmussenreports.com/publ...cation_process (accessed July 2, 2013).

This page titled [8.11: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [8.9: Chapter References](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

8.12: Chapter Solution (Practice + Homework)




Figure 8.12.23

This page titled [8.12: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

CHAPTER OVERVIEW

9: Hypothesis Testing with One Sample

- [9.1: Introduction to Hypothesis Testing](#)
- [9.2: Null and Alternative Hypotheses](#)
- [9.3: Outcomes and the Type I and Type II Errors](#)
- [9.4: Distribution Needed for Hypothesis Testing](#)
- [9.5: Full Hypothesis Test Examples](#)
- [9.6: Chapter Key Terms](#)
- [9.7: Chapter Review](#)
- [9.8: Chapter Formula Review](#)
- [9.9: Chapter Practice](#)
- [9.10: Chapter Homework](#)
- [9.11: Chapter References](#)
- [9.12: Chapter Solution \(Practice + Homework\)](#)

This page titled [9: Hypothesis Testing with One Sample](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

9.1: Introduction to Hypothesis Testing

Now we are down to the bread and butter work of the statistician: developing and testing hypotheses. It is important to put this material in a broader context so that the method by which a hypothesis is formed is understood completely. Using textbook examples often clouds the real source of statistical hypotheses.

Statistical testing is part of a much larger process known as the scientific method. This method was developed more than two centuries ago as the accepted way that new knowledge could be created. Until then, and unfortunately even today, among some, "knowledge" could be created simply by some authority saying something was so, *ipso dicta*. Superstition and conspiracy theories were (are?) accepted uncritically.



Figure 9.1.1 You can use a hypothesis test to decide if a dog breeder's claim that every Dalmatian has 35 spots is statistically sound. (Credit: Robert Neff)

The scientific method, briefly, states that only by following a careful and specific process can some assertion be included in the accepted body of knowledge. This process begins with a set of assumptions upon which a theory, sometimes called a model, is built. This theory, if it has any validity, will lead to predictions; what we call hypotheses.

As an example, in Microeconomics the theory of consumer choice begins with certain assumption concerning human behavior. From these assumptions a theory of how consumers make choices using indifference curves and the budget line. This theory gave rise to a very important prediction, namely, that there was an inverse relationship between price and quantity demanded. This relationship was known as the demand curve. The negative slope of the demand curve is really just a prediction, or a hypothesis, that can be tested with statistical tools.

Unless hundreds and hundreds of statistical tests of this hypothesis had not confirmed this relationship, the so-called Law of Demand would have been discarded years ago. This is the role of statistics, to test the hypotheses of various theories to determine if they should be admitted into the accepted body of knowledge; how we understand our world. Once admitted, however, they may be later discarded if new theories come along that make better predictions.

Not long ago two scientists claimed that they could get more energy out of a process than was put in. This caused a tremendous stir for obvious reasons. They were on the cover of *Time* and were offered extravagant sums to bring their research work to private industry and any number of universities. It was not long until their work was subjected to the rigorous tests of the scientific method and found to be a failure. No other lab could replicate their findings. Consequently they have sunk into obscurity and their theory discarded. It may surface again when someone can pass the tests of the hypotheses required by the scientific method, but until then it is just a curiosity. Many pure frauds have been attempted over time, but most have been found out by applying the process of the scientific method.

This discussion is meant to show just where in this process statistics falls. Statistics and statisticians are not necessarily in the business of developing theories, but in the business of testing others' theories. Hypotheses come from these theories based upon an explicit set of assumptions and sound logic. The hypothesis comes first, before any data are gathered. Data do not create hypotheses; they are used to test them. If we bear this in mind as we study this section the process of forming and testing hypotheses will make more sense.

One job of a statistician is to make statistical inferences about populations based on samples taken from the population. **Confidence intervals** are one way to estimate a population parameter. Another way to make a statistical inference is to make a decision about the value of a specific parameter. For instance, a car dealer advertises that its new small truck gets 35 miles per gallon, on average. A tutoring service claims that its method of tutoring helps 90% of its students get an A or a B. A company says that women managers in their company earn an average of \$60,000 per year.

A statistician will make a decision about these claims. This process is called "**hypothesis testing**." A hypothesis test involves collecting data from a sample and evaluating the data. Then, the statistician makes a decision as to whether or not there is sufficient evidence, based upon analyses of the data, to reject the null hypothesis.

In this chapter, you will conduct hypothesis tests on single means and single proportions. You will also learn about the errors associated with these tests.

This page titled [9.1: Introduction to Hypothesis Testing](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [9.0: Introduction to Hypothesis Testing](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.2: Null and Alternative Hypotheses

The actual test begins by considering two **hypotheses**. They are called the **null hypothesis** and the **alternative hypothesis**. These hypotheses contain opposing viewpoints.

- H_0 : **The null hypothesis**: It is a statement of no difference between a sample mean or proportion and a population mean or proportion. In other words, the difference equals 0. This can often be considered the status quo and as a result if you cannot accept the null it requires some action.
- H_a : **The alternative hypothesis**: It is a claim about the population that is contradictory to H_0 and what we conclude when we cannot accept H_0 . The alternative hypothesis is the contender and must win with significant evidence to overthrow the status quo. This concept is sometimes referred to the tyranny of the status quo because as we will see later, to overthrow the null hypothesis takes usually 90 or greater confidence that this is the proper decision.

Since the null and alternative hypotheses are contradictory, you must examine evidence to decide if you have enough evidence to reject the null hypothesis or not. The evidence is in the form of sample data.

After you have determined which hypothesis the sample supports, you make a **decision**. There are two options for a decision. They are "cannot accept H_0 " if the sample information favors the alternative hypothesis or "do not reject H_0 " or "decline to reject H_0 " if the sample information is insufficient to reject the null hypothesis. These conclusions are all based upon a level of probability, a significance level, that is set by the analyst.

Table 9.1 presents the various hypotheses in the relevant pairs. For example, if the null hypothesis is equal to some value, the alternative has to be not equal to that value.

Table 9.1

H_0	H_a
equal (=)	not equal (\neq)
greater than or equal to (\geq)	less than ($<$)
less than or equal to (\leq)	more than ($>$)

Note

As a mathematical convention H_0 always has a symbol with an equal in it. H_a never has a symbol with an equal in it. The choice of symbol depends on the wording of the hypothesis test.

Example 9.1

H_0 : No more than 30% of the registered voters in Santa Clara County voted in the primary election. $p \leq 30$
 H_a : More than 30% of the registered voters in Santa Clara County voted in the primary election. $p > 30$

Example 9.2

We want to test whether the mean GPA of students in American colleges is different from 2.0 (out of 4.0). The null and alternative hypotheses are:

$H_0 : \mu = 2.0$
 $H_a : \mu \neq 2.0$

Example 9.3

We want to test if college students take less than five years to graduate from college, on the average. The null and alternative hypotheses are:

$H_0 : \mu \geq 5$
 $H_a : \mu < 5$

This page titled [9.2: Null and Alternative Hypotheses](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **9.1: Null and Alternative Hypotheses** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.3: Outcomes and the Type I and Type II Errors

When you perform a hypothesis test, there are four possible outcomes depending on the actual truth (or falseness) of the null hypothesis H_0 and the decision to reject or not. The outcomes are summarized in the following table:

Table 9.2

Statistical Decision	H_0 is actually...	
	True	False
Cannot accept H_0	Type I error	Correct outcome
Cannot reject H_0	Correct outcome	Type II error

The four possible outcomes in the table are:

1. The decision is **cannot reject H_0** when **H_0 is true (correct decision)**.
2. The decision is **cannot accept H_0** when **H_0 is true** (incorrect decision known as a **Type I error**). This case is described as "rejecting a good null". As we will see later, it is this type of error that we will guard against by setting the probability of making such an error. The goal is to NOT take an action that is an error.
3. The decision is **cannot reject H_0** when, in fact, **H_0 is false** (incorrect decision known as a **Type II error**). This is called "accepting a false null". In this situation you have allowed the status quo to remain in force when it should be overturned. As we will see, the null hypothesis has the advantage in competition with the alternative.
4. The decision is **cannot accept H_0** when **H_0 is false (correct decision)**.

Each of the errors occurs with a particular probability. The Greek letters α and β represent the probabilities.

- α = probability of a Type I error = **P(Type I error)** = probability of rejecting the null hypothesis when the null hypothesis is true: rejecting a good null.
- β = probability of a Type II error = **P(Type II error)** = probability of not rejecting the null hypothesis when the null hypothesis is false. $(1 - \beta)$ is called the **Power of the Test**.

α and β should be as small as possible because they are probabilities of errors.

Statistics allows us to set the probability that we are making a Type I error. The probability of making a Type I error is α . Recall that the confidence intervals in the last unit were set by choosing a value called Z_α (or t_α) and the alpha value determined the confidence level of the estimate because it was the probability of the interval failing to capture the true mean (or proportion parameter p). This alpha and that one are the same.

The easiest way to see the relationship between the alpha error and the level of confidence is with the following figure.




Figure 9.2

In the center of Figure 9.2 is a normally distributed sampling distribution marked H_0 . This is a sampling distribution of \bar{X} and by the Central Limit Theorem it is normally distributed. The distribution in the center is marked H_0 and represents the distribution for the null hypotheses $H_0: \mu = 100$. This is the value that is being tested. The formal statements of the null and alternative hypotheses are listed below the figure.

The distributions on either side of the H_0 distribution represent distributions that would be true if H_0 is false, under the alternative hypothesis listed as H_a . We do not know which is true, and will never know. There are, in fact, an infinite number of distributions

from which the data could have been drawn if H_a is true, but only two of them are on Figure 9.2 representing all of the others.

To test a hypothesis we take a sample from the population and determine if it could have come from the hypothesized distribution with an acceptable level of significance. This level of significance is the alpha error and is marked on Figure 9.2 as the shaded areas in each tail of the H_0 distribution. (Each area is actually $\alpha/2$ because the distribution is symmetrical and the alternative hypothesis allows for the possibility for the value to be either greater than or less than the hypothesized value--called a two-tailed test).

If the sample mean marked as \bar{X}_1 is in the tail of the distribution of H_0 , we conclude that the probability that it could have come from the H_0 distribution is less than alpha. We consequently state, "the null hypothesis cannot be accepted with (α) level of significance". The truth **may** be that this \bar{X}_1 did come from the H_0 distribution, but from out in the tail. If this is so then we have falsely rejected a true null hypothesis and have made a Type I error. What statistics has done is provide an estimate about what we know, and what we control, and that is the probability of us being wrong, α .

We can also see in Figure 9.2 that the sample mean could be really from an H_a distribution, but within the boundary set by the alpha level. Such a case is marked as \bar{X}_2 . There is a probability that \bar{X}_2 actually came from H_a but shows up in the range of H_0 between the two tails. This probability is the beta error, the probability of accepting a false null.

Our problem is that we can only set the alpha error because there are an infinite number of alternative distributions from which the mean could have come that are not equal to H_0 . As a result, the statistician places the burden of proof on the alternative hypothesis. That is, we will not reject a null hypothesis unless there is a greater than 90, or 95, or even 99 percent probability that the null is false: the burden of proof lies with the alternative hypothesis. This is why we called this the tyranny of the status quo earlier.

By way of example, the American judicial system begins with the concept that a defendant is "presumed innocent". This is the status quo and is the null hypothesis. The judge will tell the jury that they can not find the defendant guilty unless the evidence indicates guilt beyond a "reasonable doubt" which is usually defined in criminal cases as 95% certainty of guilt. If the jury cannot accept the null, innocent, then action will be taken, jail time. The burden of proof always lies with the alternative hypothesis. (In civil cases, the jury needs only to be more than 50% certain of wrongdoing to find culpability, called "a preponderance of the evidence").

The example above was for a test of a mean, but the same logic applies to tests of hypotheses for all statistical parameters one may wish to test.

The following are examples of Type I and Type II errors.

Example 9.4

Suppose the null hypothesis, H_0 , is: Frank's rock climbing equipment is safe.

Type I error: Frank thinks that his rock climbing equipment may not be safe when, in fact, it really is safe.

Type II error: Frank thinks that his rock climbing equipment may be safe when, in fact, it is not safe.

$\alpha =$ **probability** that Frank thinks his rock climbing equipment may not be safe when, in fact, it really is safe. $\beta =$ **probability** that Frank thinks his rock climbing equipment may be safe when, in fact, it is not safe.

Notice that, in this case, the error with the greater consequence is the Type II error. (If Frank thinks his rock climbing equipment is safe, he will go ahead and use it.)

This is a situation described as "accepting a false null".

Example 9.5

Suppose the null hypothesis, H_0 , is: The victim of an automobile accident is alive when he arrives at the emergency room of a hospital. This is the status quo and requires no action if it is true. If the null hypothesis cannot be accepted then action is required and the hospital will begin appropriate procedures.

Type I error: The emergency crew thinks that the victim is dead when, in fact, the victim is alive. **Type II error:** The emergency crew does not know if the victim is alive when, in fact, the victim is dead.

$\alpha =$ **probability** that the emergency crew thinks the victim is dead when, in fact, he is really alive = P(Type I error). $\beta =$ **probability** that the emergency crew does not know if the victim is alive when, in fact, the victim is dead = P(Type II error).

The error with the greater consequence is the Type I error. (If the emergency crew thinks the victim is dead, they will not treat him.)

Exercise 9.5

Suppose the null hypothesis, H_0 , is: a patient is not sick. Which type of error has the greater consequence, Type I or Type II?

Example 9.6

It's a Boy Genetic Labs claim to be able to increase the likelihood that a pregnancy will result in a boy being born. Statisticians want to test the claim. Suppose that the null hypothesis, H_0 , is: It's a Boy Genetic Labs has no effect on gender outcome. The status quo is that the claim is false. The burden of proof always falls to the person making the claim, in this case the Genetics Lab.

Type I error: This results when a true null hypothesis is rejected. In the context of this scenario, we would state that we believe that It's a Boy Genetic Labs influences the gender outcome, when in fact it has no effect. The probability of this error occurring is denoted by the Greek letter alpha, α .

Type II error: This results when we fail to reject a false null hypothesis. In context, we would state that It's a Boy Genetic Labs does not influence the gender outcome of a pregnancy when, in fact, it does. The probability of this error occurring is denoted by the Greek letter beta, β .

The error of greater consequence would be the Type I error since couples would use the It's a Boy Genetic Labs product in hopes of increasing the chances of having a boy.

Exercise 9.6

"Red tide" is a bloom of poison-producing algae—a few different species of a class of plankton called dinoflagellates. When the weather and water conditions cause these blooms, shellfish such as clams living in the area develop dangerous levels of a paralysis-inducing toxin. In Massachusetts, the Division of Marine Fisheries (DMF) monitors levels of the toxin in shellfish by regular sampling of shellfish along the coastline. If the mean level of toxin in clams exceeds 800 μg (micrograms) of toxin per kg of clam meat in any area, clam harvesting is banned there until the bloom is over and levels of toxin in clams subside. Describe both a Type I and a Type II error in this context, and state which error has the greater consequence.

Example 9.7

A certain experimental drug claims a cure rate of at least 75% for males with prostate cancer. Describe both the Type I and Type II errors in context. Which error is the more serious?

Type I: A cancer patient believes the cure rate for the drug is less than 75% when it actually is at least 75%.

Type II: A cancer patient believes the experimental drug has at least a 75% cure rate when it has a cure rate that is less than 75%.

In this scenario, the Type II error contains the more severe consequence. If a patient believes the drug works at least 75% of the time, this most likely will influence the patient's (and doctor's) choice about whether to use the drug as a treatment option.

This page titled [9.3: Outcomes and the Type I and Type II Errors](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [9.2: Outcomes and the Type I and Type II Errors](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.4: Distribution Needed for Hypothesis Testing

Earlier, we discussed sampling distributions. Particular distributions are associated with hypothesis testing. We will perform hypotheses tests of a population mean using a normal distribution or a Student's t -distribution. (Remember, use a Student's t -distribution when the population standard deviation is unknown and the sample size is small, where small is considered to be less than 30 observations.) We perform tests of a population proportion using a normal distribution when we can assume that the distribution is normally distributed. We consider this to be true if the sample proportion, p' , times the sample size is greater than 5 and $1 - p'$ times the sample size is also greater than 5. This is the same rule of thumb we used when developing the formula for the confidence interval for a population proportion.

Hypothesis Test for the Mean

Going back to the standardizing formula we can derive the **test statistic** for testing hypotheses concerning means.

$$Z_c = \frac{\bar{x} - \mu_0}{\sigma / \sqrt{n}}$$

The standardizing formula can not be solved as it is because we do not have μ , the population mean. However, if we substitute in the hypothesized value of the mean, μ_0 in the formula as above, we can compute a Z value. This is the test statistic for a test of hypothesis for a mean and is presented in Figure 9.3. We interpret this Z value as the associated probability that a sample with a sample mean of \bar{X} could have come from a distribution with a population mean of H_0 and we call this Z value Z_c for "calculated". Figure 9.3 and Figure 9.4 show this process.




Figure 9.3

In Figure 9.3 two of the three possible outcomes are presented. \bar{X}_1 and \bar{X}_3 are in the tails of the hypothesized distribution of H_0 . Notice that the horizontal axis in the top panel is labeled \bar{X} 's. This is the same theoretical distribution of \bar{X} 's, the sampling distribution, that the Central Limit Theorem tells us is normally distributed. This is why we can draw it with this shape. The horizontal axis of the bottom panel is labeled Z and is the standard normal distribution. $Z_{\frac{\alpha}{2}}$ and $-Z_{\frac{\alpha}{2}}$, called the **critical values**, are marked on the bottom panel as the Z values associated with the probability the analyst has set as the level of significance in the test, (α). The probabilities in the tails of both panels are, therefore, the same.

Notice that for each \bar{X} there is an associated Z_c , called the calculated Z , that comes from solving the equation above. This calculated Z is nothing more than the number of standard deviations that the **hypothesized** mean is from the sample mean. If the sample mean falls "too many" standard deviations from the hypothesized mean we conclude that the **sample** mean could not have come from the distribution with the hypothesized mean, given our pre-set required level of significance. It **could** have come from H_0 , but it is deemed just too unlikely. In Figure 9.3 both \bar{X}_1 and \bar{X}_3 are in the tails of the distribution. They are deemed "too far" from the hypothesized value of the mean given the chosen level of alpha. If in fact this sample mean it did come from H_0 , but from

in the tail, we have made a Type I error: we have rejected a good null. Our only real comfort is that we know the probability of making such an error, α , and we can control the size of α .

Figure 9.4 shows the third possibility for the location of the sample mean, \bar{x} . Here the sample mean is within the two critical values. That is, within the probability of $(1 - \alpha)$ and we cannot reject the null hypothesis.




Figure 9.4

This gives us the decision rule for testing a hypothesis for a two-tailed test:

Table 9.3

Decision rule: two-tail test
If $ Z_c < Z_{\frac{\alpha}{2}}$: then do not REJECT H_0
If $ Z_c > Z_{\frac{\alpha}{2}}$: then REJECT H_0

This rule will always be the same no matter what hypothesis we are testing or what formulas we are using to make the test. The only change will be to change the Z_c to the appropriate symbol for the test statistic for the parameter being tested. Stating the decision rule another way: if the sample mean is unlikely to have come from the distribution with the hypothesized mean we cannot accept the null hypothesis. Here we define "unlikely" as having a probability less than alpha of occurring.

P-value Approach

An alternative decision rule can be developed by calculating the probability that a sample mean could be found that would give a test statistic larger than the test statistic found from the current sample data assuming that the null hypothesis is true. Here the notion of "likely" and "unlikely" is defined by the probability of drawing a sample with a mean from a population with the hypothesized mean that is either larger or smaller than that found in the sample data. Simply stated, the p -value approach compares the desired significance level, α , to the p -value which is the probability of drawing a sample mean further from the hypothesized value than the actual sample mean. A large p -value calculated from the data indicates that we should not reject the **null hypothesis**. The smaller the p -value, the more unlikely the outcome, and the stronger the evidence is against the null hypothesis. We would reject the null hypothesis if the evidence is strongly against it. The relationship between the decision rule of comparing the calculated test statistics, Z_c , and the Critical Value, Z_{α} , and using the p -value can be seen in Figure 9.5.




Figure 9.5

The calculated value of the test statistic is Z_c in this example and is marked on the bottom graph of the standard normal distribution because it is a Z value. In this case the calculated value is in the tail and thus we cannot accept the null hypothesis, the associated \bar{X} is just too unusually large to believe that it came from the distribution with a mean of μ_0 with a significance level of α .

If we use the p -value decision rule we need one more step. We need to find in the standard normal table the probability associated with the calculated test statistic, Z_c . We then compare that to the α associated with our selected level of confidence. In Figure 9.5 we see that the p -value is less than α and therefore we cannot accept the null. We know that the p -value is less than α because the area under the p -value is smaller than $\alpha/2$. It is important to note that two researchers drawing randomly from the same population may find two different p -values from their samples. This occurs because the p -value is calculated as the probability in the tail beyond the sample mean assuming that the null hypothesis is correct. Because the sample means will in all likelihood be different this will create two different p -values. Nevertheless, the conclusions as to the null hypothesis should be different with only the level of probability of α .

Here is a systematic way to make a decision of whether you cannot accept or cannot reject a null **hypothesis** if using the **p-value** and a **preset or preconceived** α (the "**significance level**"). A preset α is the probability of a **Type I** error (rejecting the null hypothesis when the null hypothesis is true). It may or may not be given to you at the beginning of the problem. In any case, the value of α is the decision of the analyst. When you make a decision to reject or not reject H_0 , do as follows:

- If $\alpha > p$ -value, cannot accept H_0 . The results of the sample data are significant. There is sufficient evidence to conclude that H_0 is an incorrect belief and that the **alternative hypothesis**, H_a , may be correct.
- If $\alpha \leq p$ -value, cannot reject H_0 . The results of the sample data are not significant. There is not sufficient evidence to conclude that the alternative hypothesis, H_a , may be correct. In this case the status quo stands.
- When you "cannot reject H_0 ", it does not mean that you should believe that H_0 is true. It simply means that the sample data have **failed** to provide sufficient evidence to cast serious doubt about the truthfulness of H_0 . Remember that the null is the status quo and it takes high probability to overthrow the status quo. This bias in favor of the null hypothesis is what gives rise to the statement "tyranny of the status quo" when discussing hypothesis testing and the scientific method.

Both decision rules will result in the same decision and it is a matter of preference which one is used.

One and Two-tailed Tests

The discussion of Figure 9.3-Figure 9.5 was based on the null and alternative hypothesis presented in Figure 9.3. This was called a two-tailed test because the alternative hypothesis allowed that the mean could have come from a population which was either larger or smaller than the hypothesized mean in the null hypothesis. This could be seen by the statement of the alternative hypothesis as $\mu \neq 100$, in this example.

It may be that the analyst has no concern about the value being "too" high or "too" low from the hypothesized value. If this is the case, it becomes a one-tailed test and all of the alpha probability is placed in just one tail and not split into $\alpha/2$ as in the above case of a two-tailed test. Any test of a claim will be a one-tailed test. For example, a car manufacturer claims that their Model 17B provides gas mileage of greater than 25 miles per gallon. The null and alternative hypothesis would be:

- $H_0 : \mu \leq 25$

- $H_a : \mu > 25$

The claim would be in the alternative hypothesis. The burden of proof in hypothesis testing is carried in the alternative. This is because failing to reject the null, the status quo, must be accomplished with 90 or 95 percent significance that it cannot be maintained. Said another way, we want to have only a 5 or 10 percent probability of making a Type I error, rejecting a good null; overthrowing the status quo.

This is a one-tailed test and all of the alpha probability is placed in just one tail and not split into $\alpha/2$ as in the above case of a two-tailed test.

Figure 9.6 shows the two possible cases and the form of the null and alternative hypothesis that give rise to them.




Figure 9.6

where μ_0 is the hypothesized value of the population mean.

Table 9.4 Test Statistics for Test of Means, Varying Sample Size, Population Standard Deviation Known or Unknown

Sample size	Test statistic
< 30 (σ unknown)	$t_c = \frac{\bar{X} - \mu_0}{s/\sqrt{n}}$
< 30 (σ known)	$Z_c = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}}$
> 30 (σ unknown)	$Z_c = \frac{\bar{X} - \mu_0}{s/\sqrt{n}}$
> 30 (σ known)	$Z_c = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}}$

Effects of Sample Size on Test Statistic

In developing the confidence intervals for the mean from a sample, we found that most often we would not have the population standard deviation, σ . If the sample size were less than 30, we could simply substitute the point estimate for σ , the sample standard deviation, s , and use the student's t -distribution to correct for this lack of information.

When testing hypotheses we are faced with this same problem and the solution is exactly the same. Namely: If the population standard deviation is unknown, and the sample size is less than 30, substitute s , the point estimate for the population standard deviation, σ , in the formula for the test statistic and use the student's t distribution. All the formulas and figures above are unchanged except for this substitution and changing the Z distribution to the student's t distribution on the graph. Remember that the student's t distribution can only be computed knowing the proper degrees of freedom for the problem. In this case, the degrees of freedom is computed as before with confidence intervals: $df = (n - 1)$. The calculated t -value is compared to the t -value associated with the pre-set level of confidence required in the test, $t_{\alpha, df}$ found in the student's t tables. If we do not know σ , but the sample size is 30 or more, we simply substitute s for σ and use the normal distribution.

Table 9.4 summarizes these rules.

A Systematic Approach for Testing A Hypothesis

A systematic approach to hypothesis testing follows the following steps and in this order. This template will work for all hypotheses that you will ever test.

- Set up the null and alternative hypothesis. This is typically the hardest part of the process. Here the question being asked is reviewed. What parameter is being tested, a mean, a proportion, differences in means, etc. Is this a one-tailed test or two-tailed test? Remember, if someone is making a claim it will always be a one-tailed test.
- Decide the level of significance required for this particular case and determine the critical value. These can be found in the appropriate statistical table. The levels of confidence typical for businesses are 80, 90, 95, 98, and 99. However, the level of significance is a policy decision and should be based upon the risk of making a Type I error, rejecting a good null. Consider the consequences of making a Type I error.

Next, on the basis of the hypotheses and sample size, select the appropriate test statistic and find the relevant critical value: Z_{α} , t_{α} , etc. Drawing the relevant probability distribution and marking the critical value is always big help. Be sure to match the graph with the hypothesis, especially if it is a one-tailed test.

- Take a sample(s) and calculate the relevant parameters: sample mean, standard deviation, or proportion. Using the formula for the test statistic from above in step 2, now calculate the test statistic for this particular case using the parameters you have just calculated.
- Compare the calculated test statistic and the critical value. Marking these on the graph will give a good visual picture of the situation. There are now only two situations:
 1. The test statistic is in the tail: Cannot Accept the null, the probability that this sample mean (proportion) came from the hypothesized distribution is too small to believe that it is the real home of these sample data.
 2. The test statistic is not in the tail: Cannot Reject the null, the sample data are compatible with the hypothesized population parameter.
- Reach a conclusion. It is best to articulate the conclusion two different ways. First a formal statistical conclusion such as “With a 5 % level of significance we cannot accept the null hypotheses that the population mean is equal to XX (units of measurement)”. The second statement of the conclusion is less formal and states the action, or lack of action, required. If the formal conclusion was that above, then the informal one might be, “The machine is broken and we need to shut it down and call for repairs”.

All hypotheses tested will go through this same process. The only changes are the relevant formulas and those are determined by the hypothesis required to answer the original question.

This page titled [9.4: Distribution Needed for Hypothesis Testing](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **9.3: Distribution Needed for Hypothesis Testing** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.5: Full Hypothesis Test Examples

Tests on Means

Example 9.5.8

Jeffrey, as an eight-year old, **established a mean time of 16.43 seconds** for swimming the 25-yard freestyle, with a **standard deviation of 0.8 seconds**. His dad, Frank, thought that Jeffrey could swim the 25-yard freestyle faster using goggles. Frank bought Jeffrey a new pair of expensive goggles and timed Jeffrey for **15 25-yard freestyle swims**. For the 15 swims, **Jeffrey's mean time was 16 seconds**. **Frank thought that the goggles helped Jeffrey to swim faster than the 16.43 seconds**. Conduct a hypothesis test using a preset $\alpha = 0.05$.

Answer

Set up the Hypothesis Test:

Since the problem is about a mean, this is a **test of a single population mean**.

Set the null and alternative hypothesis:

In this case there is an implied challenge or claim. This is that the goggles will reduce the swimming time. The effect of this is to set the hypothesis as a one-tailed test. The claim will always be in the alternative hypothesis because the burden of proof always lies with the alternative. Remember that the status quo must be defeated with a high degree of confidence, in this case 95 % confidence. The null and alternative hypotheses are thus:

$$H_0 : \mu \geq 16.43 \quad H_a : \mu < 16.43$$

For Jeffrey to swim faster, his time will be less than 16.43 seconds. The "<" tells you this is left-tailed.

Determine the distribution needed:

Random variable: \bar{X} = the mean time to swim the 25-yard freestyle.

Distribution for the test statistic:

The sample size is less than 30 and we do not know the population standard deviation so this is a t-test. and the proper formula is: $t_c = \frac{\bar{X} - \mu_0}{s/\sqrt{n}}$

$\mu_0 = 16.43$ comes from H_0 and not the data. $\bar{X} = 16$. $s = 0.8$, and $n = 15$.

Our step 2, setting the level of significance, has already been determined by the problem, .05 for a 95 % significance level. It is worth thinking about the meaning of this choice. The Type I error is to conclude that Jeffrey swims the 25-yard freestyle, on average, in less than 16.43 seconds when, in fact, he actually swims the 25-yard freestyle, on average, in 16.43 seconds. (Reject the null hypothesis when the null hypothesis is true.) For this case the only concern with a Type I error would seem to be that Jeffrey's dad may fail to bet on his son's victory because he does not have appropriate confidence in the effect of the goggles.

To find the critical value we need to select the appropriate test statistic. We have concluded that this is a t-test on the basis of the sample size and that we are interested in a population mean. We can now draw the graph of the t-distribution and mark the critical value. For this problem the degrees of freedom are n-1, or 14. Looking up 14 degrees of freedom at the 0.05 column of the t-table we find 1.761. This is the critical value and we can put this on our graph.

Step 3 is the calculation of the test statistic using the formula we have selected. We find that the calculated test statistic is 2.08, meaning that the sample mean is 2.08 standard deviations away from the hypothesized mean of 16.43.

$$t_c = \frac{\bar{x} - \mu_0}{s/\sqrt{n}} = \frac{16 - 16.43}{.8/\sqrt{15}} = -2.08$$




Figure 9.5.7

Step 4 has us compare the test statistic and the critical value and mark these on the graph. We see that the test statistic is in the tail and thus we move to step 4 and reach a conclusion. The probability that an average time of 16 minutes could come from a distribution with a population mean of 16.43 minutes is too unlikely for us to accept the null hypothesis. We cannot accept the null.

Step 5 has us state our conclusions first formally and then less formally. A formal conclusion would be stated as: “With a 95% level of significance we cannot accept the null hypothesis that the swimming time with goggles comes from a distribution with a population mean time of 16.43 minutes.” Less formally, “With 95% significance we believe that the goggles improves swimming speed”

If we wished to use the p -value system of reaching a conclusion we would calculate the statistic and take the additional step to find the probability of being 2.08 standard deviations from the mean on a t -distribution. This value is .0187. Comparing this to the α -level of .05 we see that we cannot accept the null. The p -value has been put on the graph as the shaded area beyond -2.08 and it shows that it is smaller than the hatched area which is the alpha level of 0.05. Both methods reach the same conclusion that we cannot accept the null hypothesis.

Exercise 9.5.8

The mean throwing distance of a football for Marco, a high school freshman quarterback, is 40 yards, with a standard deviation of two yards. The team coach tells Marco to adjust his grip to get more distance. The coach records the distances for 20 throws. For the 20 throws, Marco’s mean distance was 45 yards. The coach thought the different grip helped Marco throw farther than 40 yards. Conduct a hypothesis test using a preset $\alpha = 0.05$. Assume the throw distances for footballs are normal.

First, determine what type of test this is, set up the hypothesis test, find the p -value, sketch the graph, and state your conclusion.

Example 9.5.9

Jane has just begun her new job as on the sales force of a very competitive company. In a sample of 16 sales calls it was found that she closed the contract for an average value of 108 dollars with a standard deviation of 12 dollars. Test at 5% significance that the population mean is at least 100 dollars against the alternative that it is less than 100 dollars. Company policy requires that new members of the sales force must exceed an average of \$100 per contract during the trial employment period. Can we conclude that Jane has met this requirement at the significance level of 95%?

Answer

1. $H_0 : \mu \leq 100$

$H_a : \mu > 100$

The null and alternative hypothesis are for the parameter μ because the number of dollars of the contracts is a continuous random variable. Also, this is a one-tailed test because the company has only an interested if the number of dollars per contact is below a particular number not "too high" a number. This can be thought of as making a claim that the requirement is being met and thus the claim is in the alternative hypothesis.

2. Test statistic: $t_c = \frac{\bar{x} - \mu_0}{\frac{s}{\sqrt{n}}} = \frac{108 - 100}{\left(\frac{12}{\sqrt{16}}\right)} = 2.67$

3. Critical value: $t_a = 1.753$ with $n - 1$ degrees of freedom = 15

The test statistic is a Student's t because the sample size is below 30; therefore, we cannot use the normal distribution. Comparing the calculated value of the test statistic and the critical value of tt (ta)(ta) at a 5% significance level, we see that the calculated value is in the tail of the distribution. Thus, we conclude that 108 dollars per contract is significantly larger than the hypothesized value of 100 and thus we cannot accept the null hypothesis. There is evidence that supports Jane's performance meets company standards.




Figure 9.5.8

Exercise 9.5.9

It is believed that a stock price for a particular company will grow at a rate of \$5 per week with a standard deviation of \$1. An investor believes the stock won't grow as quickly. The changes in stock price is recorded for ten weeks and are as follows: \$4, \$3, \$2, \$3, \$1, \$7, \$2, \$1, \$1, \$2. Perform a hypothesis test using a 5% level of significance. State the null and alternative hypotheses, state your conclusion, and identify the Type I errors.

Example 9.5.10

A manufacturer of salad dressings uses machines to dispense liquid ingredients into bottles that move along a filling line. The machine that dispenses salad dressings is working properly when 8 ounces are dispensed. Suppose that the average amount dispensed in a particular sample of 35 bottles is 7.91 ounces with a variance of 0.03 ounces squared, s^2 . Is there evidence that the machine should be stopped and production wait for repairs? The lost production from a shutdown is potentially so great that management feels that the level of significance in the analysis should be 99%.

Again we will follow the steps in our analysis of this problem.

Answer

STEP 1: Set the Null and Alternative Hypothesis. The random variable is the quantity of fluid placed in the bottles. This is a continuous random variable and the parameter we are interested in is the mean. Our hypothesis therefore is about the mean. In this case we are concerned that the machine is not filling properly. From what we are told it does not matter if the machine is over-filling or under-filling, both seem to be an equally bad error. This tells us that this is a two-tailed test: if the machine is malfunctioning it will be shutdown regardless if it is from over-filling or under-filling. The null and alternative hypotheses are thus:

$$H_0 : \mu = 8$$

$$H_a : \mu \neq 8$$

STEP 2: Decide the level of significance and draw the graph showing the critical value.

This problem has already set the level of significance at 99%. The decision seems an appropriate one and shows the thought process when setting the significance level. Management wants to be very certain, as certain as probability will allow, that they are not shutting down a machine that is not in need of repair. To draw the distribution and the critical value, we need to know which distribution to use. Because this is a continuous random variable and we are interested in the mean, and the sample size is greater than 30, the appropriate distribution is the normal distribution and the relevant critical value is 2.575 from the normal table or the t-table at 0.005 column and infinite degrees of freedom. We draw the graph and mark these points.




Figure 9.5.9

STEP 3: Calculate sample parameters and the test statistic. The sample parameters are provided, the sample mean is 7.91 and the sample variance is .03 and the sample size is 35. We need to note that the sample variance was provided not the sample standard deviation, which is what we need for the formula. Remembering that the standard deviation is simply the square root of the variance, we therefore know the sample standard deviation, s , is 0.173. With this information we calculate the test statistic as -3.07, and mark it on the graph.

$$Z_c = \frac{\bar{x} - \mu_0}{s/\sqrt{n}} = \frac{7.91 - 8}{.173/\sqrt{35}} = -3.07$$

STEP 4: Compare test statistic and the critical values Now we compare the test statistic and the critical value by placing the test statistic on the graph. We see that the test statistic is in the tail, decidedly greater than the critical value of 2.575. We note that even the very small difference between the hypothesized value and the sample value is still a large number of standard deviations. The sample mean is only 0.08 ounces different from the required level of 8 ounces, but it is 3 plus standard deviations away and thus we cannot accept the null hypothesis.

STEP 5: Reach a Conclusion

Three standard deviations of a test statistic will guarantee that the test will fail. The probability that anything is within three standard deviations is almost zero. Actually it is 0.0026 on the normal distribution, which is certainly almost zero in a practical sense. Our formal conclusion would be “ At a 99% level of significance we cannot accept the hypothesis that the sample mean came from a distribution with a mean of 8 ounces” Or less formally, and getting to the point, “At a 99% level of significance we conclude that the machine is under filling the bottles and is in need of repair”.

Hypothesis Test for Proportions

Just as there were confidence intervals for proportions, or more formally, the population parameter p of the binomial distribution, there is the ability to test hypotheses concerning p .

The population parameter for the binomial is p . The estimated value (point estimate) for p is p' where $p' = x/n$, x is the number of successes in the sample and n is the sample size.

When you perform a hypothesis test of a population proportion p , you take a simple random sample from the population. The conditions for a **binomial distribution** must be met, which are: there are a certain number n of independent trials meaning random sampling, the outcomes of any trial are binary, success or failure, and each trial has the same probability of a success p . The shape of the binomial distribution needs to be similar to the shape of the normal distribution. To ensure this, the quantities np' and nq' must both be greater than five ($np' > 5$ and $nq' > 5$). In this case the binomial distribution of a sample (estimated) proportion can be approximated by the normal distribution with $\mu = np$ and $\sigma = \sqrt{npq}$. Remember that $q = 1 - p$. There is no distribution that can correct for this small sample bias and thus if these conditions are not met we simply cannot test the hypothesis with the data available at that time. We met this condition when we first were estimating confidence intervals for p .

Again, we begin with the standardizing formula modified because this is the distribution of a binomial.

$$Z = \frac{p' - p}{\sqrt{\frac{pq}{n}}}$$

Substituting p_0 , the hypothesized value of p , we have:

$$Z_c = \frac{p' - p_0}{\sqrt{\frac{p_0q_0}{n}}}$$

This is the test statistic for testing hypothesized values of p , where the null and alternative hypotheses take one of the following forms:

Table 9.5.5

Two-tailed test	One-tailed test	One-tailed test
$H_0 : p = p_0$	$H_0 : p \leq p_0$	$H_0 : p \geq p_0$
$H_a : p \neq p_0$	$H_a : p > p_0$	$H_a : p < p_0$

The decision rule stated above applies here also: if the calculated value of Z_c shows that the sample proportion is "too many" standard deviations from the hypothesized proportion, the null hypothesis cannot be accepted. The decision as to what is "too many" is pre-determined by the analyst depending on the level of significance required in the test.

Example 9.5.11

The mortgage department of a large bank is interested in the nature of loans of first-time borrowers. This information will be used to tailor their marketing strategy. They believe that 50% of first-time borrowers take out smaller loans than other borrowers. They perform a hypothesis test to determine if the percentage is **the same or different from 50%**. They sample **100 first-time borrowers** and find **53** of these loans are smaller than the other borrowers. For the hypothesis test, they choose a 5% level of significance.

Answer

STEP 1: Set the null and alternative hypothesis.

$$H_0 : p = 0.50 \quad H_a : p \neq 0.50$$

The words "**is the same or different from**" tell you this is a two-tailed test. The Type I and Type II errors are as follows: The Type I error is to conclude that the proportion of borrowers is different from 50% when, in fact, the proportion is actually 50%. (Reject the null hypothesis when the null hypothesis is true). The Type II error is there is not enough evidence to conclude that the proportion of first time borrowers differs from 50% when, in fact, the proportion does differ from 50%. (You fail to reject the null hypothesis when the null hypothesis is false.)

STEP 2: Decide the level of significance and draw the graph showing the critical value

The level of significance has been set by the problem at the 95% level. Because this is two-tailed test one-half of the alpha value will be in the upper tail and one-half in the lower tail as shown on the graph. The critical value for the normal distribution at the 95% level of confidence is 1.96. This can easily be found on the student's t-table at the very bottom at infinite degrees of freedom remembering that at infinity the t-distribution is the normal distribution. Of course the value can also be found on the normal table but you have go looking for one-half of 95 (0.475) inside the body of the table and then read out to the sides and top for the number of standard deviations.




Figure 9.5.10

STEP 3: Calculate the sample parameters and critical value of the test statistic.

The test statistic is a normal distribution, Z , for testing proportions and is:

$$Z = \frac{p' - p_0}{\sqrt{\frac{p_0 q_0}{n}}} = \frac{.53 - .50}{\sqrt{\frac{.5(.5)}{100}}} = 0.60$$

For this case, the sample of 100 found 53 first-time borrowers were different from other borrowers. The sample proportion, $p' = 53/100 = 0.53$. The test question, therefore, is: "Is 0.53 significantly different from .50?" Putting these values into the formula for the test statistic we find that 0.53 is only 0.60 standard deviations away from .50. This is barely off of the mean of the standard normal distribution of zero. There is virtually no difference from the sample proportion and the hypothesized proportion in terms of standard deviations.

STEP 4: Compare the test statistic and the critical value.

The calculated value is well within the critical values of ± 1.96 standard deviations and thus we cannot reject the null hypothesis. To reject the null hypothesis we need significant evidence of difference between the hypothesized value and the sample value. In this case the sample value is very nearly the same as the hypothesized value measured in terms of standard deviations.

STEP 5: Reach a conclusion

The formal conclusion would be "At a 95% level of significance we cannot reject the null hypothesis that 50% of first-time borrowers have the same size loans as other borrowers". Less formally we would say that "There is no evidence that one-half of first-time borrowers are significantly different in loan size from other borrowers". Notice the length to which the conclusion goes to include all of the conditions that are attached to the conclusion. Statisticians for all the criticism they receive, are careful to be very specific even when this seems trivial. Statisticians cannot say more than they know and the data constrain the conclusion to be within the metes and bounds of the data.

Exercise 9.5.11

A teacher believes that 85% of students in the class will want to go on a field trip to the local zoo. She performs a hypothesis test to determine if the percentage is the same or different from 85%. The teacher samples 50 students and 39 reply that they would want to go to the zoo. For the hypothesis test, use a 1% level of significance.

Example 9.5.12

Suppose a consumer group suspects that the proportion of households that have three or more cell phones is 30%. A cell phone company has reason to believe that the proportion is not 30%. Before they start a big advertising campaign, they conduct a hypothesis test. Their marketing people survey 150 households with the result that 43 of the households have three or more cell phones.

Answer

Here is an abbreviated version of the system to solve hypothesis tests applied to a test on a proportions.


$$H_0 : p = 0.3$$

$$H_a : p \neq 0.3$$

$$n = 150$$

$$p' = \frac{x}{n} = \frac{43}{150} = 0.287$$

$$Z_c = \frac{p' - p_0}{\sqrt{\frac{p_0 q_0}{n}}} = \frac{0.287 - 0.3}{\sqrt{\frac{3(7)}{150}}} = 0.347$$



At a significance level of 90%
we cannot reject H_0 ;
the consumer group is correct.

Figure 9.5.11

Example 9.5.13

The National Institute of Standards and Technology provides exact data on conductivity properties of materials. Following are conductivity measurements for 11 randomly selected pieces of a particular type of glass.

1.11; 1.07; 1.11; 1.07; 1.12; 1.08; .98; .98 1.02; .95; .95

Is there convincing evidence that the average conductivity of this type of glass is greater than one? Use a significance level of 0.05.

Answer

Let's follow a four-step process to answer this statistical question.

State the Question: We need to determine if, at a 0.05 significance level, the average conductivity of the selected glass is greater than one. Our hypotheses will be

- $H_0 : \mu \leq 1$
- $H_a : \mu > 1$

Plan: We are testing a sample mean without a known population standard deviation with less than 30 observations. Therefore, we need to use a Student's-t distribution. Assume the underlying population is normal. **Do the calculations and draw the graph. State the Conclusions:** We cannot accept the null hypothesis. It is reasonable to state that the data supports the claim that the average conductivity level is greater than one.

Example 9.5.14

In a study of 420,019 cell phone users, 172 of the subjects developed brain cancer. Test the claim that cell phone users developed brain cancer at a greater rate than that for non-cell phone users (the rate of brain cancer for non-cell phone users is 0.0340%). Since this is a critical issue, use a 0.005 significance level. Explain why the significance level should be so low in terms of a Type I error.

Answer

- We need to conduct a hypothesis test on the claimed cancer rate. Our hypotheses will be

a. $H_0 : p \leq 0.00034$

b. $H_a : p > 0.00034$

If we commit a Type I error, we are essentially accepting a false claim. Since the claim describes cancer-causing environments, we want to minimize the chances of incorrectly identifying causes of cancer.

2. We will be testing a sample proportion with $x = 172$ and $n = 420,019$. The sample is sufficiently large because we have $np' = 420,019(0.00034) = 142.8$ and $nq' = 420,019(0.99966) = 419,876.2$, two independent outcomes, and a fixed probability of success $p' = 0.00034$. Thus we will be able to generalize our results to the population.

This page titled [9.5: Full Hypothesis Test Examples](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [9.4: Full Hypothesis Test Examples](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.6: Chapter Key Terms

Binomial Distribution

a discrete random variable (RV) that arises from Bernoulli trials. There are a fixed number, n , of independent trials.

“Independent” means that the result of any trial (for example, trial 1) does not affect the results of the following trials, and all trials are conducted under the same conditions. Under these circumstances the binomial RV X is defined as the number of successes in n trials. The notation is: $X \sim B(n, p)$ $\mu = np$ and the standard deviation is $\sigma = \sqrt{npq}$. The probability of exactly

x successes in n trials is $P(X = x) = \binom{n}{x} p^x q^{n-x}$.

Central Limit Theorem

Given a random variable (RV) with known mean μ and known standard deviation σ . We are sampling with size n and we are interested in two new RVs - the sample mean, \bar{X} . If the size n of the sample is sufficiently large, then $\bar{X} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$. If the size n of the sample is sufficiently large, then the distribution of the sample means will approximate a normal distribution regardless of the shape of the population. The expected value of the mean of the sample means will equal the population mean. The standard deviation of the distribution of the sample means, $\frac{\sigma}{\sqrt{n}}$, is called the standard error of the mean.

Confidence Interval (CI)

an interval estimate for an unknown population parameter. This depends on:

- The desired confidence level.
- Information that is known about the distribution (for example, known standard deviation).
- The sample and its size.

Critical Value

The t or Z value set by the researcher that measures the probability of a Type I error, σ .

Hypothesis

a statement about the value of a population parameter, in case of two hypotheses, the statement assumed to be true is called the null hypothesis (notation H_0) and the contradictory statement is called the alternative hypothesis (notation H_a).

Hypothesis Testing

Based on sample evidence, a procedure for determining whether the hypothesis stated is a reasonable statement and should not be rejected, or is unreasonable and should be rejected.

Normal Distribution

a continuous random variable (RV) with pdf $f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$, where μ is the mean of the distribution, and σ is the standard deviation, notation: $X \sim N(\mu, \sigma)$. If $\mu = 0$ and $\sigma = 1$, the RV is called **the standard normal distribution**.

Standard Deviation

a number that is equal to the square root of the variance and measures how far data values are from their mean; notation: s for sample standard deviation and σ for population standard deviation.

Student's t-Distribution

investigated and reported by William S. Gossett in 1908 and published under the pseudonym Student. The major characteristics of the random variable (RV) are:

- It is continuous and assumes any real values.
- The pdf is symmetrical about its mean of zero. However, it is more spread out and flatter at the apex than the normal distribution.
- It approaches the standard normal distribution as n gets larger.
- There is a "family" of t distributions: every representative of the family is completely defined by the number of degrees of freedom which is one less than the number of data items.

Test Statistic

The formula that counts the number of standard deviations on the relevant distribution that estimated parameter is away from the hypothesized value.

Type I Error

The decision is to reject the null hypothesis when, in fact, the null hypothesis is true.

Type II Error

The decision is not to reject the null hypothesis when, in fact, the null hypothesis is false.

This page titled [9.6: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [9.7: Chapter Key Terms](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.7: Chapter Review

9.1 Null and Alternative Hypotheses

In a **hypothesis test**, sample data is evaluated in order to arrive at a decision about some type of claim. If certain conditions about the sample are satisfied, then the claim can be evaluated for a population. In a hypothesis test, we:

1. Evaluate the **null hypothesis**, typically denoted with H_0 . The null is not rejected unless the hypothesis test shows otherwise. The null statement must always contain some form of equality ($=$, \leq or \geq)
2. Always write the **alternative hypothesis**, typically denoted with H_a or H_1 , using not equal, less than or greater than symbols, i.e., (*neq*, $<$, or $>$).
3. If we reject the null hypothesis, then we can assume there is enough evidence to support the alternative hypothesis.
4. Never state that a claim is proven true or false. Keep in mind the underlying fact that hypothesis testing is based on probability laws; therefore, we can talk only in terms of non-absolute certainties.

9.2 Outcomes and the Type I and Type II Errors

In every hypothesis test, the outcomes are dependent on a correct interpretation of the data. Incorrect calculations or misunderstood summary statistics can yield errors that affect the results. A **Type I** error occurs when a true null hypothesis is rejected. A **Type II error** occurs when a false null hypothesis is not rejected.

The probabilities of these errors are denoted by the Greek letters α and β , for a Type I and a Type II error respectively. The power of the test, $1 - \beta$, quantifies the likelihood that a test will yield the correct result of a true alternative hypothesis being accepted. A high power is desirable.

9.3 Distribution Needed for Hypothesis Testing

In order for a hypothesis test's results to be generalized to a population, certain requirements must be satisfied.

When testing for a single population mean:

1. A Student's t -test should be used if the data come from a simple, random sample and the population is approximately normally distributed, or the sample size is large, with an unknown standard deviation.
2. The normal test will work if the data come from a simple, random sample and the population is approximately normally distributed, or the sample size is large.

When testing a single population proportion use a normal test for a single population proportion if the data comes from a simple, random sample, fill the requirements for a binomial distribution, and the mean number of successes and the mean number of failures satisfy the conditions: $np > 5$ and $nq > 5$ where n is the sample size, p is the probability of a success, and q is the probability of a failure.

9.4 Full Hypothesis Test Examples

The **hypothesis test** itself has an established process. This can be summarized as follows:

1. Determine H_0 and H_a . Remember, they are contradictory.
2. Determine the random variable.
3. Determine the distribution for the test.
4. Draw a graph and calculate the test statistic.
5. Compare the calculated test statistic with the Z critical value determined by the level of significance required by the test and make a decision (cannot reject H_0 or cannot accept H_0), and write a clear conclusion using English sentences.

This page titled [9.7: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [9.10: Chapter Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.8: Chapter Formula Review

9.3 Distribution Needed for Hypothesis Testing

Table 9.8.6 Test Statistics for Test of Means, Varying Sample Size, Population Known or Unknown

Sample size	Test statistic
< 30 (σ unknown)	$t_c = \frac{\bar{X} - \mu_0}{s/\sqrt{n}}$
< 30 (σ known)	$Z_c = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}}$
> 30 (σ unknown)	$Z_c = \frac{\bar{X} - \mu_0}{s/\sqrt{n}}$
> 30 (σ known)	$Z_c = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}}$

This page titled [9.8: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [9.5: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.9: Chapter Practice

Exercise 9.9.12.15.

A group of doctors is deciding whether or not to perform an operation. Suppose the null hypothesis, H_0 , is: the surgical procedure will go well. State the Type I and Type II errors in complete sentences.

16.

A group of doctors is deciding whether or not to perform an operation. Suppose the null hypothesis, H_0 , is: the surgical procedure will go well. Which is the error with the greater consequence?

17.

The power of a test is 0.981. What is the probability of a Type II error?

18.

A group of divers is exploring an old sunken ship. Suppose the null hypothesis, H_0 , is: the sunken ship does not contain buried treasure. State the Type I and Type II errors in complete sentences.

19.

A microbiologist is testing a water sample for E-coli. Suppose the null hypothesis, H_0 , is: the sample does not contain E-coli. The probability that the sample does not contain E-coli, but the microbiologist thinks it does is 0.012. The probability that the sample does contain E-coli, but the microbiologist thinks it does not is 0.002. What is the power of this test?

20.

A microbiologist is testing a water sample for E-coli. Suppose the null hypothesis, H_0 , is: the sample contains E-coli. Which is the error with the greater consequence?

9.3 Distribution Needed for Hypothesis Testing

21.

Which two distributions can you use for hypothesis testing for this chapter?

22.

Which distribution do you use when you are testing a population mean and the population standard deviation is known? Assume sample size is large. Assume a normal distribution with $n \geq 30$.

23.

Which distribution do you use when the standard deviation is not known and you are testing one population mean? Assume a normal distribution, with $n \geq 30$.

24.

A population mean is 13. The sample mean is 12.8, and the sample standard deviation is two. The sample size is 20. What distribution should you use to perform a hypothesis test? Assume the underlying population is normal.

25.

A population has a mean is 25 and a standard deviation of five. The sample mean is 24, and the sample size is 108. What distribution should you use to perform a hypothesis test?

26.

It is thought that 42% of respondents in a taste test would prefer Brand A. In a particular test of 100 people, 39% preferred Brand A. What distribution should you use to perform a hypothesis test?

27.

You are performing a hypothesis test of a single population mean using a Student's t -distribution. What must you assume about the distribution of the data?

28.

You are performing a hypothesis test of a single population mean using a Student's t -distribution. The data are not from a simple random sample. Can you accurately perform the hypothesis test?

29.

You are performing a hypothesis test of a single population proportion. What must be true about the quantities of np and nq ?

30.

You are performing a hypothesis test of a single population proportion. You find out that np is less than five. What must you do to be able to perform a valid hypothesis test?

31.

You are performing a hypothesis test of a single population proportion. The data come from which distribution?

9.4 Full Hypothesis Test Examples

32.

Assume $H_0 : \mu = 9$ and $H_a : \mu < 9$. Is this a left-tailed, right-tailed, or two-tailed test?

33.

Assume $H_0 : \mu \leq 6$ and $H_a : \mu > 6$. Is this a left-tailed, right-tailed, or two-tailed test?

34.

Assume $H_0 : p = 0.25$ and $H_a : p \neq 0.25$. Is this a left-tailed, right-tailed, or two-tailed test?

35.

Draw the general graph of a left-tailed test.

36.

Draw the graph of a two-tailed test.

37.

A bottle of water is labeled as containing 16 fluid ounces of water. You believe it is less than that. What type of test would you use?

38.

Your friend claims that his mean golf score is 63. You want to show that it is higher than that. What type of test would you use?

39.

A bathroom scale claims to be able to identify correctly any weight within a pound. You think that it cannot be that accurate. What type of test would you use?

40.

You flip a coin and record whether it shows heads or tails. You know the probability of getting heads is 50%, but you think it is less for this particular coin. What type of test would you use?

41.

If the alternative hypothesis has a not equals (\neq) symbol, you know to use which type of test?

42.

Assume the null hypothesis states that the mean is at least 18. Is this a left-tailed, right-tailed, or two-tailed test?

43.

Assume the null hypothesis states that the mean is at most 12. Is this a left-tailed, right-tailed, or two-tailed test?

44.

Assume the null hypothesis states that the mean is equal to 88. The alternative hypothesis states that the mean is not equal to 88. Is this a left-tailed, right-tailed, or two-tailed test?

This page titled [9.9: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **9.8: Chapter Practice** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.10: Chapter Homework

Exercise **9.10.63**. Conduct a hypothesis test to see if your decision and conclusion would change if your belief were that the brown trout's mean I.Q. is **not** four.**65**.

According to an article in *Newsweek*, the natural ratio of girls to boys is 100:105. In China, the birth ratio is 100: 114 (46.7% girls). Suppose you don't believe the reported figures of the percent of girls born in China. You conduct a study. In this study, you count the number of girls and boys born in 150 randomly chosen recent births. There are 60 girls and 90 boys born of the 150. Based on your study, do you believe that the percent of girls born in China is 46.7?

66.

A poll done for *Newsweek* found that 13% of Americans have seen or sensed the presence of an angel. A contingent doubts that the percent is really that high. It conducts its own survey. Out of 76 Americans surveyed, only two had seen or sensed the presence of an angel. As a result of the contingent's survey, would you agree with the *Newsweek* poll? In complete sentences, also give three reasons why the two polls might give different results.

67.

The mean work week for engineers in a start-up company is believed to be about 60 hours. A newly hired engineer hopes that it's shorter. She asks ten engineering friends in start-ups for the lengths of their mean work weeks. Based on the results that follow, should she count on the mean work week to be shorter than 60 hours?

Data (length of mean work week): 70; 45; 55; 60; 65; 55; 55; 60; 50; 55.

68.

Sixty-eight percent of online courses taught at community colleges nationwide were taught by full-time faculty. To test if 68% also represents California's percent for full-time faculty teaching the online classes, Long Beach City College (LBCC) in California, was randomly selected for comparison. In the same year, 34 of the 44 online courses LBCC offered were taught by full-time faculty. Conduct a hypothesis test to determine if 68% represents California. NOTE: For more accurate results, use more California community colleges and this past year's data.

69.

According to an article in *Bloomberg Businessweek*, New York City's most recent adult smoking rate is 14%. Suppose that a survey is conducted to determine this year's rate. Nine out of 70 randomly chosen N.Y. City residents reply that they smoke. Conduct a hypothesis test to determine if the rate is still 14% or if it has decreased.

70.

The mean age of De Anza College students in a previous term was 26.6 years old. An instructor thinks the mean age for online students is older than 26.6. She randomly surveys 56 online students and finds that the sample mean is 29.4 with a standard deviation of 2.1. Conduct a hypothesis test.

71.

Registered nurses earned an average annual salary of \$69,110. For that same year, a survey was conducted of 41 California registered nurses to determine if the annual salary is higher than \$69,110 for California nurses. The sample average was \$71,121 with a sample standard deviation of \$7,489. Conduct a hypothesis test.

72.

La Leche League International reports that the mean age of weaning a child from breastfeeding is age four to five worldwide. In America, most nursing mothers wean their children much earlier. Suppose a random survey is conducted of 21 U.S. mothers who recently weaned their children. The mean weaning age was nine months ($3/4$ year) with a standard deviation of 4 months. Conduct a hypothesis test to determine if the mean weaning age in the U.S. is less than four years old.

73.

Over the past few decades, public health officials have examined the link between weight concerns and teen girls' smoking. Researchers surveyed a group of 273 randomly selected teen girls living in Massachusetts (between 12 and 15 years old). After four years the girls were surveyed again. Sixty-three said they smoked to stay thin. Is there good evidence that more than thirty percent of the teen girls smoke to stay thin?

After conducting the test, your decision and conclusion are

a. Instructions: For the following ten exercises,

Hypothesis testing: For the following ten exercises, answer each question.

1. 3; 2; 1; 3; 7; 2; 9; 4; 6; 6; 8; 0; 5; 6; 4; 2; 1; 3; 4; 1

At the $\alpha = 0.05$ level can it be concluded that the sample mean is higher than 5.8 visits per year?

84.

According to the *N.Y. Times Almanac* the mean family size in the U.S. is 3.18. A sample of a college math class resulted in the following family sizes:

5; 4; 5; 4; 4; 3; 6; 4; 3; 3; 5; 5; 6; 3; 3; 2; 7; 4; 5; 2; 2; 3; 2

At $\alpha = 0.05$ level, is the class' mean family size greater than the national average? Does the Almanac result remain valid? Why?

85.

The student academic group on a college campus claims that freshman students study at least 2.5 hours per day, on average. One Introduction to Statistics class was skeptical. The class took a random sample of 30 freshman students and found a mean study time of 137 minutes with a standard deviation of 45 minutes. At $\alpha = 0.01$ level, is the student academic group's claim correct?

This page titled 9.10: Chapter Homework is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- 9.6: Chapter Homework by OpenStax is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.11: Chapter References

9.1 Null and Alternative Hypotheses

Data from the National Institute of Mental Health. Available online at <http://www.nimh.nih.gov/publicat/depression.cfm>.

9.4 Full Hypothesis Test Examples

Data from Amit Schitai. Director of Instructional Technology and Distance Learning. LBCC.

Data from *Bloomberg Businessweek*. Available online at <http://www.businessweek.com/news/2011-09-15/nyc-smoking-rate-falls-to-record-low-of-14-bloomberg-says.html>.

Data from energy.gov. Available online at <http://energy.gov> (accessed June 27, 2013).

Data from Gallup®. Available online at www.gallup.com (accessed June 27, 2013).

Data from *Growing by Degrees* by Allen and Seaman.

Data from La Leche League International. Available online at www.lalecheleague.org/Law/BAFeb01.html.

Data from the American Automobile Association. Available online at www.aaa.com (accessed June 27, 2013).

Data from the American Library Association. Available online at www.ala.org (accessed June 27, 2013).

Data from the Bureau of Labor Statistics. Available online at <http://www.bls.gov/oes/current/oes291111.htm>.

Data from the Centers for Disease Control and Prevention. Available online at www.cdc.gov (accessed June 27, 2013).

Data from the U.S. Census Bureau, available online at quickfacts.census.gov/qfd/states/00000.html (accessed June 27, 2013).

Data from the United States Census Bureau. Available online at www.census.gov/hhes/socdemo/language/.

Data from Toastmasters International. Available online at toastmasters.org/artisan/details?eID=429&Page=1.

Data from Weather Underground. Available online at www.wunderground.com (accessed June 27, 2013).

Federal Bureau of Investigations. “Uniform Crime Reports and Index of Crime in Daviess in the State of Kentucky enforced by Daviess County from 1985 to 2005.” Available online at <http://www.disastercenter.com/kentucky/crime/3868.htm> (accessed June 27, 2013).

“Foothill-De Anza Community College District.” De Anza College, Winter 2006. Available online at research.fhda.edu/factbook/DA...t_da_2006w.pdf.

Johansen, C., J. Boice, Jr., J. McLaughlin, J. Olsen. “Cellular Telephones and Cancer—a Nationwide Cohort Study in Denmark.” *Institute of Cancer Epidemiology and the Danish Cancer Society*, 93(3):203-7. Available online at <http://www.ncbi.nlm.nih.gov/pubmed/11158188> (accessed June 27, 2013).

Rape, Abuse & Incest National Network. “How often does sexual assault occur?” RAINN, 2009. Available online at <http://www.rainn.org/get-information...sexual-assault> (accessed June 27, 2013).

This page titled **9.11: Chapter References** is shared under a **CC BY 4.0** license and was authored, remixed, and/or curated by **OpenStax** via **source content** that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **9.9: Chapter References** by **OpenStax** is licensed **CC BY 4.0**. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

9.12: Chapter Solution (Practice + Homework)

Figure 9.12.12

38.

a right-tailed test

40.

a left-tailed test

42.

This is a left-tailed test.

44.

This is a two-tailed test.

45.

1. 47.

c

1. 51.

b

d

55.

d

56.

1. 58.

1. 60.

1. 62.

1. 64.

1. 66.

1. 69.

1. 71.

1. 73.

c

c

77.

1. 79.

1. 81.

1. 83.

1. 85.

1. $H_0 : \mu \geq 150 H_a : \mu < 150$

2. p -value = 0.0622

3. alpha = 0.01

4. Do not reject the null hypothesis.

5. At the 1% significance level, there is not enough evidence to conclude that freshmen students study less than 2.5 hours per day, on average.

6. The student academic group's claim appears to be correct.

This page titled [9.12: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **9.11: Chapter Solution (Practice + Homework)** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

10: Hypothesis Testing with Two Samples

- 10.1: Introduction
- 10.2: Comparing Two Independent Population Means
- 10.3: Cohen's Standards for Small, Medium, and Large Effect Sizes
- 10.4: Test for Differences in Means- Assuming Equal Population Variances
- 10.5: Comparing Two Independent Population Proportions
- 10.6: Two Population Means with Known Standard Deviations
- 10.7: Matched or Paired Samples
- 10.8: Chapter Key Terms
- 10.9: Chapter Review
- 10.10: Chapter Formula Review
- 10.11: Chapter Practice
- 10.12: Homework
- 10.13: Chapter Homework
- 10.14: Chapter References
- 10.15: Chapter Solution (Practice + Homework)

This page titled [10: Hypothesis Testing with Two Samples](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

10.1: Introduction



Figure 10.1.1 If you want to test a claim that involves two groups (the types of breakfasts eaten east and west of the Mississippi River) you can use a slightly different technique when conducting a hypothesis test. (credit: Chloe Lim)

Studies often compare two groups. For example, researchers are interested in the effect aspirin has in preventing heart attacks. Over the last few years, newspapers and magazines have reported various aspirin studies involving two groups. Typically, one group is given aspirin and the other group is given a placebo. Then, the heart attack rate is studied over several years.

There are other situations that deal with the comparison of two groups. For example, studies compare various diet and exercise programs. Politicians compare the proportion of individuals from different income brackets who might vote for them. Students are interested in whether SAT or GRE preparatory courses really help raise their scores. Many business applications require comparing two groups. It may be the investment returns of two different investment strategies, or the differences in production efficiency of different management styles.

To compare two means or two proportions, you work with two groups. The groups are classified either as **independent** or **matched pairs**. **Independent groups** consist of two samples that are independent, that is, sample values selected from one population are not related in any way to sample values selected from the other population. **Matched pairs** consist of two samples that are dependent. The parameter tested using matched pairs is the population mean. The parameters tested using independent groups are either population means or population proportions of each group.

This page titled [10.1: Introduction](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.0: Introduction** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.2: Comparing Two Independent Population Means

The comparison of two independent population means is very common and provides a way to test the hypothesis that the two groups differ from each other. Is the night shift less productive than the day shift, are the rates of return from fixed asset investments different from those from common stock investments, and so on? An observed difference between two sample means depends on both the means and the sample standard deviations. Very different means can occur by chance if there is great variation among the individual samples. The test statistic will have to account for this fact. The test comparing two independent population means with unknown and possibly unequal population standard deviations is called the Aspin-Welch t -test. The degrees of freedom formula we will see later was developed by Aspin-Welch.

When we developed the hypothesis test for the mean and proportions we began with the Central Limit Theorem. We recognized that a sample mean came from a distribution of sample means, and sample proportions came from the sampling distribution of sample proportions. This made our sample parameters, the sample means and sample proportions, into random variables. It was important for us to know the distribution that these random variables came from. The Central Limit Theorem gave us the answer: the normal distribution. Our Z and t statistics came from this theorem. This provided us with the solution to our question of how to measure the probability that a sample mean came from a distribution with a particular hypothesized value of the mean or proportion. In both cases that was the question: what is the probability that the mean (or proportion) from our sample data came from a population distribution with the hypothesized value we are interested in?

Now we are interested in whether or not two samples have the same mean. Our question has not changed: Do these two samples come from the same population distribution? To approach this problem we create a new random variable. We recognize that we have two sample means, one from each set of data, and thus we have two random variables coming from two unknown distributions. To solve the problem we create a new random variable, the difference between the sample means. This new random variable also has a distribution and, again, the Central Limit Theorem tells us that this new distribution is normally distributed, regardless of the underlying distributions of the original data. A graph may help to understand this concept.




Figure 10.2.2

Pictured are two distributions of data, X_1 and X_2 , with unknown means and standard deviations. The second panel shows the sampling distribution of the newly created random variable $(\bar{X}_1 - \bar{X}_2)$. This distribution is the theoretical distribution of many many sample means from population 1 minus sample means from population 2. The Central Limit Theorem tells us that this theoretical sampling distribution of differences in sample means is normally distributed, regardless of the distribution of the actual population data shown in the top panel. Because the sampling distribution is normally distributed, we can develop a standardizing formula and calculate probabilities from the standard normal distribution in the bottom panel, the Z distribution. We have seen this same analysis before in Chapter 7 [Figure 10.2.2](#).

The Central Limit Theorem, as before, provides us with the standard deviation of the sampling distribution, and further, that the expected value of the mean of the distribution of differences in sample means is equal to the differences in the population means.

Mathematically this can be stated:

$$E(\mu_{\bar{x}_1} - \mu_{\bar{x}_2}) = \mu_1 - \mu_2$$

Because we do not know the population standard deviations, we estimate them using the two sample standard deviations from our independent samples. For the hypothesis test, we calculate the estimated standard deviation, or **standard error**, of **the difference in sample means**, $\bar{X}_1 - \bar{X}_2$.

The standard error is:

$$\sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}$$

We remember that substituting the sample variance for the population variance when we did not have the population variance was the technique we used when building the confidence interval and the test statistic for the test of hypothesis for a single mean back in [Confidence Intervals](#) and [Hypothesis Testing with One Sample](#). **The test statistic (t-score) is calculated as follows:**

$$t_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}}$$

where:

- s_1 and s_2 , the sample standard deviations, are estimates of σ_1 and σ_2 , respectively and
- σ_1 and σ_2 are the unknown population standard deviations.
- \bar{x}_1 and \bar{x}_2 are the sample means. μ_1 and μ_2 are the unknown population means.

The number of **degrees of freedom (df)** requires a somewhat complicated calculation. The df are not always a whole number. The test statistic above is approximated by the Student's t -distribution with df as follows:

Degrees of freedom

$$df = \frac{\left(\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}\right)^2}{\left(\frac{1}{n_1-1}\right)\left(\frac{(s_1)^2}{n_1}\right)^2 + \left(\frac{1}{n_2-1}\right)\left(\frac{(s_2)^2}{n_2}\right)^2}$$

When both sample sizes n_1 and n_2 are 30 or larger, the Student's t approximation is very good. If each sample has more than 30 observations then the degrees of freedom can be calculated as $n_1 + n_2 - 2$.

The format of the sampling distribution, differences in sample means, specifies that the format of the null and alternative hypothesis is:

$$H_0 : \mu_1 - \mu_2 = \delta_0$$

$$H_a : \mu_1 - \mu_2 \neq \delta_0$$

where δ_0 is the hypothesized difference between the two means. If the question is simply “is there any difference between the means?” then $\delta_0 = 0$ and the null and alternative hypotheses becomes:

$$H_0 : \mu_1 = \mu_2$$

$$H_a : \mu_1 \neq \mu_2$$

An example of when δ_0 might not be zero is when the comparison of the two groups requires a specific difference for the decision to be meaningful. Imagine that you are making a capital investment. You are considering changing from your current model machine to another. You measure the productivity of your machines by the speed they produce the product. It may be that a contender to replace the old model is faster in terms of product throughput, but is also more expensive. The second machine may also have more maintenance costs, setup costs, etc. The null hypothesis would be set up so that the new machine would have to be better than the old one by enough to cover these extra costs in terms of speed and cost of production. This form of the null and

alternative hypothesis shows how valuable this particular hypothesis test can be. For most of our work we will be testing simple hypotheses asking if there is any difference between the two distribution means.

Example 10.2.1 INDEPENDENT GROUPS

The Kona Iki Corporation produces coconut milk. They take coconuts and extract the milk inside by drilling a hole and pouring the milk into a vat for processing. They have both a day shift (called the B shift) and a night shift (called the G shift) to do this part of the process. They would like to know if the day shift and the night shift are equally efficient in processing the coconuts. A study is done sampling 9 shifts of the G shift and 16 shifts of the B shift. The results of the number of hours required to process 100 pounds of coconuts is presented in Table 10.2.1. A study is done and data are collected, resulting in the data in Table 10.2.1.

Table 10.2.1

	Sample Size	Average Number of Hours to Process 100 Pounds of Coconuts	Sample Standard Deviation
G Shift	9	2	0.8660.866
B Shift	16	3.2	1.00

Is there a difference in the mean amount of time for each shift to process 100 pounds of coconuts? Test at the 5% level of significance.

Answer

Solution 10.1

The population standard deviations are not known and cannot be assumed to equal each other. Let g be the subscript for the G Shift and b be the subscript for the B Shift. Then, μ_g is the population mean for G Shift and μ_b is the population mean for B Shift. This is a test of two **independent groups**, two population **means**.

Random variable: $\bar{X}_g - \bar{X}_b$ = difference in the sample mean amount of time between the G Shift and the B Shift takes to process the coconuts.

$$\begin{aligned} \backslash H_0 : \mu_g = \mu_b & \quad \backslash H_0 : \mu_g - \mu_b = 0 \\ H_a : \mu_g \neq \mu_b & \quad H_a : \mu_g - \mu_b \neq 0 \end{aligned}$$

The words "**the same**" tell you $\backslash H_0$ has an "=". Since there are no other words to indicate H_a , is either faster or slower. This is a two tailed test.

Distribution for the test: Use t_{df} where df is calculated using the df formula for independent groups, two population means above. Using a calculator, df is approximately 18.8462.

Graph:




Figure 10.2.3

$$t_c = \frac{(\bar{X}_1 - \bar{X}_2) - \delta_0}{\sqrt{\frac{S_1^2}{n_1} + \frac{S_2^2}{n_2}}} = -3.01$$

We next find the critical value on the t -table using the degrees of freedom from above. The critical value, 2.093, is found in the .025 column, this is $\alpha/2$, at 19 degrees of freedom. (The convention is to round up the degrees of freedom to make the conclusion more conservative.) Next we calculate the test statistic and mark this on the t -distribution graph.

Make a decision: Since the calculated t -value is in the tail we cannot accept the null hypothesis that there is no difference between the two groups. The means are different.

The graph has included the sampling distribution of the differences in the sample means to show how the t -distribution aligns with the sampling distribution data. We see in the top panel that the calculated difference in the two means is -1.2 and the bottom panel shows that this is 3.01 standard deviations from the mean. Typically we do not need to show the sampling distribution graph and can rely on the graph of the test statistic, the t -distribution in this case, to reach our conclusion.

Conclusion: At the 5% level of significance, the sample data show there is sufficient evidence to conclude that the mean number of hours that the G Shift takes to process 100 pounds of coconuts is different from the B Shift (mean number of hours for the B Shift is greater than the mean number of hours for the G Shift).

NOTE

When the sum of the sample sizes is larger than 30 ($n_1 + n_2 > 30$) you can use the normal distribution to approximate the Student's t .

Example 10.2.2

A study is done to determine if Company A retains its workers longer than Company B. It is believed that Company A has a higher retention than Company B. The study finds that in a sample of 11 workers at Company A their average time with the company is four years with a standard deviation of 1.5 years. A sample of 9 workers at Company B finds that the average time with the company was 3.5 years with a standard deviation of 1 year. Test this proposition at the 1% level of significance.

a. Is this a test of two means or two proportions?

Answer

Solution 10.2

a. two means because time is a continuous random variable.

b. Are the populations standard deviations known or unknown?

Answer

Solution 10.2

b. unknown

c. Which distribution do you use to perform the test?

Answer

Solution 10.2

c. Student's t

d. What is the random variable?

Answer

Solution 10.2

d. $\bar{X}_A - \bar{X}_B$

e. What are the null and alternate hypotheses?

Answer

Solution 10.2

e.

- $H_0 : \mu_A \leq \mu_B$
- $H_a : \mu_A > \mu_B$

f. Is this test right-, left-, or two-tailed?

Answer

Solution 10.2

f. right one-tailed test




Figure 10.2.4

g. What is the value of the test statistic?

Answer

Solution 10.2

g.

$$t_c = \frac{(\bar{X}_1 - \bar{X}_2) - \delta_0}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}} = 0.89$$

h. Can you accept/reject the null hypothesis?

Answer

Solution 10.2

h. Cannot reject the null hypothesis that there is no difference between the two groups. Test statistic is not in the tail. The critical value of the t distribution is 2.764 with 10 degrees of freedom. This example shows how difficult it is to reject a null hypothesis with a very small sample. The critical values require very large test statistics to reach the tail.

i. Conclusion:

Answer

Solution 10.2

i. At the 1% level of significance, from the sample data, there is not sufficient evidence to conclude that the retention of workers at Company A is longer than Company B, on average.

Example 10.2.3

An interesting research question is the effect, if any, that different types of teaching formats have on the grade outcomes of students. To investigate this issue one sample of students' grades was taken from a hybrid class and another sample taken from a standard lecture format class. Both classes were for the same subject. The mean course grade in percent for the 35 hybrid students is 74 with a standard deviation of 16. The mean grades of the 40 students from the standard lecture class was 76 percent with a standard deviation of 9. Test at 5% to see if there is any significant difference in the population mean grades between standard lecture course and hybrid class.

Answer

Solution 10.3

We begin by noting that we have two groups, students from a hybrid class and students from a standard lecture format class. We also note that the random variable, what we are interested in, is students' grades, a continuous random variable. We could have asked the research question in a different way and had a binary random variable. For example, we could have studied the percentage of students with a failing grade, or with an A grade. Both of these would be binary and thus a test of proportions and not a test of means as is the case here. Finally, there is no presumption as to which format might lead to higher grades so the hypothesis is stated as a two-tailed test.

$$H_0 : \mu_1 = \mu_2$$

$$H_a : \mu_1 \neq \mu_2$$

As would virtually always be the case, we do not know the population variances of the two distributions and thus our test statistic is:

$$t_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{\frac{s^2}{n_1} + \frac{s^2}{n_2}}} = \frac{(74 - 76) - 0}{\sqrt{\frac{16^2}{35} + \frac{9^2}{40}}} = -0.65$$

To determine the critical value of the Student's t we need the degrees of freedom. For this case we use: $df = n_1 + n_2 - 2 = 35 + 40 - 2 = 73$. This is large enough to consider it the normal distribution thus $t_{\alpha/2} = 1.96$. Again as always we determine if the calculated value is in the tail determined by the critical value. In this case we do not even need to look up the critical value: the calculated value of the difference in these two average grades is not even one standard deviation apart. Certainly not in the tail.

Conclusion: Cannot reject the null at $\alpha = 5\%$. Therefore, evidence does not exist to prove that the grades in hybrid and standard classes differ.

This page titled [10.2: Comparing Two Independent Population Means](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [10.1: Comparing Two Independent Population Means](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.3: Cohen's Standards for Small, Medium, and Large Effect Sizes

Cohen's d is a measure of "effect size" based on the differences between two means. Cohen's d , named for United States statistician Jacob Cohen, measures the relative strength of the differences between the means of two populations based on sample data. The calculated value of effect size is then compared to Cohen's standards of small, medium, and large effect sizes.

Table 10.2 Cohen's Standard Effect Sizes

Size of effect	d
Small	0.2
Medium	0.5
Large	0.8

Cohen's d is the measure of the difference between two means divided by the pooled standard deviation: $d = \frac{\bar{x}_1 - \bar{x}_2}{s_{pooled}}$ where

$$s_{pooled} = \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}}$$

It is important to note that Cohen's d does not provide a level of confidence as to the magnitude of the size of the effect comparable to the other tests of hypothesis we have studied. The sizes of the effects are simply indicative.

The effect is small because 0.384 is between Cohen's value of 0.2 for small effect size and 0.5 for medium effect size. The size of the differences of the means for the two companies is small indicating that there is not a significant difference between them.

This page titled [10.3: Cohen's Standards for Small, Medium, and Large Effect Sizes](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [10.2: Cohen's Standards for Small, Medium, and Large Effect Sizes](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.4: Test for Differences in Means- Assuming Equal Population Variances

Typically we can never expect to know any of the population parameters, mean, proportion, or standard deviation. When testing hypotheses concerning differences in means we are faced with the difficulty of two unknown variances that play a critical role in the test statistic. We have been substituting the sample variances just as we did when testing hypotheses for a single mean. And as we did before, we used a Student's t to compensate for this lack of information on the population variance. There may be situations, however, when we do not know the population variances, but we can assume that the two populations have the same variance. If this is true then the pooled sample variance will be smaller than the individual sample variances. This will give more precise estimates and reduce the probability of discarding a good null. The null and alternative hypotheses remain the same, but the test statistic changes to:

$$t_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{S_p^2 \left(\frac{1}{n_1} + \frac{1}{n_2} \right)}}$$

where S_p^2 is the pooled variance given by the formula:

$$S_p^2 = \frac{(n_1 - 1) s_1^2 + (n_2 - 1) s_2^2}{n_1 + n_2 - 2}$$

The test statistic is clearly in the tail, 2.31 is larger than the critical value of 1.703, and therefore we cannot maintain the null hypothesis. Thus, we conclude that there is significant evidence at the 95% level of confidence that the new medicine produces the effect desired.

This page titled [10.4: Test for Differences in Means- Assuming Equal Population Variances](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [10.3: Test for Differences in Means- Assuming Equal Population Variances](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.5: Comparing Two Independent Population Proportions

When conducting a hypothesis test that compares two independent population proportions, the following characteristics should be present:

1. The two independent samples are random samples that are independent.
2. The number of successes is at least five, and the number of failures is at least five, for each of the samples.
3. Growing literature states that the population must be at least ten or even perhaps 20 times the size of the sample. This keeps each population from being over-sampled and causing biased results.

Comparing two proportions, like comparing two means, is common. If two estimated proportions are different, it may be due to a difference in the populations or it may be due to chance in the sampling. A hypothesis test can help determine if a difference in the estimated proportions reflects a difference in the two population proportions.

Like the case of differences in sample means, we construct a sampling distribution for differences in sample proportions: $(p'_A - p'_B)$ where $p'_A = \frac{X_A}{n_A}$ and $p'_B = \frac{X_B}{n_B}$ are the sample proportions for the two sets of data in question. X_A and X_B are the number of successes in each sample group respectively, and n_A and n_B are the respective sample sizes from the two groups. Again we go the Central Figure 10.5.5




Figure 10.5.5

Generally, the null hypothesis allows for the test of a difference of a particular value, δ_0 , just as we did for the case of differences in means.

$$H_0 : p_1 - p_2 = \delta_0$$

$$H_1 : p_1 - p_2 \neq \delta_0$$

Most common, however, is the test that the two proportions are the same. That is,

$$H_0 : p_A = p_B$$

$$H_a : p_A \neq p_B$$

To conduct the test, we use a pooled proportion, p_c .

The pooled proportion is calculated as follows:

$$p_c = \frac{x_A + x_B}{n_A + n_B}$$

The test statistic (z-score) is:

$$Z_c = \frac{(p'_A - p'_B) - \delta_0}{\sqrt{p_c(1-p_c)\left(\frac{1}{n_A} + \frac{1}{n_B}\right)}}$$

where δ_0 is the hypothesized differences between the two proportions and p_c is the pooled variance from the formula above.

Example 10.5.6

A bank has recently acquired a new branch and thus has customers in this new territory. They are interested in the default rate in their new territory. They wish to test the hypothesis that the default rate is different from their current customer base. They sample 200 files in area A, their current customers, and find that 20 have defaulted. In area B, the new customers, another sample of 200 files shows 12 have defaulted on their loans. At a 10% level of significance can we say that the default rates are the same or different?

Answer

Solution 10.6

This is a test of proportions. We know this because the underlying random variable is binary, default or not default. Further, we know it is a test of differences in proportions because we have two sample groups, the current customer base and the newly acquired customer base. Let A and B be the subscripts for the two customer groups. Then p_A and p_B are the two population proportions we wish to test.

Random Variable:

$P'_A - P'_B$ = difference in the proportions of customers who defaulted in the two groups.

$$H_0 : p_A = p_B$$

$$H_a : p_A \neq p_B$$

The words "is a difference" tell you the test is two-tailed.

Distribution for the test: Since this is a test of two binomial population proportions, the distribution is normal:

$$p_c = \frac{x_A + x_B}{n_A + n_B} = \frac{20 + 12}{200 + 200} = 0.08 \quad 1 - p_c = 0.92$$

$(p'_A - p'_B) = 0.04$ follows an approximate normal distribution.

$$\text{Estimated proportion for group A: } p'_A = \frac{x_A}{n_A} = \frac{20}{200} = 0.1$$

$$\text{Estimated proportion for group B: } p'_B = \frac{x_B}{n_B} = \frac{12}{200} = 0.06$$

The estimated difference between the two groups is : $p'_A - p'_B = 0.1 - 0.06 = 0.04$.




Figure 10.5.6

$$Z_c = \frac{(P'_A - P'_B) - \delta_0}{P_c (1 - P_c) \left(\frac{1}{n_A} + \frac{1}{n_B} \right)} = 0.54$$

The calculated test statistic is .54 and is not in the tail of the distribution.

Make a decision: Since the calculate test statistic is not in the tail of the distribution we cannot reject H_0 .

Conclusion: At a 1% level of significance, from the sample data, there is not sufficient evidence to conclude that there is a difference between the proportions of customers who defaulted in the two groups.

Exercise 10.5.6

Two types of valves are being tested to determine if there is a difference in pressure tolerances. Fifteen out of a random sample of 100 of Valve A cracked under 4,500 psi. Six out of a random sample of 100 of Valve B cracked under 4,500 psi. Test at a 5% level of significance.

This page titled [10.5: Comparing Two Independent Population Proportions](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [10.4: Comparing Two Independent Population Proportions](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.6: Two Population Means with Known Standard Deviations

Even though this situation is not likely (knowing the population standard deviations is very unlikely), the following example illustrates hypothesis testing for independent means with known population standard deviations. The sampling distribution for the difference between the means is normal in accordance with the central limit theorem. The random variable is $\bar{X}_1 - \bar{X}_2$. The normal distribution has the following format:

The standard deviation is:

$$\sqrt{\frac{(\sigma_1)^2}{n_1} + \frac{(\sigma_2)^2}{n_2}}$$

The test statistic (z-score) is:

$$Z_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{\frac{(\sigma_1)^2}{n_1} + \frac{(\sigma_2)^2}{n_2}}}$$

At the 5% level of significance, from the sample data, there is not sufficient evidence to conclude that the mean age of Democratic senators is greater than the mean age of the Republican senators.

This page titled [10.6: Two Population Means with Known Standard Deviations](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [10.5: Two Population Means with Known Standard Deviations](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.7: Matched or Paired Samples

In most cases of economic or business data we have little or no control over the process of how the data are gathered. In this sense the data are not the result of a planned controlled experiment. In some cases, however, we can develop data that are part of a controlled experiment. This situation occurs frequently in quality control situations. Imagine that the production rates of two machines built to the same design, but at different manufacturing plants, are being tested for differences in some production metric such as speed of output or meeting some production specification such as strength of the product. The test is the same in format to what we have been testing, but here we can have matched pairs for which we can test if differences exist. Each observation has its matched pair against which differences are calculated. First, the differences in the metric to be tested between the two lists of observations must be calculated, and this is typically labeled with the letter "d." Then, the average of these matched differences, \bar{X}_d is calculated as is its standard deviation, S_d . We expect that the standard deviation of the differences of the matched pairs will be smaller than unmatched pairs because presumably fewer differences should exist because of the correlation between the two groups.

When using a hypothesis test for matched or paired samples, the following characteristics may be present:

1. In a hypothesis test for matched or paired samples, subjects are matched in pairs and differences are calculated. The differences are the data. The population mean for the differences, μ_d , is then tested using a Student's-t test for a single population mean with $n-1$ degrees of freedom, where n is the number of differences, that is, the number of pairs not the number of observations.

The null and alternative hypotheses for this test are:

$$H_a : \mu_d \neq 0$$

The test statistic is:

$$t_c = \frac{\bar{x}_d - \mu_d}{\left(\frac{s_d}{\sqrt{n}}\right)}$$

At a 5% level of significance, from the sample data, there is not sufficient evidence to conclude that the strength development class helped to make the players stronger, on average.

This page titled [10.7: Matched or Paired Samples](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.6: Matched or Paired Samples** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.8: Chapter Key Terms

Cohen's d

a measure of effect size based on the differences between two means. If d is between 0 and 0.2 then the effect is small. If d approaches 0.5, then the effect is medium, and if d approaches 0.8, then it is a large effect.

Independent Groups

two samples that are selected from two populations, and the values from one population are not related in any way to the values from the other population.

Matched Pairs

two samples that are dependent. Differences between a before and after scenario are tested by testing one population mean of differences.

Pooled Variance

a weighted average of two variances that can then be used when calculating standard error.

This page titled [10.8: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

10.9: Chapter Review

10.1 Comparing Two Independent Population Means

Two population means from independent samples where the population standard deviations are not known

- Random Variable: $\bar{X}_1 - \bar{X}_2$ = the difference of the sampling means
- Distribution: Student's t -distribution with degrees of freedom (variances not pooled)

10.2 Cohen's Standards for Small, Medium, and Large Effect Sizes

Cohen's d is a measure of “effect size” based on the differences between two means.

It is important to note that Cohen's d does not provide a level of confidence as to the magnitude of the size of the effect comparable to the other tests of hypothesis we have studied. The sizes of the effects are simply indicative.

10.3 Test for Differences in Means: Assuming Equal Population Variances

In situations when we do not know the population variances but assume the variances are the same, the pooled sample variance will be smaller than the individual sample variances.

This will give more precise estimates and reduce the probability of discarding a good null.

10.4 Comparing Two Independent Population Proportions

Test of two population proportions from independent samples.

- Random variable: $\mathbf{p}'_A - \mathbf{p}'_B$ = difference between the two estimated proportions
- Distribution: normal distribution

10.5 Two Population Means with Known Standard Deviations

A hypothesis test of two population means from independent samples where the population standard deviations are known (typically approximated with the sample standard deviations), will have these characteristics:

- Random variable: $\bar{X}_1 - \bar{X}_2$ = the difference of the means
- Distribution: normal distribution

10.6 Matched or Paired Samples

A hypothesis test for matched or paired samples (t-test) has these characteristics:

- Test the differences by subtracting one measurement from the other measurement
- Random Variable: \bar{x}_d = mean of the differences
- Distribution: Student's t -distribution with $n-1$ degrees of freedom
- If the number of differences is small (less than 30), the differences must follow a normal distribution.
- Two samples are drawn from the same set of objects.
- Samples are dependent.

This page titled [10.9: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.13: Chapter Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.10: Chapter Formula Review

10.1 Comparing Two Independent Population Means

Standard error: $SE = \sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}$

Test statistic (t-score): $t_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}}}$

Degrees of freedom:

$$df = \frac{\left(\frac{(s_1)^2}{n_1} + \frac{(s_2)^2}{n_2}\right)^2}{\left(\frac{1}{n_1-1}\right)\left(\frac{(s_1)^2}{n_1}\right)^2 + \left(\frac{1}{n_2-1}\right)\left(\frac{(s_2)^2}{n_2}\right)^2}$$

where:

s_1 and s_2 are the sample standard deviations, and n_1 and n_2 are the sample sizes.

\bar{x}_1 and \bar{x}_2 are the sample means.

10.2 Cohen's Standards for Small, Medium, and Large Effect Sizes

Cohen's d is the measure of effect size:

$$d = \frac{\bar{x}_1 - \bar{x}_2}{s_{\text{pooled}}}$$

where $s_{\text{pooled}} = \sqrt{\frac{(n_1-1)s_1^2 + (n_2-1)s_2^2}{n_1+n_2-2}}$

10.3 Test for Differences in Means: Assuming Equal Population Variances

$$t_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{S_p^2 \left(\frac{1}{n_1} + \frac{1}{n_2}\right)}}$$

where S_p^2 is the pooled variance given by the formula:

$$S_p^2 = \frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}$$

10.4 Comparing Two Independent Population Proportions

Pooled Proportion: $p_c = \frac{x_A + x_B}{n_A + n_B}$

Test Statistic (z-score): $Z_c = \frac{(p'_A - p'_B)}{\sqrt{p_c(1-p_c)\left(\frac{1}{n_A} + \frac{1}{n_B}\right)}}$

where

p'_A and p'_B are the sample proportions, p_A and p_B are the population proportions,

P_c is the pooled proportion, and n_A and n_B are the sample sizes.

10.5 Two Population Means with Known Standard Deviations

Test Statistic (z-score):

$$Z_c = \frac{(\bar{x}_1 - \bar{x}_2) - \delta_0}{\sqrt{\frac{(\sigma_1)^2}{n_1} + \frac{(\sigma_2)^2}{n_2}}}$$

where:

σ_1 and σ_2 are the known population standard deviations. n_1 and n_2 are the sample sizes. \bar{x}_1 and \bar{x}_2 are the sample means. μ_1 and

μ_2 are the population means.

10.6 Matched or Paired Samples

Test Statistic (t-score): $t_c = \frac{\bar{x}_d - \mu_d}{\left(\frac{s_d}{\sqrt{n}}\right)}$

where:

\bar{x}_d is the mean of the sample differences. μ_d is the mean of the population differences. s_d is the sample standard deviation of the differences. n is the sample size.

This page titled [10.10: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.8: Chapter Formula Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.11: Chapter Practice

10.1 Comparing Two Independent Population Means

Use the following information to answer the next 15 exercises: Indicate if the hypothesis test is for

- a. Use the following information to answer the next three exercises: A study is done to determine which of two soft drinks has more sugar. There are 13 cans of Beverage A in a sample and six cans of Beverage B. The mean amount of sugar in Beverage A is 36 grams with a standard deviation of 0.6 grams. The mean amount of sugar in Beverage B is 38 grams with a standard deviation of 0.8 grams. The researchers believe that Beverage B has more sugar than Beverage A, on average. Both populations have normal distributions.**16.**

Are standard deviations known or unknown?

17.

What is the random variable?

18.

Is this a one-tailed or two-tailed test?

19.

Is this a test of means or proportions?

20.

State the null and alternative hypotheses.

1. Table **10.844**.

What is the random variable?

45.

State the null and alternative hypotheses.

46.

What is the test statistic?

47.

At the 1% significance level, what is your conclusion?

Table 10.11.9

Plant group	Sample mean height of plants (inches)	Population standard deviation
Food	16	2.5
No food	14	1.5

48.

Is the population standard deviation known or unknown?

49.

State the null and alternative hypotheses.

50.

At the 1% significance level, what is your conclusion?

Use the following information to answer the next five exercises. Two metal alloys are being considered as material for ball bearings. The mean melting point of the two alloys is to be compared. 15 pieces of each metal are being tested. Both populations have normal distributions. The following table is the result. It is believed that Alloy Zeta has a different melting point.

	Sample mean melting temperatures (°F)	Population standard deviation
Alloy Gamma	800	95
Alloy Zeta	900	105

Table 10.10

51.

State the null and alternative hypotheses.

52.

Is this a right-, left-, or two-tailed test?

53.

At the 1% significance level, what is your conclusion?

10.6 Matched or Paired Samples

Use the following information to answer the next five exercises. A study was conducted to test the effectiveness of a software patch in reducing system failures over a six-month period. Results for randomly selected installations are shown in Table 10.11.11 The “before” value is matched to an “after” value, and the differences are calculated. The differences have a normal distribution. Test at the 1% significance level.

Installation	A	B	C	D	E	F	G	H
Before	3	6	4	2	5	8	2	6
After	1	5	2	0	1	0	2	2

Table 10.11

54.

What is the random variable?

55.

State the null and alternative hypotheses.

56.

What conclusion can you draw about the software patch?

Use the following information to answer next five exercises. A study was conducted to test the effectiveness of a juggling class. Before the class started, six subjects juggled as many balls as they could at once. After the class, the same six subjects juggled as many balls as they could. The differences in the number of balls are calculated. The differences have a normal distribution. Test at the 1% significance level.

Table 10.11.12

Subject	A	B	C	D	E	F
Before	3	4	3	2	4	5
After	4	5	6	4	5	7

57.

State the null and alternative hypotheses.

58.

What is the sample mean difference?

59.

What conclusion can you draw about the juggling class?

Use the following information to answer the next five exercises. A doctor wants to know if a blood pressure medication is effective. Six subjects have their blood pressures recorded. After twelve weeks on the medication, the same six subjects have their blood pressure recorded again. For this test, only systolic pressure is of concern. Test at the 1% significance level.

Patient	A	B	C	D	E	F
Before	161	162	165	162	166	171
After	158	159	166	160	167	169

Table 10.13

60.

State the null and alternative hypotheses.

61.

What is the test statistic?

62.

What is the sample mean difference?

63.

What is the conclusion?

This page titled [10.11: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [10.11: Chapter Practice](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.12: Homework

Use the following information to answer the next ten exercises. indicate which of the following choices best identifies the hypothesis test.

a. Table **10.25****115**.

University of Michigan researchers reported in the *Journal of the National Cancer Institute* that quitting smoking is especially beneficial for those under age 49. In this American Cancer Society study, the risk (probability) of dying of lung cancer was about the same as for those who had never smoked.

116.

Lesley E. Tan investigated the relationship between left-handedness vs. right-handedness and motor competence in preschool children. Random samples of 41 left-handed preschool children and 41 right-handed preschool children were given several tests of motor skills to determine if there is evidence of a difference between the children based on this experiment. The experiment produced the means and standard deviations shown Table 10.12.26 Determine the appropriate test and best distribution to use for that test.

Table 10.12.26

	Left-handed	Right-handed
Sample size	41	41
Sample mean	97.5	98.1
Sample standard deviation	17.5	19.2

a. This is:

- a test of two independent means.
- a test of two proportions.
- a test of a single mean.
- a test of a single proportion.

This page titled [10.12: Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- 10.7: Homework** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.13: Chapter Homework

10.1 Comparing Two Independent Population Means

64.

The mean number of English courses taken in a two-year time period by male and female college students is believed to be about the same. An experiment is conducted and data are collected from 29 males and 16 females. The males took an average of three English courses with a standard deviation of 0.8. The females took an average of four English courses with a standard deviation of 1.0. Are the means statistically the same?

65.

A student at a four-year college claims that mean enrollment at four-year colleges is higher than at two-year colleges in the United States. Two surveys are conducted. Of the 35 two-year colleges surveyed, the mean enrollment was 5,068 with a standard deviation of 4,777. Of the 35 four-year colleges surveyed, the mean enrollment was 5,466 with a standard deviation of 8,191.

66.

At Rachel's 11th birthday party, eight girls were timed to see how long (in seconds) they could hold their breath in a relaxed position. After a two-minute rest, they timed themselves while jumping. The girls thought that the mean difference between their jumping and relaxed times would be zero. Test their hypothesis.

Relaxed time (seconds)	Jumping time (seconds)
26	21
47	40
30	28
22	21
23	25
45	43
37	35
29	32

Table 10.14

67.

Mean entry-level salaries for college graduates with mechanical engineering degrees and electrical engineering degrees are believed to be approximately the same. A recruiting office thinks that the mean mechanical engineering salary is actually lower than the mean electrical engineering salary. The recruiting office randomly surveys 50 entry level mechanical engineers and 60 entry level electrical engineers. Their mean salaries were \$46,100 and \$46,700, respectively. Their standard deviations were \$3,450 and \$4,210, respectively. Conduct a hypothesis test to determine if you agree that the mean entry-level mechanical engineering salary is lower than the mean entry-level electrical engineering salary.

68.

Marketing companies have collected data implying that teenage girls use more ring tones on their cellular phones than teenage boys do. In one particular study of 40 randomly chosen teenage girls and boys (20 of each) with cellular phones, the mean number of ring tones for the girls was 3.2 with a standard deviation of 1.5. The mean for the boys was 1.7 with a standard deviation of 0.8. Conduct a hypothesis test to determine if the means are approximately the same or if the girls' mean is higher than the boys' mean.

Use the information from [Appendix C: Data Sets](#) to answer the next four exercises.

69.

Using the data from Lap 1 only, conduct a hypothesis test to determine if the mean time for completing a lap in races is the same as it is in practices.

70.

Repeat the test in Table 10.13.16 Test at the 1% level of significance.

	Number who are obese	Sample size
Men	42,769	155,525
Women	67,169	248,775

Table 10.16

87.

Two computer users were discussing tablet computers. A higher proportion of people ages 16 to 29 use tablets than the proportion of people age 30 and older. Table 10.13.17 details the number of tablet owners for each age group. Test at the 1% level of significance.

Table 10.13.17

	16–29 year olds	30 years old and older
Own a tablet	69	231
Sample size	628	2,309

88.

A group of friends debated whether more men use smartphones than women. They consulted a research study of smartphone use among adults. The results of the survey indicate that of the 973 men randomly sampled, 379 use smartphones. For women, 404 of the 1,304 who were randomly sampled use smartphones. Test at the 5% level of significance.

89.

While her husband spent 2½ hours picking out new speakers, a statistician decided to determine whether the percent of men who enjoy shopping for electronic equipment is higher than the percent of women who enjoy shopping for electronic equipment. The population was Saturday afternoon shoppers. Out of 67 men, 24 said they enjoyed the activity. Eight of the 24 women surveyed claimed to enjoy the activity. Interpret the results of the survey.

90.

We are interested in whether children’s educational computer software costs less, on average, than children’s entertainment software. Thirty-six educational software titles were randomly picked from a catalog. The mean cost was \$31.14 with a standard deviation of \$4.69. Thirty-five entertainment software titles were randomly picked from the same catalog. The mean cost was \$33.86 with a standard deviation of \$10.87. Decide whether children’s educational software costs less, on average, than children’s entertainment software.

91.

Joan Nguyen recently claimed that the proportion of college-age males with at least one pierced ear is as high as the proportion of college-age females. She conducted a survey in her classes. Out of 107 males, 20 had at least one pierced ear. Out of 92 females, 47 had at least one pierced ear. Do you believe that the proportion of males has reached the proportion of females?

92.

"To Breakfast or Not to Breakfast?" by Richard Ayore

In the American society, birthdays are one of those days that everyone looks forward to. People of different ages and peer groups gather to mark the 18th, 20th, ..., birthdays. During this time, one looks back to see what he or she has achieved for the past year and also focuses ahead for more to come.

If, by any chance, I am invited to one of these parties, my experience is always different. Instead of dancing around with my friends while the music is booming, I get carried away by memories of my family back home in Kenya. I remember the good times I had with my brothers and sister while we did our daily routine.

Every morning, I remember we went to the shamba (garden) to weed our crops. I remember one day arguing with my brother as to why he always remained behind just to join us an hour later. In his defense, he said that he preferred waiting for breakfast before he came to weed. He said, "This is why I always work more hours than you guys!"

And so, to prove him wrong or right, we decided to give it a try. One day we went to work as usual without breakfast, and recorded the time we could work before getting tired and stopping. On the next day, we all ate breakfast before going to work. We recorded how long we worked again before getting tired and stopping. Of interest was our mean increase in work time. Though not sure, my brother insisted that it was more than two hours. Using the data in Table 10.13.18 solve our problem.

Table 10.13.18

Work hours with breakfast	Work hours without breakfast
8	6
7	5
9	5
5	4
9	7
8	7
10	7
7	5
6	6
9	5

10.5 Two Population Means with Known Standard Deviations

NOTE

If you are using a Student's t -distribution for one of the following homework problems, including for paired data, you may assume that the underlying population is normally distributed. (When using these tests in a real situation, you must first prove that assumption, however.)

93.

A study is done to determine if students in the California state university system take longer to graduate, on average, than students enrolled in private universities. One hundred students from both the California state university system and private universities are surveyed. Suppose that from years of research, it is known that the population standard deviations are 1.5811 years and 1 year, respectively. The following data are collected. The California state university system students took on average 4.5 years with a standard deviation of 0.8. The private university students took on average 4.1 years with a standard deviation of 0.3.

94.

Parents of teenage boys often complain that auto insurance costs more, on average, for teenage boys than for teenage girls. A group of concerned parents examines a random sample of insurance bills. The mean annual cost for 36 teenage boys was \$679. For 23 teenage girls, it was \$559. From past years, it is known that the population standard deviation for each group is \$180. Determine whether or not you believe that the mean cost for auto insurance for teenage boys is greater than that for teenage girls.

95.

A group of transfer bound students wondered if they will spend the same mean amount on texts and supplies each year at their four-year university as they have at their community college. They conducted a random survey of 54 students at their community college and 66 students at their local four-year university. The sample means were \$947 and \$1,011, respectively. The population standard deviations are known to be \$254 and \$87, respectively. Conduct a hypothesis test to determine if the means are statistically the same.

96.

Some manufacturers claim that non-hybrid sedan cars have a lower mean miles-per-gallon (mpg) than hybrid ones. Suppose that consumers test 21 hybrid sedans and get a mean of 31 mpg with a standard deviation of seven mpg. Thirty-one non-hybrid sedans get a mean of 22 mpg with a standard deviation of four mpg. Suppose that the population standard deviations are known to be six and three, respectively. Conduct a hypothesis test to evaluate the manufacturers claim.

97.

A baseball fan wanted to know if there is a difference between the number of games played in a World Series when the American League won the series versus when the National League won the series. From 1922 to 2012, the population standard deviation of games won by the American League was 1.14, and the population standard deviation of games won by the National League was 1.11. Of 19 randomly selected World Series games won by the American League, the mean number of games won was 5.76. The mean number of 17 randomly selected games won by the National League was 5.42. Conduct a hypothesis test.

98.

One of the questions in a study of marital satisfaction of dual-career couples was to rate the statement “I’m pleased with the way we divide the responsibilities for childcare.” The ratings went from one (strongly agree) to five (strongly disagree). Table 10.13.19 contains ten of the paired responses for husbands and wives. Conduct a hypothesis test to see if the mean difference in the husband’s versus the wife’s satisfaction level is negative (meaning that, within the partnership, the husband is happier than the wife).

Wife’s score	2	2	3	3	4	2	1	1	2	4
Husband’s score	2	2	1	3	2	1	1	1	2	4

Table 10.19

10.6 Matched or Paired Samples

99.

Ten individuals went on a low-fat diet for 12 weeks to lower their cholesterol. The data are recorded in Table 10.13.20 Do you think that their cholesterol levels were significantly lowered?

Table 10.13.20

Starting cholesterol level	Ending cholesterol level
140	140
220	230
110	120
240	220
200	190
180	150
190	200
360	300
280	300
260	240

Use the following information to answer the next two exercises. A new AIDS prevention drug was tried on a group of 224 HIV positive patients. Forty-five patients developed AIDS after four years. In a control group of 224 HIV positive patients, 68 developed AIDS after four years. We want to test whether the method of treatment reduces the proportion of patients that develop AIDS after four years or if the proportions of the treated group and the untreated group stay the same.

Let the subscript t = treated patient and ut = untreated patient.

100.

The appropriate hypotheses are:

- Use the following information to answer the next two exercises. An experiment is conducted to show that blood pressure can be consciously reduced in people trained in a “biofeedback exercise program.” Six subjects were randomly selected and blood

pressure measurements were recorded before and after the training. The difference between blood pressures was calculated (after - before) producing the following results: $\bar{x}_d = -10.2$ $s_d = 8.4$. Using the data, test the hypothesis that the blood pressure has decreased after the training. **101.**

The distribution for the test is:

a. The correct decision is:

1. Table **10.23105**.

A politician asked his staff to determine whether the underemployment rate in the northeast decreased from 2011 to 2012. The results are in Table 10.13.24

Table 10.13.24

Northeastern states	2011	2012
Connecticut	17.3	16.4
Delaware	17.4	13.7
Maine	19.3	16.1
Maryland	16.0	15.5
Massachusetts	17.6	18.2
New Hampshire	15.4	13.5
New Jersey	19.2	18.7
New York	18.5	18.7
Ohio	18.2	18.8
Pennsylvania	16.5	16.9
Rhode Island	20.7	22.4
Vermont	14.7	12.3
West Virginia	15.5	17.3

This page titled [10.13: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.9: Chapter Homework** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.14: Chapter References

10.1 Comparing Two Independent Population Means

Data from Graduating Engineer + Computer Careers. Available online at www.graduatingengineer.com

Data from *Microsoft Bookshelf*.

Data from the United States Senate website, available online at www.Senate.gov (accessed June 17, 2013).

“List of current United States Senators by Age.” Wikipedia. Available online at en.Wikipedia.org/wiki/List_of...enators_by_age (accessed June 17, 2013).

“Sectoring by Industry Groups.” Nasdaq. Available online at www.nasdaq.com/markets/barcha...&base=industry (accessed June 17, 2013).

“Strip Clubs: Where Prostitution and Trafficking Happen.” Prostitution Research and Education, 2013. Available online at www.prostitutionresearch.com/ProsViolPosttrauStress.html (accessed June 17, 2013).

“World Series History.” Baseball-Almanac, 2013. Available online at <http://www.baseball-almanac.com/ws/wsmenu.shtml> (accessed June 17, 2013).

10.4 Comparing Two Independent Population Proportions

Data from *Educational Resources*, December catalog.

Data from Hilton Hotels. Available online at <http://www.hilton.com> (accessed June 17, 2013).

Data from Hyatt Hotels. Available online at hyatt.com (accessed June 17, 2013).

Data from Statistics, United States Department of Health and Human Services.

Data from Whitney Exhibit on loan to San Jose Museum of Art.

Data from the American Cancer Society. Available online at <http://www.cancer.org/index> (accessed June 17, 2013).

Data from the Chancellor’s Office, California Community Colleges, November 1994.

“State of the States.” Gallup, 2013. Available online at <http://www.gallup.com/poll/125066/St...ef=interactive> (accessed June 17, 2013).

“West Nile Virus.” Centers for Disease Control and Prevention. Available online at <http://www.cdc.gov/ncidod/dvbid/westnile/index.htm> (accessed June 17, 2013).

10.5 Two Population Means with Known Standard Deviations

Data from the United States Census Bureau. Available online at www.census.gov/prod/cen2010/b...c2010br-02.pdf

Hinduja, Sameer. “Sexting Research and Gender Differences.” Cyberbullying Research Center, 2013. Available online at <http://cyberbullying.us/blog/sexting...r-differences/> (accessed June 17, 2013).

“Smart Phone Users, By the Numbers.” Visually, 2013. Available online at <http://visual.ly/smart-phone-users-numbers> (accessed June 17, 2013).

Smith, Aaron. “35% of American adults own a Smartphone.” Pew Internet, 2013. Available online at www.pewinternet.org/~media/F...martphones.pdf (accessed June 17, 2013).

“State-Specific Prevalence of Obesity Among Adults—United States, 2007.” MMWR, CDC. Available online at <http://www.cdc.gov/mmwr/preview/mmwrhtml/mm5728a1.htm> (accessed June 17, 2013).

“Texas Crime Rates 1960–2012.” FBI, Uniform Crime Reports, 2013. Available online at: <http://www.disastercenter.com/crime/txcrime.htm> (accessed June 17, 2013).

This page titled [10.14: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.12: Chapter References** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

10.15: Chapter Solution (Practice + Homework)

1.

two proportions

3.

matched or paired samples

5.

single mean

7.

independent group means, population standard deviations and/or variances unknown

9.

two proportions

11.

independent group means, population standard deviations and/or variances unknown

13.

independent group means, population standard deviations and/or variances unknown

15.

two proportions

17.

The random variable is the difference between the mean amounts of sugar in the two soft drinks.

19.

means

21.

two-tailed

23.

the difference between the mean life spans of whites and nonwhites

25.

This is a comparison of two population means with unknown population standard deviations.

27.

Check student's solution.

28.

1. 31.

$P'_{OS1} - P'_{OS2}$ = difference in the proportions of phones that had system failures within the first eight hours of operation with OS_1 and OS_2 .

proportions

36.

right-tailed

38.

The random variable is the difference in proportions (percents) of the populations that are of two or more races in Nevada and North Dakota.

40.

Our sample sizes are much greater than five each, so we use the normal for two proportions distribution for this hypothesis test.

42.

1. 44.

The difference in mean speeds of the fastball pitches of the two pitchers

-2.46

47.

At the 1% significance level, we can reject the null hypothesis. There is sufficient data to conclude that the mean speed of Rodriguez's fastball is faster than Wesley's.

49.

Subscripts: 1 = Food, 2 = No Food

$$H_0 : \mu_1 \leq \mu_2$$

$$H_a : \mu_1 > \mu_2$$

51.

Subscripts: 1 = Gamma, 2 = Zeta

$$H_0 : \mu_1 = \mu_2$$

$$H_a : \mu_1 \neq \mu_2$$

53.

There is sufficient evidence so we cannot accept the null hypothesis. The data support that the melting point for Alloy Zeta is different from the melting point of Alloy Gamma.

54.

the mean difference of the system failures

56.

With a p -value 0.0067, we cannot accept the null hypothesis. There is enough evidence to support that the software patch is effective in reducing the number of system failures.

60.

$$H_0 : \mu_d \geq 0$$

$$H_a : \mu_d < 0$$

63.

We decline to reject the null hypothesis. There is not sufficient evidence to support that the medication is effective.

65.

Subscripts: 1: two-year colleges; 2: four-year colleges

1. 67.

Subscripts: 1: mechanical engineering; 2: electrical engineering

1. 69.

1. 71.

1. 74.

c

Test: two independent sample means, population standard deviations unknown. Random variable:

$$\bar{X}_1 - \bar{X}_2$$

Distribution: $H_0 : \mu_1 = \mu_2$ $H_a : \mu_1 < \mu_2$. The mean age of entering prostitution in Canada is lower than the mean age in the United States.

Graph: left-tailed p -value : 0.0151

Decision: Cannot reject H_0 .

Conclusion: At the 1% level of significance, from the sample data, there is not sufficient evidence to conclude that the mean age of entering prostitution in Canada is lower than the mean age in the United States.

78.

d

80.

1. 82.

Subscripts: 1 = Cabrillo College, 2 = Lake Tahoe College

1. 84.

a

Test: two independent sample proportions.

Random variable: $p'_1 - p'_2$

Distribution: $H_0 : p_1 = p_2$ $H_a : p_1 \neq p_2$. The proportion of eReader users is different for the 16- to 29-year-old users from that of the 30 and older users.

Graph: two-tailed

87.

Test: two independent sample proportions

Random variable: $p'_1 - p'_2$

Distribution: $H_0 : p_1 = p_2$ $H_a : p_1 > p_2$. A higher proportion of tablet owners are aged 16 to 29 years old than are 30 years old and older.

Graph: right-tailed

Do not reject the H_0 .

Conclusion: At the 1% level of significance, from the sample data, there is not sufficient evidence to conclude that a higher proportion of tablet owners are aged 16 to 29 years old than are 30 years old and older.

89.

Subscripts: 1: men; 2: women

1. 91.

1. 92.

1. 94.

Subscripts: 1 = boys, 2 = girls

1. 96.

Subscripts: 1 = non-hybrid sedans, 2 = hybrid sedans

1. 98.

1. 99.

p -value = 0.1494

103.

Test: two matched pairs or paired samples (t -test)

Random variable: \bar{X}_d

Distribution: t_{12}

$H_0 : \mu_d = 0$ $H_a : \mu_d > 0$

The mean of the differences of new female breast cancer cases in the south between 2013 and 2012 is greater than zero. The estimate for new female breast cancer cases in the south is higher in 2013 than in 2012.

Graph: right-tailed

p -value: 0.0004

Decision: Cannot accept H_0

Conclusion: At the 5% level of significance, from the sample data, there is sufficient evidence to conclude that there was a higher estimate of new female breast cancer cases in 2013 than in 2012.

105.

Test: matched or paired samples (t -test)

Difference

data:

$\{-0.9, -3.7, -3.2, -0.5, 0.6, -1.9, -0.5, 0.2, 0.6, 0.4, 1.7, -2.4, 1.8\}$

Random Variable: \bar{X}_d

Distribution: $H_0 : \mu_d = 0$ $H_a : \mu_d < 0$

The mean of the differences of the rate of underemployment in the northeastern states between 2012 and 2011 is less than zero. The underemployment rate went down from 2011 to 2012.

Graph: left-tailed.

Decision: Cannot reject H_0 .

Conclusion: At the 5% level of significance, from the sample data, there is not sufficient evidence to conclude that there was a decrease in the underemployment rates of the northeastern states from 2011 to 2012.

107.

e

109.

d

111.

f

113.

e

115.

f

117.

This page titled [10.15: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **10.14: Chapter Solution (Practice + Homework)** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

11: The Chi-Square Distribution

- 11.1: Prelude to the Chi-Square Distribution
- 11.2: Facts About the Chi-Square Distribution
- 11.3: Test of a Single Variance
- 11.4: Goodness-of-Fit Test
- 11.5: Test of Independence
- 11.6: Test for Homogeneity
- 11.7: Comparison of the Chi-Square Tests
- 11.8: Chapter Key Terms
- 11.9: Chapter Review
- 11.10: Chapter Formula Review
- 11.11: Chapter Practice
- 11.12: Homework
- 11.13: Chapter Homework
- 11.14: Chapter References
- 11.15: Chapter Solution (Practice + Homework)

This page titled [11: The Chi-Square Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

11.1: Prelude to the Chi-Square Distribution

Have you ever wondered if lottery winning numbers were evenly distributed or if some numbers occurred with a greater frequency? How about if the types of movies people preferred were different across different age groups? What about if a coffee machine was dispensing approximately the same amount of coffee each time? You could answer these questions by conducting a hypothesis test.



Figure 11.1.1: The chi-square distribution can be used to find relationships between two things, like grocery prices at different stores. (credit: Pete/flickr)

You will now study a new distribution, one that is used to determine the answers to such questions. This distribution is called the chi-square distribution. In this chapter, you will learn the three major applications of the chi-square distribution:

1. the goodness-of-fit test, which determines if data fit a particular distribution, such as in the lottery example
2. the test of independence, which determines if events are independent, such as in the movie example
3. the test of a single variance, which tests variability, such as in the coffee example

This page titled [11.1: Prelude to the Chi-Square Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.0: Prelude to the Chi-Square Distribution** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.2: Facts About the Chi-Square Distribution

The notation for the **chi-square distribution** is:

$$\chi \sim \chi_{df}^2$$

where df = degrees of freedom which depends on how chi-square is being used. (If you want to practice calculating chi-square probabilities then use $df = n - 1$. The degrees of freedom for the three major uses are each calculated differently.)

For the χ^2 distribution, the population mean is $\mu = df$ and the population standard deviation is $\sigma = \sqrt{2(df)}$.

The random variable is shown as χ^2 .

The random variable for a chi-square distribution with k degrees of freedom is the sum of k independent, squared standard normal variables.

$$\chi^2 = (Z_1)^2 + (Z_2)^2 + \dots + (Z_k)^2$$

1. The curve is non-symmetrical and skewed to the right.
2. There is a different chi-square curve for each df (11.2.1).
3. The test statistic for any test is always greater than or equal to zero.
4. When $df > 90$, the chi-square curve approximates the normal distribution. For $\chi \sim \chi_{1,000}^2$ the mean, $\mu = df = 1,000$ and the standard deviation, $\sigma = \sqrt{2(1,000)} = 44.7$. Therefore, $\chi \sim N(1,000, 44.7)$, approximately.
5. The mean, μ , is located just to the right of the peak.




Figure 11.2.1

This page titled [11.2: Facts About the Chi-Square Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [11.1: Facts About the Chi-Square Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.3: Test of a Single Variance

Thus far our interest has been exclusively on the population parameter μ or its counterpart in the binomial, p . Surely the mean of a population is the most critical piece of information to have, but in some cases we are interested in the variability of the outcomes of some distribution. In almost all production processes quality is measured not only by how closely the machine matches the target, but also the variability of the process. If one were filling bags with potato chips not only would there be interest in the average weight of the bag, but also how much variation there was in the weights. No one wants to be assured that the average weight is accurate when their bag has no chips. Electricity voltage may meet some average level, but great variability, spikes, can cause serious damage to electrical machines, especially computers. I would not only like to have a high mean grade in my classes, but also low variation about this mean. In short, statistical tests concerning the variance of a distribution have great value and many applications.

A **test of a single variance** assumes that the underlying distribution is **normal**. The null and alternative hypotheses are stated in terms of the **population variance**. The test statistic is:

$$\chi_c^2 = \frac{(n-1)s^2}{\sigma_0^2}$$

where:

- n = the total number of observations in the sample data
- s^2 = sample variance
- σ_0^2 = hypothesized value of the population variance
- $H_0 : \sigma^2 = \sigma_0^2$
- $H_a : \sigma^2 \neq \sigma_0^2$

You may think of s as the random variable in this test. The number of degrees of freedom is $df = n - 1$. A test of a single variance may be right-tailed, left-tailed, or two-tailed. [Example 11.3.1](#) will show you how to set up the null and alternative hypotheses. The null and alternative hypotheses contain statements about the population variance.

Example 11.3.1

Math instructors are not only interested in how their students do on exams, on average, but how the exam scores vary. To many instructors, the variance (or standard deviation) may be more important than the average.

Suppose a math instructor believes that the standard deviation for his final exam is five points. One of his best students thinks otherwise. The student claims that the standard deviation is more than five points. If the student were to conduct a hypothesis test, what would the null and alternative hypotheses be?

Answer

Even though we are given the population standard deviation, we can set up the test using the population variance as follows.

- $H_0 : \sigma^2 \leq 5^2$
- $H_a : \sigma^2 > 5^2$

Exercise 11.3.1

A SCUBA instructor wants to record the collective depths each of his students' dives during their checkout. He is interested in how the depths vary, even though everyone should have been at the same depth. He believes the standard deviation is three feet. His assistant thinks the standard deviation is less than three feet. If the instructor were to conduct a test, what would the null and alternative hypotheses be?

Example 11.3.2

With individual lines at its various windows, a post office finds that the standard deviation for waiting times for customers on Friday afternoon is 7.2 minutes. The post office experiments with a single, main waiting line and finds that for a random sample of 25 customers, the waiting times for customers have a standard deviation of 3.5 minutes on a Friday afternoon.

With a significance level of 5%, test the claim that **a single line causes lower variation among waiting times for customers**.

Answer

Since the claim is that a single line causes less variation, this is a test of a single variance. The parameter is the population variance, σ^2 .

Random Variable: The sample standard deviation, s , is the random variable. Let s = standard deviation for the waiting times.

- $H_0 : \sigma^2 \geq 7.2^2$
- $H_a : \sigma^2 < 7.2^2$

Distribution for the test: χ^2_{24} , where:

- n = the number of customers sampled
- $df = n - 1 = 25 - 1 = 24$

Calculate the test statistic:

$$\chi^2_c = \frac{(n-1)s^2}{\sigma^2} = \frac{(25-1)(3.5)^2}{7.2^2} = 5.67$$

where $n = 25$, $s = 3.5$, and $\sigma = 7.2$.




Figure 11.3.3

The graph of the Chi-square shows the distribution and marks the critical value with 24 degrees of freedom at 95% level of confidence, $\alpha = 0.05$, 13.85. The critical value of 13.85 came from the Chi squared table which is read very much like the students t table. The difference is that the students t distribution is symmetrical and the Chi squared distribution is not. At the top of the Chi squared table we see not only the familiar 0.05, 0.10, etc. but also 0.95, 0.975, etc. These are the columns used to find the left hand critical value. The graph also marks the calculated χ^2 test statistic of 5.67. Comparing the test statistic with the critical value, as we have done with all other hypothesis tests, we reach the conclusion.

The word "less" tells you this is a left-tailed test.

Make a decision: Because the calculated test statistic is in the tail we cannot accept H_0 . This means that you reject $\sigma^2 \geq 7.2^2$. In other words, you do not think the variation in waiting times is 7.2 minutes or more; you think the variation in waiting times is less.

Conclusion: At a 5% level of significance, from the data, there is sufficient evidence to conclude that a single line causes a lower variation among the waiting times **or** with a single line, the customer waiting times vary less than 7.2 minutes.

Example 11.3.3

Professor Hadley has a weakness for cream filled donuts, but he believes that some bakeries are not properly filling the donuts. A sample of 24 donuts reveals a mean amount of filling equal to 0.04 cups, and the sample standard deviation is 0.11 cups. Professor Hadley has an interest in the average quantity of filling, of course, but he is particularly distressed if one donut is radically different from another. Professor Hadley does not like surprises.

Test at 95% the null hypothesis that the population variance of donut filling is significantly different from the average amount of filling.

Answer

This is clearly a problem dealing with variances. In this case we are testing a single sample rather than comparing two samples from different populations. The null and alternative hypotheses are thus:

$$H_0 : \sigma^2 = 0.04$$

$$H_0 : \sigma^2 \neq 0.04$$

The test is set up as a two-tailed test because Professor Hadley has shown concern with too much variation in filling as well as too little: his dislike of a surprise is any level of filling outside the expected average of 0.04 cups. The test statistic is calculated to be:

$$\chi_c^2 = \frac{(n-1)s^2}{\sigma_o^2} = \frac{(24-1)0.11^2}{0.04^2} = 6.9575$$

The calculated χ^2 test statistic, 6.96, is in the tail therefore at a 0.05 level of significance, we cannot accept the null hypothesis that the variance in the donut filling is equal to 0.04 cups. It seems that Professor Hadley is destined to meet disappointment with each bit.




Figure 11.3.4

Exercise 11.3.3

The FCC conducts broadband speed tests to measure how much data per second passes between a consumer's computer and the internet. As of August of 2012, the standard deviation of Internet speeds across Internet Service Providers (ISPs) was 12.2 percent. Suppose a sample of 15 ISPs is taken, and the standard deviation is 13.2. An analyst claims that the standard deviation of speeds is more than what was reported. State the null and alternative hypotheses, compute the degrees of freedom, the test statistic, sketch the graph of the distribution and mark the area associated with the level of confidence, and draw a conclusion. Test at the 1% significance level.

This page titled [11.3: Test of a Single Variance](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [11.2: Test of a Single Variance](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.4: Goodness-of-Fit Test

In this type of hypothesis test, you determine whether the data "fit" a particular distribution or not. For example, you may suspect your unknown data fit a binomial distribution. You use a chi-square test (meaning the distribution for the hypothesis test is chi-square) to determine if there is a fit or not. **The null and the alternative hypotheses for this test may be written in sentences or may be stated as equations or inequalities.**

The test statistic for a goodness-of-fit test is:

$$\sum_k \frac{(O - E)^2}{E}$$

where:

- O = **observed values** (data)
- E = **expected values** (from theory)
- k = the number of different data cells or categories

The observed values are the data values and the expected values are the values you would expect to get if the null hypothesis were true. There are n terms of the form $\frac{(O-E)^2}{E}$.

The number of degrees of freedom is $df = (\text{number of categories} - 1)$.

The goodness-of-fit test is almost always right-tailed. If the observed values and the corresponding expected values are not close to each other, then the test statistic can get very large and will be way out in the right tail of the chi-square curve.

NOTE

The number of expected values inside each cell needs to be at least five in order to use this test.

Example 11.4.4

Absenteeism of college students from math classes is a major concern to math instructors because missing class appears to increase the drop rate. Suppose that a study was done to determine if the actual student absenteeism rate follows faculty perception. The faculty expected that a group of 100 students would miss class according to Table 11.4.1.

Table 11.4.1

Number of absences per term	Expected number of students
0–2	50
3–5	30
6–8	12
9–11	6
12+	2

A random survey across all mathematics courses was then done to determine the actual number (**observed**) of absences in a course. The chart in Table 11.4.2 displays the results of that survey.

Table 11.4.2

Number of absences per term	Actual number of students
0–2	35
3–5	40
6–8	20
9–11	1
12+	4

Determine the null and alternative hypotheses needed to conduct a goodness-of-fit test.

H_a : Student absenteeism **fits** faculty perception.

The alternative hypothesis is the opposite of the null hypothesis.

H_a : Student absenteeism **does not fit** faculty perception.

a. Can you use the information as it appears in the charts to conduct the goodness-of-fit test?

Answer

Solution 11.4

a. **No.** Notice that the expected number of absences for the "12+" entry is less than five (it is two). Combine that group with the "9–11" group to create new tables where the number of students for each entry are at least five. The new results are in Table 11.4.3 and Table 11.4.4

Number of absences per term	Expected number of students
0–2	50
3–5	30
6–8	12
9+	8

Table 11.3

Table 11.4.4

Number of absences per term	Actual number of students
0–2	35
3–5	40
6–8	20
9+	5

b. What is the number of degrees of freedom (df)?

Answer

Solution 11.4

b. There are four "cells" or categories in each of the new tables.

$$df = \text{number of cells} - 1 = 4 - 1 = 3$$

Exercise 11.4.4

A factory manager needs to understand how many products are defective versus how many are produced. The number of expected defects is listed in Table 11.4.5

Table 11.4.5

Number produced	Number defective
0–100	5
101–200	6
201–300	7
301–400	8
401–500	10

A random sample was taken to determine the actual number of defects. Table 11.4.6 shows the results of the survey.

Table 11.4.6

Number produced	Number defective
0–100	5
101–200	7
201–300	8
301–400	9
401–500	11

State the null and alternative hypotheses needed to conduct a goodness-of-fit test, and state the degrees of freedom.

Example 11.4.5

Employers want to know which days of the week employees are absent in a five-day work week. Most employers would like to believe that employees are absent equally during the week. Suppose a random sample of 60 managers were asked on which day of the week they had the highest number of employee absences. The results were distributed as in Table 11.4.7. For the population of employees, do the days for the highest number of absences occur with equal frequencies during a five-day work week? Test at a 5% significance level.

Table 11.4.7 Day of the Week Employees were Most Absent

	Monday	Tuesday	Wednesday	Thursday	Friday
Number of absences	15	12	9	9	15

Answer

Solution 11.5

The null and alternative hypotheses are:

- H_0 : The absent days occur with equal frequencies, that is, they fit a uniform distribution.
- H_a : The absent days occur with unequal frequencies, that is, they do not fit a uniform distribution.

If the absent days occur with equal frequencies, then, out of 60 absent days (the total in the sample: $15 + 12 + 9 + 9 + 15 = 60$), there would be 12 absences on Monday, 12 on Tuesday, 12 on Wednesday, 12 on Thursday, and 12 on Friday. These numbers are the **expected** (E) values. The values in the table are the **observed** (O) values or data.

This time, calculate the χ^2 test statistic by hand. Make a chart with the following headings and fill in the columns:

- Expected (E) values (12, 12, 12, 12, 12)
- Observed (O) values (15, 12, 9, 9, 15)
- $(O - E)$
- $(O - E)^2$
- $\frac{(O - E)^2}{E}$

Now add (sum) the last column. The sum is three. This is the χ^2 test statistic.

The calculated test statistics is 3 and the critical value of the χ^2 distribution at 4 degrees of freedom the 0.05 level of confidence is 9.48. This value is found in the χ^2 table at the 0.05 column on the degrees of freedom row 4.

The degrees of freedom are the number of cells $- 1 = 5 - 1 = 4$

Next, complete a graph like the following one with the proper labeling and shading. (You should shade the right tail.)




Figure 11.4.5

$$\chi_c^2 = \sum_k \frac{(O - E)^2}{E} = 3$$

The decision is not to reject the null hypothesis because the calculated value of the test statistic is not in the tail of the distribution.

Conclusion: At a 5% level of significance, from the sample data, there is not sufficient evidence to conclude that the absent days do not occur with equal frequencies.

Exercise 11.4.5

Teachers want to know which night each week their students are doing most of their homework. Most teachers think that students do homework equally throughout the week. Suppose a random sample of 56 students were asked on which night of the week they did the most homework. The results were distributed as in Table 11.4.8

Table 11.4.8

	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Number of students	11	8	10	7	10	5	5

From the population of students, do the nights for the highest number of students doing the majority of their homework occur with equal frequencies during a week? What type of hypothesis test should you use?

Example 11.4.6

One study indicates that the number of televisions that American families have is distributed (this is the **given** distribution for the American population) as in Table 11.4.9

Table 11.4.9

Number of Televisions	Percent
0	10
1	16
2	55
3	11
4+	8

The table contains expected (E) percents.

A random sample of 600 families in the far western United States resulted in the data in Table 11.4.10

Table 11.4.10

Number of Televisions	Frequency
	Total = 600

Number of Televisions	Frequency
0	66
1	119
2	340
3	60
4+	15
	Total = 600

The table contains observed (O) frequency values.

At the 1% significance level, does it appear that the distribution "number of televisions" of far western United States families is different from the distribution for the American population as a whole?

Answer

Solution 11.6

This problem asks you to test whether the far western United States families distribution fits the distribution of the American families. This test is always right-tailed.

The first table contains expected percentages. To get expected (E) frequencies, multiply the percentage by 600. The expected frequencies are shown in Table 11.4.11

Table 11.4.11

Number of televisions	Percent	Expected frequency
0	10	$(0.10)(600) = 60$
1	16	$(0.16)(600) = 96$
2	55	$(0.55)(600) = 330$
3	11	$(0.11)(600) = 66$
over 3	8	$(0.08)(600) = 48$

Therefore, the expected frequencies are 60, 96, 330, 66, and 48.

H_0 : The "number of televisions" distribution of far western United States families is the same as the "number of televisions" distribution of the American population.

H_a : The "number of televisions" distribution of far western United States families is different from the "number of televisions" distribution of the American population.

Distribution for the test: χ^2_4 where $df = (\text{the number of cells}) - 1 = 5 - 1 = 4$.

Calculate the test statistic: $\chi^2 = 29.65$

Graph:




Figure 11.4.6

The graph of the Chi-square shows the distribution and marks the critical value with four degrees of freedom at 99% level of confidence, $\alpha = .01$, 13.277. The graph also marks the calculated chi squared test statistic of 29.65. Comparing the test statistic with the critical value, as we have done with all other hypothesis tests, we reach the conclusion.

Make a decision: Because the test statistic is in the tail of the distribution we cannot accept the null hypothesis.

This means you reject the belief that the distribution for the far western states is the same as that of the American population as a whole.

Conclusion: At the 1% significance level, from the data, there is sufficient evidence to conclude that the "number of televisions" distribution for the far western United States is different from the "number of televisions" distribution for the American population as a whole.

Exercise 11.4.6

The expected percentage of the number of pets students have in their homes is distributed (this is the given distribution for the student population of the United States) as in Table 11.4.12

Table 11.4.12

Number of pets	Percent
0	18
1	25
2	30
3	18
4+	9

A random sample of 1,000 students from the Eastern United States resulted in the data in Table 11.4.13

Table 11.4.13

Number of pets	Frequency
0	210
1	240
2	320
3	140
4+	90

At the 1% significance level, does it appear that the distribution "number of pets" of students in the Eastern United States is different from the distribution for the United States student population as a whole?

Example 11.4.7

Suppose you flip two coins 100 times. The results are 20HH, 27HT, 30TH, and 23TT. Are the coins fair? Test at a 5% significance level.

Answer

Solution 11.7

This problem can be set up as a goodness-of-fit problem. The sample space for flipping two fair coins is $\{HH, HT, TH, TT\}$. Out of 100 flips, you would expect 25 HH, 25HT, 25TH and 25TT. This is the expected distribution from the binomial probability distribution. The question, "Are the coins fair?" is the same as saying, "Does the distribution of the coins (20HH, 27HT, 30TH, 23TT) fit the expected distribution?"

Random Variable: Let X = the number of heads in one flip of the two coins. X takes on the values 0, 1, 2. (There are 0, 1, or 2 heads in the flip of two coins.) Therefore, the **number of cells is three**. Since X = the number of heads, the observed frequencies are 20 (for two heads), 57 (for one head), and 23 (for zero heads or both tails). The expected frequencies are 25 (for two heads), 50 (for one head), and 25 (for zero heads or both tails). This test is right-tailed.

H_0 : The coins are fair.

H_a : The coins are not fair.

Distribution for the test: χ^2_2 where $df = 3 - 1 = 2$.

Calculate the test statistic: $\chi^2 = 2.14$.

Graph:




Figure 11.4.7

The graph of the Chi-square shows the distribution and marks the critical value with two degrees of freedom at 95% level of confidence, $\alpha = 0.05$, 5.991. The graph also marks the calculated χ^2 test statistic of 2.14. Comparing the test statistic with the critical value, as we have done with all other hypothesis tests, we reach the conclusion.

Conclusion: There is insufficient evidence to conclude that the coins are not fair: we cannot reject the null hypothesis that the coins are fair.

This page titled [11.4: Goodness-of-Fit Test](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [11.3: Goodness-of-Fit Test](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.5: Test of Independence

Tests of independence involve using a **contingency table** of observed (data) values. The test statistic for a **test of independence** is similar to that of a goodness-of-fit test:

$$\sum_{(i,j)} \frac{(O - E)^2}{E}$$

where:

- O = observed values
- E = expected values
- i = the number of rows in the table
- j = the number of columns in the table

There are $i \cdot j$ terms of the form $\frac{(O-E)^2}{E}$.

A **test of independence determines whether two factors are independent or not**. You first encountered the term independence in [Table 3.1](#) earlier. As a review, consider the following example.

Note

The expected value inside each cell needs to be at least five in order for you to use this test.

Example 11.8

Suppose A = a speeding violation in the last year and B = a cell phone user while driving. If A and B are independent then $P(A \cap B) = P(A)P(B)$. $A \cap B$ is the event that a driver received a speeding violation last year and also used a cell phone while driving. Suppose, in a study of drivers who received speeding violations in the last year, and who used cell phone while driving, that 755 people were surveyed. Out of the 755, 70 had a speeding violation and 685 did not; 305 used cell phones while driving and 450 did not.

Let y = expected number of drivers who used a cell phone while driving and received speeding violations.

If A and B are independent, then $P(A \cap B) = P(A)P(B)$. By substitution,

$$\frac{y}{755} = \left(\frac{70}{755} \right) \left(\frac{305}{755} \right)$$

Solve for y : $y = \frac{(70)(305)}{755} = 28.3$

About 28 people from the sample are expected to use cell phones while driving and to receive speeding violations.

In a test of independence, we state the null and alternative hypotheses in words. Since the contingency table consists of **two factors**, the null hypothesis states that the factors are **independent** and the alternative hypothesis states that they are **not independent (dependent)**. If we do a test of independence using the example, then the null hypothesis is:

H_0 : Being a cell phone user while driving and receiving a speeding violation are independent events; in other words, they have no effect on each other.

If the null hypothesis were true, we would expect about 28 people to use cell phones while driving and to receive a speeding violation.

The test of independence is always right-tailed because of the calculation of the test statistic. If the expected and observed values are not close together, then the test statistic is very large and way out in the right tail of the chi-square curve, as it is in a goodness-of-fit.

The number of degrees of freedom for the test of independence is:

$$df = (\text{number of columns} - 1)(\text{number of rows} - 1)$$

The following formula calculates the **expected number (E)**:

$$E = \frac{(\text{row total})(\text{column total})}{\text{total number surveyed}}$$

Exercise 11.8

A sample of 300 students is taken. Of the students surveyed, 50 were music students, while 250 were not. Ninety-seven of the 300 surveyed were on the honor roll, while 203 were not. If we assume being a music student and being on the honor roll are independent events, what is the expected number of music students who are also on the honor roll?

Example 11.9

A volunteer group, provides from one to nine hours each week with disabled senior citizens. The program recruits among community college students, four-year college students, and nonstudents. In Table 11.14 is a **sample** of the adult volunteers and the number of hours they volunteer per week.

The table contains **observed (O)** values (data).
Table 11.14 Number of Hours Worked Per Week by Volunteer Type (Observed)

Type of volunteer	1–3 Hours	4–6 Hours	7–9 Hours	Row total
Community college students	111	96	48	255
Four-year college students	96	133	61	290
Nonstudents	91	150	53	294
Column total	298	379	162	839

Is the number of hours volunteered **independent** of the type of volunteer?

Answer

Solution 11.9

The **observed table** and the question at the end of the problem, "Is the number of hours volunteered independent of the type of volunteer?" tell you this is a test of independence. The two factors are **number of hours volunteered** and **type of volunteer**. This test is always right-tailed.

H_0 : The number of hours volunteered is **independent** of the type of volunteer.

H_a : The number of hours volunteered is **dependent** on the type of volunteer.

The expected result are in Table 11.15.

The table contains **expected (E)** values (data).
Table 11.15 Number of Hours Worked Per Week by Volunteer Type (Expected)

Type of volunteer	1-3 Hours	4-6 Hours	7-9 Hours
Community college students	90.57	115.19	49.24
Four-year college students	103.00	131.00	56.00
Nonstudents	104.42	132.81	56.77

For example, the calculation for the expected frequency for the top left cell is

$$E = \frac{(\text{row total})(\text{column total})}{\text{total number surveyed}} = \frac{(255)(298)}{839} = 90.57$$

Calculate the test statistic: $\chi^2 = 12.99$ (calculator or computer)

Distribution for the test: χ_4^2

$$df = (3 \text{ columns} - 1)(3 \text{ rows} - 1) = (2)(2) = 4$$

Graph:




Figure 11.8

The graph of the Chi-square shows the distribution and marks the critical value with four degrees of freedom at 95% level of confidence, $\alpha = 0.05$, 9.488. The graph also marks the calculated χ^2_c test statistic of 12.99. Comparing the test statistic with the critical value, as we have done with all other hypothesis tests, we reach the conclusion.

Make a decision: Because the calculated test statistic is in the tail we cannot accept H_0 . This means that the factors are not independent.

Conclusion: At a 5% level of significance, from the data, there is sufficient evidence to conclude that the number of hours volunteered and the type of volunteer are dependent on one another.

For the example in Table 11.15, if there had been another type of volunteer, teenagers, what would the degrees of freedom be?

Exercise 11.9

The Bureau of Labor Statistics gathers data about employment in the United States. A sample is taken to calculate the number of U.S. citizens working in one of several industry sectors over time. Table 11.16 shows the results:

Table 11.16

Industry sector	2000	2010	2020	Total
Nonagriculture wage and salary	13,243	13,044	15,018	41,305
Goods-producing, excluding agriculture	2,457	1,771	1,950	6,178
Services-providing	10,786	11,273	13,068	35,127
Agriculture, forestry, fishing, and hunting	240	214	201	655
Nonagriculture self-employed and unpaid family worker	931	894	972	2,797
Secondary wage and salary jobs in agriculture and private household industries	14	11	11	36
Secondary jobs as a self-employed or unpaid family worker	196	144	152	492
Total	27,867	27,351	31,372	86,590

We want to know if the change in the number of jobs is independent of the change in years. State the null and alternative hypotheses and the degrees of freedom.

Example 11.10

De Anza College is interested in the relationship between anxiety level and the need to succeed in school. A random sample of 400 students took a test that measured anxiety level and need to succeed in school. Table 11.17 shows the results. De Anza College wants to know if anxiety level and need to succeed in school are independent events.

Table 11.17 Need to Succeed in School vs. Anxiety Level

Need to succeed in school	High anxiety	Med-high anxiety	Medium anxiety	Med-low anxiety	Low anxiety	Row total
High need	35	42	53	15	10	155
Medium need	18	48	63	33	31	193
Low need	4	5	11	15	17	52
Column total	57	95	127	63	58	400

a. How many high anxiety level students are expected to have a high need to succeed in school?

Answer

Solution 11.10

a. The column total for a high anxiety level is 57. The row total for high need to succeed in school is 155. The sample size or total surveyed is 400.

$$E = \frac{(\text{row total})(\text{column total})}{\text{total surveyed}} = \frac{155 \cdot 57}{400} = 22.09$$

The expected number of students who have a high anxiety level and a high need to succeed in school is about 22.

b. If the two variables are independent, how many students do you expect to have a low need to succeed in school and a med-low level of anxiety?

Answer

Solution 11.10

b. The column total for a med-low anxiety level is 63. The row total for a low need to succeed in school is 52. The sample size or total surveyed is 400.

c. $E = \frac{(\text{row total})(\text{column total})}{\text{total surveyed}} = \underline{\hspace{2cm}}$

Answer

Solution 11.10

c. $E = \frac{(\text{row total})(\text{column total})}{\text{total surveyed}} = 8.19$

d. The expected number of students who have a med-low anxiety level and a low need to succeed in school is about _____.

Answer

Solution 11.10

d. 8

- **11.4: Test of Independence** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.6: Test for Homogeneity

The goodness-of-fit test can be used to decide whether a population fits a given distribution, but it will not suffice to decide whether two populations follow the same unknown distribution. A different test, called the **test for homogeneity**, can be used to draw a conclusion about whether two populations have the same distribution. To calculate the test statistic for a test for homogeneity, follow the same procedure as with the test of independence.

NOTE

The expected value inside each cell needs to be at least five in order for you to use this test.

Hypotheses

- H_0 : The distributions of the two populations are the same.
- H_a : The distributions of the two populations are not the same.

Test Statistic

Use a χ^2 test statistic. It is computed in the same way as the test for independence.

Degrees of Freedom (df)

$$df = \text{number of columns} - 1$$

Requirements

All values in the table must be greater than or equal to five.

Common Uses

Comparing two populations. For example: men vs. women, before vs. after, east vs. west. The variable is categorical with more than two possible response values.

Example 11.6.1

Do male and female college students have the same distribution of living arrangements? Use a level of significance of 0.05. Suppose that 250 randomly selected male college students and 300 randomly selected female college students were asked about their living arrangements: dormitory, apartment, with parents, other. The results are shown in Table 11.6.18 Do male and female college students have the same distribution of living arrangements?

Table 11.6.18 Distribution of living arrangements for college males and college females

	Dormitory	Apartment	With Parents	Other
Males	72	84	49	45
Females	91	86	88	35

Answer

Solution 11.11

H_0 : The distribution of living arrangements for male college students is the same as the distribution of living arrangements for female college students.

H_a : The distribution of living arrangements for male college students is not the same as the distribution of living arrangements for female college students.

Degrees of Freedom (df):

$$df = \text{number of columns} - 1 = 4 - 1 = 3$$

Distribution for the test: χ^2_3

Calculate the test statistic: $\chi^2_c = 10.129$




Figure 11.6.9

The graph of the Chi-square shows the distribution and marks the critical value with three degrees of freedom at 95% level of confidence, $\alpha = 0.05$, 7.815. The graph also marks the calculated χ^2 test statistic of 10.129. Comparing the test statistic with the critical value, as we have done with all other hypothesis tests, we reach the conclusion.

Make a decision: Because the calculated test statistic is in the tail we cannot accept H_0 . This means that the distributions are not the same.

Conclusion: At a 5% level of significance, from the data, there is sufficient evidence to conclude that the distributions of living arrangements for male and female college students are not the same.

Notice that the conclusion is only that the distributions are not the same. We cannot use the test for homogeneity to draw any conclusions about how they differ.

Exercise 11.6.1A

Do families and singles have the same distribution of cars? Use a level of significance of 0.05. Suppose that 100 randomly selected families and 200 randomly selected singles were asked what type of car they drove: sport, sedan, hatchback, truck, van/SUV. The results are shown in Table 11.6.19. Do families and singles have the same distribution of cars? Test at a level of significance of 0.05.

Table 11.6.19

	Sport	Sedan	Hatchback	Truck	Van/SUV
Family	5	15	35	17	28
Single	45	65	37	46	7

Exercise 11.6.1B

Ivy League schools receive many applications, but only some can be accepted. At the schools listed in Table 11.6.20 two types of applications are accepted: regular and early decision.

Table 11.6.20

Application type accepted	Brown	Columbia	Cornell	Dartmouth	Penn	Yale
Regular	2,115	1,792	5,306	1,734	2,685	1,245
Early decision	577	627	1,228	444	1,195	761

We want to know if the number of regular applications accepted follows the same distribution as the number of early applications accepted. State the null and alternative hypotheses, the degrees of freedom and the test statistic, sketch the graph of the χ^2 distribution and show the critical value and the calculated value of the test statistic, and draw a conclusion about the test of homogeneity.

This page titled [11.6: Test for Homogeneity](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [11.5: Test for Homogeneity](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.7: Comparison of the Chi-Square Tests

Above the χ^2 test statistic was used in three different circumstances. The following list is a summary of which χ^2 test is the appropriate one to use in different circumstances.

Test for Goodness-of-Fit

Use the goodness-of-fit test to decide whether a population with an unknown distribution "fits" a known distribution. In this case there will be a single qualitative survey question or a single outcome of an experiment from a single population. Goodness-of-Fit is typically used to see if the population is uniform (all outcomes occur with equal frequency), the population is normal, or the population is the same as another population with a known distribution. The null and alternative hypotheses are:

- H_0 : The population fits the given distribution.
- H_a : The population does not fit the given distribution.

Test for Independence

Use the test for independence to decide whether two variables (factors) are independent or dependent. In this case there will be two qualitative survey questions or experiments and a contingency table will be constructed. The goal is to see if the two variables are unrelated (independent) or related (dependent). The null and alternative hypotheses are:

- H_0 : The two variables (factors) are independent.
- H_a : The two variables (factors) are dependent.

Test for Homogeneity

Use the test for homogeneity to decide if two populations with unknown distributions have the same distribution as each other. In this case there will be a single qualitative survey question or experiment given to two different populations. The null and alternative hypotheses are:

- H_0 : The two populations follow the same distribution.
- H_a : The two populations have different distributions.

This page titled [11.7: Comparison of the Chi-Square Tests](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.6: Comparison of the Chi-Square Tests** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.8: Chapter Key Terms

Contingency Table

a table that displays sample values for two different factors that may be dependent or contingent on one another; it facilitates determining conditional probabilities.

Goodness-of-Fit

a hypothesis test that compares expected and observed values in order to look for significant differences within one non-parametric variable. The degrees of freedom used equals the (number of categories – 1).

Test for Homogeneity

a test used to draw a conclusion about whether two populations have the same distribution. The degrees of freedom used equals the (number of columns – 1).

Test of Independence

a hypothesis test that compares expected and observed values for contingency tables in order to test for independence between two variables. The degrees of freedom used equals the (number of columns – 1) multiplied by the (number of rows – 1).

This page titled [11.8: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

11.9: Chapter Review

11.1 Facts About the Chi-Square Distribution

The chi-square distribution is a useful tool for assessment in a series of problem categories. These problem categories include primarily (i) whether a data set fits a particular distribution, (ii) whether the distributions of two populations are the same, (iii) whether two events might be independent, and (iv) whether there is a different variability than expected within a population.

An important parameter in a chi-square distribution is the degrees of freedom df in a given problem. The random variable in the chi-square distribution is the sum of squares of df standard normal variables, which must be independent. The key characteristics of the chi-square distribution also depend directly on the degrees of freedom.

The chi-square distribution curve is skewed to the right, and its shape depends on the degrees of freedom df . For $df > 90$, the curve approximates the normal distribution. Test statistics based on the chi-square distribution are always greater than or equal to zero. Such application tests are almost always right-tailed tests.

11.2 Test of a Single Variance

To test variability, use the chi-square test of a single variance. The test may be left-, right-, or two-tailed, and its hypotheses are always expressed in terms of the variance (or standard deviation).

11.3 Goodness-of-Fit Test

To assess whether a data set fits a specific distribution, you can apply the goodness-of-fit hypothesis test that uses the chi-square distribution. The null hypothesis for this test states that the data come from the assumed distribution. The test compares observed values against the values you would expect to have if your data followed the assumed distribution. The test is almost always right-tailed. Each observation or cell category must have an expected value of at least five.

11.4 Test of Independence

To assess whether two factors are independent or not, you can apply the test of independence that uses the chi-square distribution. The null hypothesis for this test states that the two factors are independent. The test compares observed values to expected values. The test is right-tailed. Each observation or cell category must have an expected value of at least 5.

11.5 Test for Homogeneity

To assess whether two data sets are derived from the same distribution—which need not be known, you can apply the test for homogeneity that uses the chi-square distribution. The null hypothesis for this test states that the populations of the two data sets come from the same distribution. The test compares the observed values against the expected values if the two populations followed the same distribution. The test is right-tailed. Each observation or cell category must have an expected value of at least five.

11.6 Comparison of the Chi-Square Tests

The goodness-of-fit test is typically used to determine if data fits a particular distribution. The test of independence makes use of a contingency table to determine the independence of two factors. The test for homogeneity determines whether two populations come from the same distribution, even if this distribution is unknown.

This page titled [11.9: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.13: Chapter Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.10: Chapter Formula Review

Facts About the Chi-Square Distribution

$\chi^2 = (Z_1)^2 + (Z_2)^2 + \dots + (Z_{df})^2$ chi-square distribution random variable

$\mu_{\chi^2}^2 = df$ chi-square distribution population mean

$\sigma_{\chi^2} = \sqrt{2(df)}$ Chi-Square distribution population standard deviation

Test of a Single Variance

$\chi^2 = \frac{(n-1)s^2}{\sigma_0^2}$ Test of a single variance statistic where:

n : sample size

s : sample standard deviation

σ_0 : hypothesized value of the population standard deviation

$df = n - 1$ Degrees of freedom

Test of a Single Variance

- Use the test to determine variation.
- The degrees of freedom is the number of samples - 1.
- The test statistic is $\frac{(n-1)s^2}{\sigma_0^2}$, where n = sample size, s^2 = sample variance, and σ^2 = population variance.
- The test may be left-, right-, or two-tailed.

Goodness-of-Fit Test

$\sum_k \frac{(O-E)^2}{E}$ goodness-of-fit test statistic where:

O : observed values

E : expected values

k : number of different data cells or categories

$df = k - 1$ degrees of freedom

Test of Independence

Test of Independence

- The number of degrees of freedom is equal to (number of columns - 1)(number of rows - 1).
- The test statistic is $\sum_{i,j} \frac{(O-E)^2}{E}$ where O = observed values, E = expected values, i = the number of rows in the table, and j = the number of columns in the table.
- If the null hypothesis is true, the expected number $E = \frac{(\text{row total})(\text{column total})}{\text{total surveyed}}$.

Test for Homogeneity

$\sum_{i,j} \frac{(O-E)^2}{E}$ Homogeneity test statistic where: O = observed values

E = expected values

i = number of rows in data contingency table

j = number of columns in data contingency table

$df = (i - 1)(j - 1)$ Degrees of freedom

This page titled [11.10: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.8: Chapter Formula Review** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.11: Chapter Practice

11.1 Facts About the Chi-Square Distribution

1.



Figure 11.11.10

11.2 Test of a Single Variance

Use the following information to answer the next three exercises: An archer's standard deviation for his hits is six (data is measured in distance from the center of the target). An observer claims the standard deviation is less.

6.

What type of test should be used?

7.

State the null and alternative hypotheses.

8.

Is this a right-tailed, left-tailed, or two-tailed test?

Figure 11.11.11

Let $\alpha = 0.05$

Decision: _____

Reason for the Decision: _____

Conclusion (write out in complete sentences): _____

29.

Does it appear that the pattern of AIDS cases in Santa Clara County corresponds to the distribution of ethnic groups in this county? Why or why not?

11.4 Test of Independence

Determine the appropriate test to be used in the next three exercises.

30.

A pharmaceutical company is interested in the relationship between age and presentation of symptoms for a common viral infection. A random sample is taken of 500 people with the infection across different age groups.

31.

The owner of a baseball team is interested in the relationship between player salaries and team winning percentage. He takes a random sample of 100 players from different organizations.

32.

A marathon runner is interested in the relationship between the brand of shoes runners wear and their run times. She takes a random sample of 50 runners and records their run times as well as the brand of shoes they were wearing.

Use the following information to answer the next seven exercises: Transit Railroads is interested in the relationship between travel distance and the ticket class purchased. A random sample of 200 passengers is taken. Table 11.11.25 shows the results. The railroad wants to know if a passenger's choice in ticket class is independent of the distance they must travel.

Table 11.11.25

Traveling distance	Third class	Second class	First class	Total
1–100 miles	21	14	6	41
101–200 miles	18	16	8	42
201–300 miles	16	17	15	48
301–400 miles	12	14	21	47
401–500 miles	6	6	10	22
Total	73	67	60	200

33.

State the hypotheses.

H_0 : _____

H_a : _____

34.

$df =$ _____

35.

How many passengers are expected to travel between 201 and 300 miles and purchase second-class tickets?

36.

How many passengers are expected to travel between 401 and 500 miles and purchase first-class tickets?

37.

What is the test statistic?

38.

What can you conclude at the 5% level of significance?

Use the following information to answer the next eight exercises: An article in the New England Journal of Medicine, discussed a study on smokers in California and Hawaii. In one part of the report, the self-reported ethnicity and smoking levels per day were given. Of the people smoking at most ten cigarettes per day, there were 9,886 African Americans, 2,745 Native Hawaiians, 12,831 Latinos, 8,378 Japanese Americans and 7,650 whites. Of the people smoking 11 to 20 cigarettes per day, there were 6,514 African Americans, 3,062 Native Hawaiians, 4,932 Latinos, 10,680 Japanese Americans, and 9,877 whites. Of the people smoking 21 to 30 cigarettes per day, there were 1,671 African Americans, 1,419 Native Hawaiians, 1,406 Latinos, 4,715 Japanese Americans, and 6,062 whites. Of the people smoking at least 31 cigarettes per day, there were 759 African Americans, 788 Native Hawaiians, 800 Latinos, 2,305 Japanese Americans, and 3,970 whites.

39.

Complete the table.

Table 11.11.26 Smoking Levels by Ethnicity (Observed)

Smoking level per day	African American	Native Hawaiian	Latino	Japanese Americans	White	Totals
1-10						
11-20						
21-30						
31+						

Smoking level per day	African American	Native Hawaiian	Latino	Japanese Americans	White	Totals
Totals						

40.

State the hypotheses.

H_0 : _____

H_a : _____

41.

Enter expected values in Table 11.11.26 Round to two decimal places.

Calculate the following values:

42.

df = _____

43.

χ^2 test statistic = _____

44.

Is this a right-tailed, left-tailed, or two-tailed test? Explain why.

45.

Graph the situation. Label and scale the horizontal axis. Mark the mean and test statistic. Shade in the region corresponding to the confidence level.



Figure 11.11.12

State the decision and conclusion (in a complete sentence) for the following preconceived levels of α .

46.

$\alpha = 0.05$

1. Decision: _____

2. Reason for the decision: _____

3. Conclusion (write out in a complete sentence): _____

47.

$\alpha = 0.01$

1. Decision: _____

2. Reason for the decision: _____

3. Conclusion (write out in a complete sentence): _____

11.5 Test for Homogeneity

48.

A math teacher wants to see if two of her classes have the same distribution of test scores. What test should she use?

49.

What are the null and alternative hypotheses for Table 11.11.27

Table 11.11.27

	20–30	30–40	40–50	50–60
Private practice	16	40	38	6
Hospital	8	44	59	39

53.

State the null and alternative hypotheses.

54.

$df =$ _____

55.

What is the test statistic?

56.

What can you conclude at the 5% significance level?

11.6 Comparison of the Chi-Square Tests

57.

Which test do you use to decide whether an observed distribution is the same as an expected distribution?

58.

What is the null hypothesis for the type of test from [Exercise 11.11.57](#)?

59.

Which test would you use to decide whether two factors have a relationship?

60.

Which test would you use to decide if two populations have the same distribution?

61.

How are tests of independence similar to tests for homogeneity?

62.

How are tests of independence different from tests for homogeneity?

This page titled [11.11: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [11.11: Chapter Practice](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.12: Homework

122.

1. Explain why a goodness-of-fit test and a test of independence are generally right-tailed tests.
 2. If you did a left-tailed test, what would you be testing?
-

This page titled [11.12: Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [11.7: Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.13: Chapter Homework

Facts About the Chi-Square Distribution

Decide whether the following statements are true or false.

63.

As the number of degrees of freedom increases, the graph of the chi-square distribution looks more and more symmetrical.

64.

The standard deviation of the chi-square distribution is twice the mean.

65.

The mean and the median of the chi-square distribution are the same if $df = 24$.

Test of a Single Variance

Use the following information to answer the next twelve exercises: Suppose an airline claims that its flights are consistently on time with an average delay of at most 15 minutes. It claims that the average delay is so consistent that the variance is no more than 150 minutes. Doubting the consistency part of the claim, a disgruntled traveler calculates the delays for his next 25 flights. The average delay for those 25 flights is 22 minutes with a standard deviation of 15 minutes.

66.

Is the traveler disputing the claim about the average or about the variance?

67.

A sample standard deviation of 15 minutes is the same as a sample variance of _____ minutes.

68.

Is this a right-tailed, left-tailed, or two-tailed test?

69.

H_0 : _____

70.

$df =$ _____

71.

chi-square test statistic = _____

72.

Graph the situation. Label and scale the horizontal axis. Mark the mean and test statistic. Shade the area associated with the level of confidence.

73.

Let $\alpha = 0.05$

Decision: _____

Conclusion (write out in a complete sentence.): _____

74.

How did you know to test the variance instead of the mean?

75.

If an additional test were done on the claim of the average delay, which distribution would you use?

76.

If an additional test were done on the claim of the average delay, but 45 flights were surveyed, which distribution would you use?

77.

A plant manager is concerned her equipment may need recalibrating. It seems that the actual weight of the 15 oz. cereal boxes it fills has been fluctuating. The standard deviation should be at most 0.5 oz. In order to determine if the machine needs to be recalibrated, 84 randomly selected boxes of cereal from the next day's production were weighed. The standard deviation of the 84 boxes was 0.54. Does the machine need to be recalibrated?

78.

Consumers may be interested in whether the cost of a particular calculator varies from store to store. Based on surveying 43 stores, which yielded a sample mean of \$84 and a sample standard deviation of \$12, test the claim that the standard deviation is greater than \$15.

79.

Isabella, an accomplished **Bay to Breakers** runner, claims that the standard deviation for her time to run the 7.5 mile race is at most three minutes. To test her claim, Rupinder looks up five of her race times. They are 55 minutes, 61 minutes, 58 minutes, 63 minutes, and 57 minutes.

80.

Airline companies are interested in the consistency of the number of babies on each flight, so that they have adequate safety equipment. They are also interested in the variation of the number of babies. Suppose that an airline executive believes the average number of babies on flights is six with a variance of nine at most. The airline conducts a survey. The results of the 18 flights surveyed give a sample average of 6.4 with a sample standard deviation of 3.9. Conduct a hypothesis test of the airline executive's belief.

81.

The number of births per woman in China is 1.6 down from 5.91 in 1966. This fertility rate has been attributed to the law passed in 1979 restricting births to one per woman. Suppose that a group of students studied whether or not the standard deviation of births per woman was greater than 0.75. They asked 50 women across China the number of births they had had. The results are shown in Table 11.13.28 Does the students' survey indicate that the standard deviation is greater than 0.75?

Table 11.13.28

# of births	Frequency
0	5
1	30
2	10
3	5

82.

According to an avid aquarist, the average number of fish in a 20-gallon tank is 10, with a standard deviation of two. His friend, also an aquarist, does not believe that the standard deviation is two. She counts the number of fish in 15 other 20-gallon tanks. Based on the results that follow, do you think that the standard deviation is different from two? Data: 11; 10; 9; 10; 10; 11; 11; 10; 12; 9; 7; 9; 11; 10; 11

83.

The manager of "Frenchies" is concerned that patrons are not consistently receiving the same amount of French fries with each order. The chef claims that the standard deviation for a ten-ounce order of fries is at most 1.5 oz., but the manager thinks that it may be higher. He randomly weighs 49 orders of fries, which yields a mean of 11 oz. and a standard deviation of two oz.

84.

You want to buy a specific computer. A sales representative of the manufacturer claims that retail stores sell this computer at an average price of \$1,249 with a very narrow standard deviation of \$25. You find a website that has a price comparison for the same computer at a series of stores as follows: \$1,299; \$1,229.99; \$1,193.08; \$1,279; \$1,224.95; \$1,229.99; \$1,269.95; \$1,249. Can you

argue that pricing has a larger standard deviation than claimed by the manufacturer? Use the 5% significance level. As a potential buyer, what would be the practical conclusion from your analysis?

85.

A company packages apples by weight. One of the weight grades is Class A apples. Class A apples have a mean weight of 150 g, and there is a maximum allowed weight tolerance of 5% above or below the mean for apples in the same consumer package. A batch of apples is selected to be included in a Class A apple package. Given the following apple weights of the batch, does the fruit comply with the Class A grade weight tolerance requirements. Conduct an appropriate hypothesis test.

- a. at the 5% significance level
- b. at the 1% significance level

Weights in selected apple batch (in grams): 158; 167; 149; 169; 164; 139; 154; 150; 157; 171; 152; 161; 141; 166; 172;

11.3 Goodness-of-Fit Test

86.

A six-sided die is rolled 120 times. Fill in the expected frequency column. Then, conduct a hypothesis test to determine if the die is fair. The data in Table 11.13.29 are the result of the 120 rolls.

Table 11.13.29

Face value	Frequency	Expected frequency
1	15	
2	29	
3	16	
4	15	
5	30	
6	15	

87.

The marital status distribution of the U.S. male population, ages 15 and older, is as shown in Table 11.13.30

Table 11.13.30

Marital status	Percent	Expected frequency
Never married	31.3	
Married	56.1	
Widowed	2.5	
Divorced/Separated	10.1	

Suppose that a random sample of 400 U.S. young adult males, 18 to 24 years old, yielded the following frequency distribution. We are interested in whether this age group of males fits the distribution of the U.S. adult population. Calculate the frequency one would expect when surveying 400 people. Fill in Table 11.13.31 rounding to two decimal places.

Table 11.13.31

Marital status	Frequency
Never married	140
Married	238
Widowed	2
Divorced/Separated	20

Use the following information to answer the next two exercises: The columns in Table 11.13.32 contain the Race/Ethnicity of U.S. Public Schools for a recent year, the percentages for the Advanced Placement Examinee Population for that class, and the Overall Student Population. Suppose the right column contains the result of a survey of 1,000 local students from that year who took an AP Exam.

Table 11.13.32

Race/Ethnicity	AP examinee population	Overall student population	Survey frequency
Asian, Asian American, or Pacific Islander	10.2%	5.4%	113
Black or African-American	8.2%	14.5%	94
Hispanic or Latino	15.5%	15.9%	136
American Indian or Alaska Native	0.6%	1.2%	10
White	59.4%	61.6%	604
Not reported/other	6.1%	1.4%	43

88.

Perform a goodness-of-fit test to determine whether the local results follow the distribution of the U.S. overall student population based on ethnicity.

89.

Perform a goodness-of-fit test to determine whether the local results follow the distribution of U.S. AP examinee population, based on ethnicity.

90.

The City of South Lake Tahoe, CA, has an Asian population of 1,419 people, out of a total population of 23,609. Suppose that a survey of 1,419 self-reported Asians in the Manhattan, NY, area yielded the data in Table 11.13.33 Conduct a goodness-of-fit test to determine if the self-reported sub-groups of Asians in the Manhattan area fit that of the Lake Tahoe area.

Table 11.13.33

Race	Lake Tahoe frequency	Manhattan frequency
Asian Indian	131	174
Chinese	118	557
Filipino	1,045	518
Japanese	80	54
Korean	12	29
Vietnamese	9	21
Other	24	66

Use the following information to answer the next two exercises: UCLA conducted a survey of more than 263,000 college freshmen from 385 colleges in fall 2005. The results of students' expected majors by gender were reported in *The Chronicle of Higher Education* (2/2/2006). Suppose a survey of 5,000 graduating females and 5,000 graduating males was done as a follow-up last year to determine what their actual majors were. The results are shown in the tables for Table 11.13.36 shows the business categories in the survey, the sample size of each category, and the number of businesses in each category that recycle one commodity. Based on the study, on average half of the businesses were expected to be recycling one commodity. As a result, the last column shows the expected number of businesses in each category that recycle one commodity. At the 5% significance level, perform a hypothesis test to determine if the observed number of businesses that recycle one commodity follows the uniform distribution of the expected values.

Table 11.13.36

--	--	--	--

Business type	Number in class	Observed number that recycle one commodity	Expected number that recycle one commodity
Office	35	19	17.5
Retail/Wholesale	48	27	24
Food/Restaurants	53	35	26.5
Manufacturing/Medical	52	21	26
Hotel/Mixed	24	9	12

98.

Table 11.13.37 contains information from a survey among 499 participants classified according to their age groups. The second column shows the percentage of obese people per age class among the study participants. The last column comes from a different study at the national level that shows the corresponding percentages of obese people in the same age classes in the USA. Perform a hypothesis test at the 5% significance level to determine whether the survey participants are a representative sample of the USA obese population.

Table 11.13.37

Age class (years)	Obese (percentage)	Expected USA average (percentage)
20–30	75.0	32.6
31–40	26.5	32.6
41–50	13.6	36.6
51–60	21.9	36.6
61–70	21.0	39.7

11.4 Test of Independence

99.

A recent debate about where in the United States skiers believe the skiing is best prompted the following survey. Test to see if the best ski area is independent of the level of the skier.

U.S. ski area	Beginner	Intermediate	Advanced
Tahoe	20	30	40
Utah	10	30	60
Colorado	10	40	50

Table 11.38

100.

Car manufacturers are interested in whether there is a relationship between the size of car an individual drives and the number of people in the driver's family (that is, whether car size and family size are independent). To test this, suppose that 800 car owners were randomly surveyed with the results in Table 11.13.39. Conduct a test of independence.

Table 11.13.39

Family Size	Sub & Compact	Mid-size	Full-size	Van & Truck
1	20	35	40	35
2	20	50	70	80
3–4	20	50	100	90
5+	20	30	70	70

101.

College students may be interested in whether or not their majors have any effect on starting salaries after graduation. Suppose that 300 recent graduates were surveyed as to their majors in college and their starting salaries after graduation. Table 11.13.40 shows the data. Conduct a test of independence.

Major	< \$50,000	\$50,000 – \$68,999	\$69,000 +
English	5	20	5
Engineering	10	30	60
Nursing	10	15	15
Business	10	20	30
Psychology	20	30	20

Table 11.40

102.

Some travel agents claim that honeymoon hot spots vary according to age of the bride. Suppose that 280 recent brides were interviewed as to where they spent their honeymoons. The information is given in Table 11.13.41. Conduct a test of independence.

Table 11.13.41

Location	20–29	30–39	40–49	50 and over
Niagara Falls	15	25	25	20
Poconos	15	25	25	10
Europe	10	25	15	5
Virgin Islands	20	25	15	5

103.

A manager of a sports club keeps information concerning the main sport in which members participate and their ages. To test whether there is a relationship between the age of a member and his or her choice of sport, 643 members of the sports club are randomly selected. Conduct a test of independence.

Sport	18 - 25	26 - 30	31 - 40	41 and over
Racquetball	42	58	30	46
Tennis	58	76	38	65
Swimming	72	60	65	33

Table 11.42

104.

A major food manufacturer is concerned that the sales for its skinny french fries have been decreasing. As a part of a feasibility study, the company conducts research into the types of fries sold across the country to determine if the type of fries sold is independent of the area of the country. The results of the study are shown in Table 11.13.43. Conduct a test of independence.

Table 11.13.43

Type of Fries	Northeast	South	Central	West
Skinny fries	70	50	20	25
Curly fries	100	60	15	30
Steak fries	20	40	10	10

105.

According to Dan Lenard, an independent insurance agent in the Buffalo, N.Y. area, the following is a breakdown of the amount of life insurance purchased by males in the following age groups. He is interested in whether the age of the male and the amount of life insurance purchased are independent events. Conduct a test for independence.

Age of males	None	< \$200,000	\$200,000–\$400,000	\$401,001–\$1,000,000	\$1,000,001+
20–29	40	15	40	0	5
30–39	35	5	20	20	10
40–49	20	0	30	0	30
50+	40	30	15	15	10

Table 11.44

106.

Suppose that 600 thirty-year-olds were surveyed to determine whether or not there is a relationship between the level of education an individual has and salary. Conduct a test of independence.

Table 11.13.45

Annual salary	Not a high school graduate	High school graduate	College graduate	Masters or doctorate
< \$30,000	15	25	10	5
\$30,000–\$40,000	20	40	70	30
\$40,000–\$50,000	10	20	40	55
\$50,000–\$60,000	5	10	20	60
\$60,000+	0	5	10	150

Read the statement and decide whether it is true or false.

107.

The number of degrees of freedom for a test of independence is equal to the sample size minus one.

108.

The test for independence uses tables of observed and expected data values.

109.

The test to use when determining if the college or university a student chooses to attend is related to his or her socioeconomic status is a test for independence.

110.

In a test of independence, the expected number is equal to the row total multiplied by the column total divided by the total surveyed.

111.

An ice cream maker performs a nationwide survey about favorite flavors of ice cream in different geographic areas of the U.S. Based on Table 11.13.46 do the numbers suggest that geographic location is independent of favorite ice cream flavors? Test at the 5% significance level.

U.S. region/Flavor	Strawberry	Chocolate	Vanilla	Rocky road	Mint chocolate chip	Pistachio	Row total
West	12	21	22	19	15	8	97
Midwest	10	32	22	11	15	6	96
East	8	31	27	8	15	7	96
South	15	28	30	8	15	6	102

U.S. region/Flavor	Strawberry	Chocolate	Vanilla	Rocky road	Mint chocolate chip	Pistachio	Row total
Column total	45	112	101	46	60	27	391

Table 11.46

112.

Table 11.13.47 provides a recent survey of the youngest online entrepreneurs whose net worth is estimated at one million dollars or more. Their ages range from 17 to 30. Each cell in the table illustrates the number of entrepreneurs who correspond to the specific age group and their net worth. Are the ages and net worth independent? Perform a test of independence at the 5% significance level.

Table 11.13.47

Age group \ Net worth value (in millions of US dollars)	1–5	6–24	≥25	Row total
17–25	8	7	5	20
26–30	6	5	9	20
Column total	14	12	14	40

113.

A 2013 poll in California surveyed people about taxing sugar-sweetened beverages. The results are presented in Table 11.13.48 and are classified by ethnic group and response type. Are the poll responses independent of the participants' ethnic group? Conduct a test of independence at the 5% significance level.

Opinion/Ethnicity	Asian-American	White/Non-Hispanic	African-American	Latino	Row total
Against tax	48	433	41	160	682
In favor of tax	54	234	24	147	459
No opinion	16	43	16	19	94
Column total	118	710	81	326	1235

Table 11.48

11.5 Test for Homogeneity

114.

A psychologist is interested in testing whether there is a difference in the distribution of personality types for business majors and social science majors. The results of the study are shown in Table 11.13.49 Conduct a test of homogeneity. Test at a 5% level of significance.

Table 11.13.49

	Open	Conscientious	Extrovert	Agreeable	Neurotic
Business	41	52	46	61	58
Social Science	72	75	63	80	65

115.

Do men and women select different breakfasts? The breakfasts ordered by randomly selected men and women at a popular breakfast place is shown in Table 11.13.50 Conduct a test for homogeneity at a 5% level of significance.

	French toast	Pancakes	Waffles	Omelettes

Men	47	35	28	53
Women	65	59	55	60

Table 11.50

116.

A fisherman is interested in whether the distribution of fish caught in Green Valley Lake is the same as the distribution of fish caught in Echo Lake. Of the 191 randomly selected fish caught in Green Valley Lake, 105 were rainbow trout, 27 were other trout, 35 were bass, and 24 were catfish. Of the 293 randomly selected fish caught in Echo Lake, 115 were rainbow trout, 58 were other trout, 67 were bass, and 53 were catfish. Perform a test for homogeneity at a 5% level of significance.

117.

In 2007, the United States had 1.5 million homeschooled students, according to the U.S. National Center for Education Statistics. In Table 11.13.51 you can see that parents decide to homeschool their children for different reasons, and some reasons are ranked by parents as more important than others. According to the survey results shown in the table, is the distribution of applicable reasons the same as the distribution of the most important reason? Provide your assessment at the 5% significance level. Did you expect the result you obtained?

Reasons for fomeschooling	Applicable reason (in thousands of respondents)	Most important reason (in thousands of respondents)	Row total
Concern about the environment of other schools	1,321	309	1,630
Dissatisfaction with academic instruction at other schools	1,096	258	1,354
To provide religious or moral instruction	1,257	540	1,797
Child has special needs, other than physical or mental	315	55	370
Nontraditional approach to child's education	984	99	1,083
Other reasons (e.g., finances, travel, family time, etc.)	485	216	701
Column total	5,458	1,477	6,935

Table 11.51

118.

When looking at energy consumption, we are often interested in detecting trends over time and how they correlate among different countries. The information in Table 11.13.52 shows the average energy use (in units of kg of oil equivalent per capita) in the USA and the joint European Union countries (EU) for the six-year period 2005 to 2010. Do the energy use values in these two areas come from the same distribution? Perform the analysis at the 5% significance level.

Table 11.13.52

Year	European Union	United States	Row total
2010	3,413	7,164	10,557
2009	3,302	7,057	10,359
2008	3,505	7,488	10,993
2007	3,537	7,758	11,295
2006	3,595	7,697	11,292
2005	3,613	7,847	11,460

Year	European Union	United States	Row total
Column total	20,965	45,011	65,976

119.

The Insurance Institute for Highway Safety collects safety information about all types of cars every year, and publishes a report of Top Safety Picks among all cars, makes, and models. Table 11.13.53 presents the number of Top Safety Picks in six car categories for the two years 2009 and 2013. Analyze the table data to conclude whether the distribution of cars that earned the Top Safety Picks safety award has remained the same between 2009 and 2013. Derive your results at the 5% significance level.

Year \ Car type	Small	Mid-size	Large	Small SUV	Mid-size SUV	Large SUV	Row total
2009	12	22	10	10	27	6	87
2013	31	30	19	11	29	4	124
Column total	43	52	29	21	56	10	211

Table 11.53

11.6 Comparison of the Chi-Square Tests

120.

Is there a difference between the distribution of community college statistics students and the distribution of university statistics students in what technology they use on their homework? Of some randomly selected community college students, 43 used a computer, 102 used a calculator with built in statistics functions, and 65 used a table from the textbook. Of some randomly selected university students, 28 used a computer, 33 used a calculator with built in statistics functions, and 40 used a table from the textbook. Conduct an appropriate hypothesis test using a 0.05 level of significance.

Read the statement and decide whether it is true or false.

121.

If $df = 2$, the chi-square distribution has a shape that reminds us of the exponential.

This page titled [11.13: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.9: Chapter Homework** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.14: Chapter References

11.1 Facts About the Chi-Square Distribution

Data from *Parade Magazine*.

“HIV/AIDS Epidemiology Santa Clara County.” Santa Clara County Public Health Department, May 2011.

11.2 Test of a Single Variance

“AppleInsider Price Guides.” Apple Insider, 2013. Available online at http://appleinsider.com/mac_price_guide (accessed May 14, 2013).

Data from the World Bank, June 5, 2012.

11.3 Goodness-of-Fit Test

Data from the U.S. Census Bureau

Data from the College Board. Available online at <http://www.collegeboard.com>.

Data from the U.S. Census Bureau, Current Population Reports.

Ma, Y., E.R. Bertone, E.J. Stanek III, G.W. Reed, J.R. Hebert, N.L. Cohen, P.A. Merriam, I.S. Ockene, “Association between Eating Patterns and Obesity in a Free-living US Adult Population.” *American Journal of Epidemiology* volume 158, no. 1, pages 85-92.

Ogden, Cynthia L., Margaret D. Carroll, Brian K. Kit, Katherine M. Flegal, “Prevalence of Obesity in the United States, 2009–2010.” NCHS Data Brief no. 82, January 2012. Available online at <http://www.cdc.gov/nchs/data/databriefs/db82.pdf> (accessed May 24, 2013).

Stevens, Barbara J., “Multi-family and Commercial Solid Waste and Recycling Survey.” Arlington Count, VA. Available online at www.arlingtonva.us/department.../file84429.pdf (accessed May 24, 2013).

11.4 Test of Independence

DiCamilo, Mark, Mervin Field, “Most Californians See a Direct Linkage between Obesity and Sugary Sodas. Two in Three Voters Support Taxing Sugar-Sweetened Beverages If Proceeds are Tied to Improving School Nutrition and Physical Activity Programs.” The Field Poll, released Feb. 14, 2013. Available online at field.com/fieldpollonline/sub...rs/Rls2436.pdf (accessed May 24, 2013).

Harris Interactive, “Favorite Flavor of Ice Cream.” Available online at <http://www.statisticbrain.com/favori...r-of-ice-cream> (accessed May 24, 2013)

“Youngest Online Entrepreneurs List.” Available online at <http://www.statisticbrain.com/younge...repreneur-list> (accessed May 24, 2013).

11.5 Test for Homogeneity

Data from the Insurance Institute for Highway Safety, 2013. Available online at www.iihs.org/iihs/ratings (accessed May 24, 2013).

“Energy use (kg of oil equivalent per capita).” The World Bank, 2013. Available online at <http://data.worldbank.org/indicator/...G.OE/countries> (accessed May 24, 2013).

“Parent and Family Involvement Survey of 2007 National Household Education Survey Program (NHES),” U.S. Department of Education, National Center for Education Statistics. Available online at <http://nces.ed.gov/pubsearch/pubsinf...?pubid=2009030> (accessed May 24, 2013).

“Parent and Family Involvement Survey of 2007 National Household Education Survey Program (NHES),” U.S. Department of Education, National Center for Education Statistics. Available online at http://nces.ed.gov/pubs2009/2009030_sup.pdf (accessed May 24, 2013).

This page titled [11.14: Chapter References](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.12: Chapter References** by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

11.15: Chapter Solution (Practice + Homework)

1.

mean = 25 and standard deviation = 7.0711

3.

when the number of degrees of freedom is greater than 90

5.

$$df = 2$$

6.

a test of a single variance

8.

a left-tailed test

10.

$$H_0 : \sigma^2 = 0.812;$$

$$H_a : \sigma^2 > 0.812.$$

12.

a test of a single variance

16.

a goodness-of-fit test

18.

3

20.

2.04

21.

We decline to reject the null hypothesis. There is not enough evidence to suggest that the observed test scores are significantly different from the expected test scores.

23.

H_0 : the distribution of AIDS cases follows the ethnicities of the general population of Santa Clara County.

25.

right-tailed

27.

2016.136

28.

• 30.

a test of independence

a test of independence

34.

8

36.

6.6

39.

Table 11.15.54

Smoking level per day	African American	Native Hawaiian	Latino	Japanese Americans	White	Totals
1-10	9,886	2,745	12,831	8,378	7,650	41,490
11-20	6,514	3,062	4,932	10,680	9,877	35,065
21-30	1,671	1,419	1,406	4,715	6,062	15,273
31+	759	788	800	2,305	3,970	8,622
Totals	18,830	8,014	19,969	26,078	27,559	10,0450

41.

Table 11.15.55

Smoking level per day	African American	Native Hawaiian	Latino	Japanese Americans	White
1-10	7777.57	3310.11	8248.02	10771.29	11383.01
11-20	6573.16	2797.52	6970.76	9103.29	9620.27
21-30	2863.02	1218.49	3036.20	3965.05	4190.23
31+	1616.25	687.87	1714.01	2238.37	2365.49

43.

10,301.8

44.

right

46.

1. 48.

test for homogeneity

test for homogeneity

52.

All values in the table must be greater than or equal to five.

54.

3

57.

a goodness-of-fit test

59.

a test for independence

61.

Answers will vary. Sample answer: Tests of independence and tests for homogeneity both calculate the test statistic the same way $\sum_{(ij)} \frac{(O-E)^2}{E}$. In addition, all values must be greater than or equal to five.

63.

true

65.

false

67.

225

69.

$$H_0 : \sigma^2 \leq 150$$

71.

36

72.

Check student's solution.

74.

The claim is that the variance is no more than 150 minutes.

76.

a Student's t - or normal distribution

78.

1. 80.

1. 82.

1. 84.

1. 87.

Table 11.15.56

Marital status	Percent	Expected frequency
Never married	31.3	125.2
Married	56.1	224.4
Widowed	2.5	10
Divorced/Separated	10.1	40.4

1. 89.

1. 91.

1. 94.

true

false

98.

1. 100.

1. 102.

1. 104.

1. 106.

1. 108.

true

true

112.

1. 114.

1. 116.

1. 118.

1. 120.

1. 122.

1. The test statistic is always positive and if the expected and observed values are not close together, the test statistic is large and the null hypothesis will be rejected.
2. Testing to see if the data fits the distribution “too well” or is too perfect.

This page titled [11.15: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **11.14: Chapter Solution (Practice + Homework)** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

12: F Distribution and One-Way ANOVA

[12.1: Introduction to F Distribution and One-Way ANOVA](#)

[12.2: Test of Two Variances](#)

[12.3: One-Way ANOVA](#)

[12.4: The F Distribution and the F-Ratio](#)

[12.5: Facts About the F Distribution](#)

[12.6: Chapter Key Terms](#)

[12.7: Chapter Review](#)

[12.8: Chapter Formula Review](#)

[12.9: Chapter Practice](#)

[12.10: Chapter Homework](#)

[12.11: Chapter Reference](#)

[12.12: Chapter Solution \(Practice + Homework\)](#)

This page titled [12: F Distribution and One-Way ANOVA](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

12.1: Introduction to F Distribution and One-Way ANOVA



Figure 12.1.1 One-way ANOVA is used to measure information from several groups.

Many statistical applications in psychology, social science, business administration, and the natural sciences involve several groups. For example, an environmentalist is interested in knowing if the average amount of pollution varies in several bodies of water. A sociologist is interested in knowing if the amount of income a person earns varies according to his or her upbringing. A consumer looking for a new car might compare the average gas mileage of several models.

For hypothesis tests comparing averages among more than two groups, statisticians have developed a method called "Analysis of Variance" (abbreviated ANOVA). In this chapter, you will study the simplest form of ANOVA called single factor or one-way ANOVA. You will also study the F distribution, used for one-way ANOVA, and the test for differences between two variances. This is just a very brief overview of one-way ANOVA. One-Way ANOVA, as it is presented here, relies heavily on a calculator or computer.

This page titled [12.1: Introduction to F Distribution and One-Way ANOVA](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.0: Introduction to F Distribution and One-Way ANOVA](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.2: Test of Two Variances

This chapter introduces a new probability density function, the F distribution. This distribution is used for many applications including ANOVA and for testing equality across multiple means. We begin with the F distribution and the test of hypothesis of differences in variances. It is often desirable to compare two variances rather than two averages. For instance, college administrators would like two college professors grading exams to have the same variation in their grading. In order for a lid to fit a container, the variation in the lid and the container should be approximately the same. A supermarket might be interested in the variability of check-out times for two checkers. In finance, the variance is a measure of risk and thus an interesting question would be to test the hypothesis that two different investment portfolios have the same variance, the volatility.

In order to perform a F test of two variances, it is important that the following are true:

1. The populations from which the two samples are drawn are approximately normally distributed.
2. The two populations are independent of each other.

Unlike most other hypothesis tests in this book, the F test for equality of two variances is very sensitive to deviations from normality. If the two distributions are not normal, or close, the test can give a biased result for the test statistic.

Suppose we sample randomly from two independent normal populations. Let σ_1^2 and σ_2^2 be the unknown population variances and s_1^2 and s_2^2 be the sample variances. Let the sample sizes be n_1 and n_2 . Since we are interested in comparing the two sample variances, we use the F ratio:

$$F = \frac{\left[\frac{s_1^2}{\sigma_1^2} \right]}{\left[\frac{s_2^2}{\sigma_2^2} \right]}$$

F has the distribution $F \sim F(n_1 - 1, n_2 - 1)$

where $n_1 - 1$ are the degrees of freedom for the numerator and $n_2 - 1$ are the degrees of freedom for the denominator.

If the null hypothesis is $\sigma_1^2 = \sigma_2^2$, then the F Ratio, test statistic, becomes $F_c = \frac{\left[\frac{s_1^2}{\sigma_1^2} \right]}{\left[\frac{s_2^2}{\sigma_2^2} \right]} = \frac{s_1^2}{s_2^2}$

The various forms of the hypotheses tested are:

Table 12.1

Two-Tailed Test	One-Tailed Test	One-Tailed Test
$H_0 : \sigma_1^2 = \sigma_2^2$	$H_0 : \sigma_1^2 \leq \sigma_2^2$	$H_0 : \sigma_1^2 \geq \sigma_2^2$
$H_1 : \sigma_1^2 \neq \sigma_2^2$	$H_1 : \sigma_1^2 > \sigma_2^2$	$H_1 : \sigma_1^2 < \sigma_2^2$

A more general form of the null and alternative hypothesis for a two tailed test would be :

$$H_0 : \frac{\sigma_1^2}{\sigma_2^2} = \delta_0$$

$$H_a : \frac{\sigma_1^2}{\sigma_2^2} \neq \delta_0$$

Where if $\delta_0 = 1$ it is a simple test of the hypothesis that the two variances are equal. This form of the hypothesis does have the benefit of allowing for tests that are more than for simple differences and can accommodate tests for specific differences as we did for differences in means and differences in proportions. This form of the hypothesis also shows the relationship between the F distribution and the χ^2 : the F is a ratio of two chi squared distributions a distribution we saw in the last chapter. This is helpful in determining the degrees of freedom of the resultant F distribution.

If the two populations have equal variances, then s_1^2 and s_2^2 are close in value and the test statistic, $F_c = \frac{s_1^2}{s_2^2}$ is close to one. But if the two population variances are very different, s_1^2 and s_2^2 tend to be very different, too. Choosing s_1^2 as the larger sample variance causes the ratio $\frac{s_1^2}{s_2^2}$ to be greater than one. If s_1^2 and s_2^2 are far apart, then $F_c = \frac{s_1^2}{s_2^2}$ is a large number.

Therefore, if F is close to one, the evidence favors the null hypothesis (the two population variances are equal). But if F is much larger than one, then the evidence is against the null hypothesis. In essence, we are asking if the calculated F statistic, test statistic, is significantly different from one.

To determine the critical points we have to find F_{α, df_1, df_2} . See Appendix A for the F table. This F table has values for various levels of significance from 0.1 to 0.001 designated as "p" in the first column. To find the critical value choose the desired significance level and follow down and across to find the critical value at the intersection of the two different degrees of freedom. The F distribution has two different degrees of freedom, one associated with the numerator, df_1 , and one associated with the denominator, df_2 and to complicate matters the F distribution is not symmetrical and changes the degree of skewness as the degrees of freedom change. The degrees of freedom in the numerator is $n_1 - 1$, where n_1 is the sample size for group 1, and the degrees of freedom in the denominator is $n_2 - 1$, where n_2 is the sample size for group 2. F_{α, df_1, df_2} will give the critical value on the **upper** end of the F distribution.

To find the critical value for the **lower** end of the distribution, reverse the degrees of freedom and divide the F -value from the table into one.

- Upper tail critical value : F_{α, df_1, df_2}
- Lower tail critical value : $1/F_{\alpha, df_2, df_1}$

When the calculated value of F is between the critical values, not in the tail, we cannot reject the null hypothesis that the two variances came from a population with the same variance. If the calculated F -value is in either tail we cannot accept the null hypothesis just as we have been doing for all of the previous tests of hypothesis.

An alternative way of finding the critical values of the F distribution makes the use of the F -table easier. We note in the F -table that all the values of F are greater than one therefore the critical F value for the left hand tail will always be less than one because to find the critical value on the left tail we divide an F value into the number one as shown above. We also note that if the sample variance in the numerator of the test statistic is larger than the sample variance in the denominator, the resulting F value will be greater than one. The shorthand method for this test is thus to be sure that the larger of the two sample variances is placed in the numerator to calculate the test statistic. This will mean that only the right hand tail critical value will have to be found in the F -table.

Example 12.1

Two college instructors are interested in whether or not there is any variation in the way they grade math exams. They each grade the same set of 10 exams. The first instructor's grades have a variance of 52.3. The second instructor's grades have a variance of 89.9. Test the claim that the first instructor's variance is smaller. (In most colleges, it is desirable for the variances of exam grades to be nearly the same among instructors.) The level of significance is 10%.

Answer

Solution 12.1

Let 1 and 2 be the subscripts that indicate the first and second instructor, respectively.

$$n_1 = n_2 = 10 .$$

$$H_0 : \sigma_1^2 \geq \sigma_2^2 \text{ and } H_a : \sigma_1^2 < \sigma_2^2$$

Calculate the test statistic: By the null hypothesis ($\sigma_1^2 \geq \sigma_2^2$), the F statistic is:

$$F_c = \frac{s_2^2}{s_1^2} = \frac{89.9}{52.3} = 1.719$$

Critical value for the test: $F_{9,9} = 5.35$ where $n_1 - 1 = 9$ and $n_2 - 1 = 9$.




Figure 12.2

Make a decision: Since the calculated F value is not in the tail we cannot reject H_0 .

Conclusion: With a 10% level of significance, from the data, there is insufficient evidence to conclude that the variance in grades for the first instructor is smaller.

Exercise 12.1

The New York Choral Society divides male singers up into four categories from highest voices to lowest: Tenor1, Tenor2, Bass1, Bass2. In the table are heights of the men in the Tenor1 and Bass2 groups. One suspects that taller men will have lower voices, and that the variance of height may go up with the lower voices as well. Do we have good evidence that the variance of the heights of singers in each of these two groups (Tenor1 and Bass2) are different?

Table 12.2

Tenor1	Bass 2	Tenor 1	Bass 2	Tenor 1	Bass 2
69	72	67	72	68	67
72	75	70	74	67	70
71	67	65	70	64	70
66	75	72	66		69
76	74	70	68		72
74	72	68	75		71
71	72	64	68		74
66	74	73	70		75
68	72	66	72		

This page titled [12.2: Test of Two Variances](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.1: Test of Two Variances](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.3: One-Way ANOVA

The purpose of a one-way ANOVA test is to determine the existence of a statistically significant difference among several group means. The test actually uses **variances** to help determine if the means are equal or not. In order to perform a one-way ANOVA test, there are five basic **assumptions** to be fulfilled:

1. The null hypothesis is simply that all the group population means are the same. The alternative hypothesis is that at least one pair of means is different. For example, if there are k groups:

$$H_0 : \mu_1 = \mu_2 = \mu_3 = \dots \mu_k$$

The graphs, a set of box plots representing the distribution of values with the group means indicated by a horizontal line through the box, help in the understanding of the hypothesis test. In the first graph (red box plots), $H_0 : \mu_1 = \mu_2 = \mu_3$ and the three populations have the same distribution if the null hypothesis is true. The variance of the combined data is approximately the same as the variance of each of the populations.

If the null hypothesis is false, then the variance of the combined data is larger which is caused by the different means as shown in the second graph (green box plots).




Figure 12.3.3 (a) H_0 is true. All means are the same; the differences are due to random variation. (b) H_0 is not true. All means are not the same; the differences are too large to be due to random variation.

This page titled [12.3: One-Way ANOVA](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.2: One-Way ANOVA](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.4: The F Distribution and the F-Ratio

The distribution used for the hypothesis test is a new one. It is called the **F-distribution**, invented by George Snedecor but named in honor of Sir Ronald Fisher, an English statistician. The F statistic is a ratio (a fraction). There are two sets of degrees of freedom; one for the numerator and one for the denominator.

For example, if F follows an F distribution and the number of degrees of freedom for the numerator is four, and the number of degrees of freedom for the denominator is ten, then $F \sim F_{4,10}$.

To calculate the **F ratio**, two estimates of the variance are made.

- Variance between samples:** An estimate of σ^2 that is the variance of the sample means multiplied by n (when the sample sizes are the same.). If the samples are different sizes, the variance between samples is weighted to account for the different sample sizes. The variance is also called **variation due to treatment or explained variation**.
 - Variance within samples:** An estimate of σ^2 that is the average of the sample variances (also known as a pooled variance). When the sample sizes are different, the variance within samples is weighted. The variance is also called the **variation due to error or unexplained variation**.
- $SS_{between}$ is the **sum of squares** that represents the variation among the different samples
 - SS_{within} is the sum of squares that represents the variation within samples that is due to chance.

To find a "sum of squares" means to add together squared quantities that, in some cases, may be weighted. We used sum of squares to calculate the sample variance and the sample standard deviation in [Table 1.19](#).

MS means "**mean square**." $MS_{between}$ is the variance between groups, and MS_{within} is the variance within groups.

Calculation of Sum of Squares and Mean Square

- k is the the number of different groups
- n_j is the size of the j^{th} group
- s_j = the sum of the values in the j^{th} group
- n is the total number of all the values combined (total sample size: Σn_j)
- x is the one value:

$$\sum x = \sum s_j$$

- Sum of squares of all values from every group combined:

$$\sum x^2$$

- Between group variability:

$$SS_{total} = \sum x^2 - \frac{(\sum x)^2}{n}$$

- Total sum of squares:

$$\sum x^2 - \frac{(\sum x)^2}{n}$$

- Explained variation: sum of squares representing variation among the different samples:

$$SS_{between} = \sum \left[\frac{(s_j)^2}{n_j} \right] - \frac{(\sum s_j)^2}{n}$$

- Unexplained variation: sum of squares representing variation within samples due to chance:

$$SS_{within} = SS_{total} - SS_{between}$$

- df 's for different groups (df 's for the numerator):

$$df = k - 1$$

- Equation for errors within samples (df 's for the denominator):

$$df_{within} = n - k$$

- Mean square (variance estimate) explained by the different groups:

$$MS_{\text{between}} = \frac{SS_{\text{between}}}{df_{\text{between}}}$$

- Mean square (variance estimate) that is due to chance (unexplained):

$$MS_{\text{within}} = \frac{SS_{\text{within}}}{df_{\text{within}}}$$

MS_{between} and MS_{within} can be written as follows:

$$MS_{\text{between}} = \frac{SS_{\text{between}}}{df_{\text{between}}} = \frac{SS_{\text{between}}}{k - 1}$$

$$MS_{\text{within}} = \frac{SS_{\text{within}}}{df_{\text{within}}} = \frac{SS_{\text{within}}}{n - k}$$

The one-way ANOVA test depends on the fact that MS_{between} can be influenced by population differences among means of the several groups. Since MS_{within} compares values of each group to its own group mean, the fact that group means might be different does not affect MS_{within} .

The null hypothesis says that all groups are samples from populations having the same normal distribution. The alternate hypothesis says that at least two of the sample groups come from populations with different normal distributions. If the null hypothesis is true, MS_{between} and MS_{within} should both estimate the same value.

Note

The null hypothesis says that all the group population means are equal. The hypothesis of equal means implies that the populations have the same normal distribution, because it is assumed that the populations are normal and that they have equal variances.

Definition: F-Ratio or F Statistic

$$F = \frac{MS_{\text{between}}}{MS_{\text{within}}} \quad (12.4.1)$$

If MS_{between} and MS_{within} estimate the same value (following the belief that H_0 is true), then the F -ratio should be approximately equal to one. Mostly, just sampling errors would contribute to variations away from one. As it turns out, MS_{between} consists of the population variance plus a variance produced from the differences between the samples. MS_{within} is an estimate of the population variance. Since variances are always positive, if the null hypothesis is false, MS_{between} will generally be larger than MS_{within} . Then the F -ratio will be larger than one. However, if the population effect is small, it is not unlikely that MS_{within} will be larger in a given sample.

The foregoing calculations were done with groups of different sizes. If the groups are the same size, the calculations simplify somewhat and the F -ratio can be written as:

F-Ratio Formula when the groups are the same size

The foregoing calculations were done with groups of different sizes. If the groups are the same size, the calculations simplify somewhat and the F -ratio can be written as

$$F = \frac{n \cdot s_{\bar{x}}^2}{s_{\text{pooled}}^2} \quad (12.4.2)$$

where

- n = the sample size
- $df_{\text{numerator}} = k - 1$
- $df_{\text{denominator}} = n - k$
- s_{pooled}^2 = the mean of the sample variances (pooled variance)
- $s_{\bar{x}}^2$ = the variance of the sample means

Data are typically put into a table for easy viewing. One-Way ANOVA results are often displayed in this manner by computer software.

Table 12.4.1

Source of variation	Sum of squares (SS)	Degrees of freedom (df)	Mean square (MS)	F
Factor (Between)	$SS(\text{Factor})$	$k-1$	$MS(\text{Factor}) = \frac{SS(\text{Factor})}{k-1}$	$F = \frac{MS(\text{Factor})}{MS(\text{Error})}$
Error (Within)	$SS(\text{Error})$	$n-k$	$MS(\text{Error}) = \frac{SS(\text{Error})}{n-k}$	
Total	$SS(\text{Total})$	$n-1$		

Example 12.2

Three different diet plans are to be tested for mean weight loss. The entries in the table are the weight losses for the different plans. The one-way ANOVA results are shown in Table 12.4.2

Table 12.4.2

Plan 1: $n_1 = 4$	Plan 2: $n_2 = 3$	Plan 3: $n_3 = 3$
5	3.5	8
4.5	7	4
4		3.5
3	4.5	

$$s_1 = 16.5, s_2 = 15, s_3 = 15.5$$

Following are the calculations needed to fill in the one-way ANOVA table. The table is used to conduct a hypothesis test.

$$\begin{aligned}
 S(\text{between}) &= \sum \left[\frac{(s_j)^2}{n_j} \right] - \frac{(\sum s_j)^2}{n} \\
 &= \frac{s_1^2}{4} + \frac{s_2^2}{3} + \frac{s_3^2}{3} - \frac{(s_1 + s_2 + s_3)^2}{10}
 \end{aligned}$$

where $n_1 = 4, n_2 = 3, n_3 = 3$ and $n = n_1 + n_2 + n_3 = 10$.

$$\begin{aligned}
 S(\text{between}) &= \frac{(16.5)^2}{4} + \frac{(15)^2}{3} + \frac{(15.5)^2}{3} - \frac{(16.5 + 15 + 15.5)^2}{10} \\
 &= 2.2458
 \end{aligned}$$

$$\begin{aligned}
 S(\text{total}) &= \sum x^2 - \frac{(\sum x)^2}{n} \\
 &= (5^2 + 4.5^2 + 4^2 + 3^2 + 3.5^2 + 7^2 + 4.5^2 + 8^2 + 4^2 + 3.5^2) - \frac{(5 + 4.5 + 4 + 3 + 3.5 + 7 + 4.5 + 8 + 4 + 3.5)^2}{10} \\
 &= 244 - \frac{47^2}{10} \\
 &= 244 - 220.9 \\
 &= 23.1
 \end{aligned}$$

$$\begin{aligned}
 S(\text{within}) &= S(\text{total}) - SS(\text{between}) \\
 &= 23.1 - 2.2458 \\
 &= 20.8542
 \end{aligned}$$

Table 12.4.3

Source of variation	Sum of squares (SS)	Degrees of freedom (df)	Mean square (MS)	F
Factor (Between)	$SS(\text{Factor}) = SS(\text{Between}) = 2.2458$	$k-1 = 3\text{ groups}-1 = 2$	$MS(\text{Factor}) = \frac{SS(\text{Factor})}{k-1} = \frac{2.2458}{2} = 1.1229$	$F = \frac{MS(\text{Factor})}{MS(\text{Error})} = \frac{1.1229}{2.9792} = 0.3769$

Source of variation	Sum of squares (SS)	Degrees of freedom (df)	Mean square (MS)	F
Error (Within)	$SS(Error) = SS(Within) = 20.8542$	$n - k = 10 - 3 = 7$	$MS(Error) = \frac{SS(Error)}{n - k} = \frac{20.8542}{7} = 2.9792$	
Total	$SS(Total) = 2.2458 + 20.8542 = 23.1$	$10 - 1 = 9$		

Exercise 12.2

As part of an experiment to see how different types of soil cover would affect slicing tomato production, Marist College students grew tomato plants under different soil cover conditions. Groups of three plants each had one of the following treatments

- bare soil
- a commercial ground cover
- black plastic
- straw
- compost

All plants grew under the same conditions and were the same variety. Students recorded the weight (in grams) of tomatoes produced by each of the $n = 15$ plants:

Table 12.4.4

Bare: $n_1 = 3$	Ground Cover: $n_2 = 3$	Plastic: $n_3 = 3$	Straw: $n_4 = 3$	Compost: $n_5 = 3$
2,625	5,348	6,583	7,285	6,277
2,997	5,682	8,560	6,897	7,818
4,915	5,482	3,830	9,230	8,677

Create the one-way ANOVA table.

The one-way ANOVA hypothesis test is **always right-tailed** because larger F -values are way out in the right tail of the F -distribution curve and tend to make us reject H_0 .

Example 12.3

Let's return to the slicing tomato exercise in [Try It](#). The means of the tomato yields under the five mulching conditions are represented by $\mu_1, \mu_2, \mu_3, \mu_4, \mu_5$. We will conduct a hypothesis test to determine if all means are the same or at least one is different. Using a significance level of 5%, test the null hypothesis that there is no difference in mean yields among the five groups against the alternative hypothesis that at least one mean is different from the rest.

Answer

The null and alternative hypotheses are:

$$H_0 : \mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5$$

$$H_a : \mu_i \neq \mu_j \text{ some } i \neq j$$

The one-way ANOVA results are shown in Table 12.4.5

Table 12.4.5

Source of variation	Sum of squares (SS)	Degrees of freedom (df)	Mean square (MS)	F
Factor (Between)	36,648,561	$5 - 1 = 4$	$\frac{36,648,561}{4} = 9,162,140$	$\frac{9,162,140}{2,044,672.6} = 4.4810$
Error (Within)	20,446,726	$15 - 5 = 10$	$\frac{20,446,726}{10} = 2,044,672.6$	
Total	57,095,287	$15 - 1 = 14$		

Distribution for the test: $F_{4,10}$

$$df(num) = 5 - 1 = 4$$

$$df(\text{denom}) = 15 - 5 = 10$$

Test statistic: $F = 4.4810$




Figure 12.4.1

Probability Statement: $p\text{-value} = P(F > 4.481) = 0.0248$.

Compare α and the p-value: $\alpha = 0.05$, $p\text{-value} = 0.0248$

Make a decision: Since $\alpha > p\text{-value}$, we cannot accept H_0 .

Conclusion: At the 5% significance level, we have reasonably strong evidence that differences in mean yields for slicing tomato plants grown under different mulching conditions are unlikely to be due to chance alone. We may conclude that at least some of mulches led to different mean yields.

Exercise 12.3

MRSA, or *Staphylococcus aureus*, can cause a serious bacterial infections in hospital patients. Table 12.4.6 shows various colony counts from different patients who may or may not have MRSA. The data from the table is plotted in Figure 12.4.2

Table 12.4.6

Conc = 0.6	Conc = 0.8	Conc = 1.0	Conc = 1.2	Conc = 1.4
9	16	22	30	27
66	93	147	199	168
98	82	120	148	132

Plot of the data for the different concentrations:




Figure 12.4.2

Test whether the mean number of colonies are the same or are different. Construct the ANOVA table, find the p-value, and state your conclusion. Use a 5% significance level.

Example 12.4

Four sororities took a random sample of sisters regarding their grade means for the past term. The results are shown in Table 12.4.7.

Table 12.4.7: Mean grades for four sororities

Sorority 1	Sorority 2	Sorority 3	Sorority 4
2.17	2.63	2.63	3.79
1.85	1.77	3.78	3.45

Sorority 1	Sorority 2	Sorority 3	Sorority 4
2.83	3.25	4.00	3.08
1.69	1.86	2.55	2.26
3.33	2.21	2.45	3.18

Using a significance level of 1%, is there a difference in mean grades among the sororities?

Answer

Let $\mu_1, \mu_2, \mu_3, \mu_4$ be the population means of the sororities. Remember that the null hypothesis claims that the sorority groups are from the same normal distribution. The alternate hypothesis says that at least two of the sorority groups come from populations with different normal distributions. Notice that the four sample sizes are each five.

Note: This is an example of a **balanced design**, because each factor (i.e., sorority) has the same number of observations.

$$H_0 : \mu_1 = \mu_2 = \mu_3 = \mu_4$$

H_a : Not all of the means $\mu_1, \mu_2, \mu_3, \mu_4$ are equal.

Distribution for the test: $F_{3,16}$

where $k = 4$ groups and $n = 20$ samples in total

$$df(num) = k - 1 = 4 - 1 = 3$$

$$df(denom) = n - k = 20 - 4 = 16$$

Calculate the test statistic: $F = 2.23$

Graph:




Figure 12.4.3

Probability statement: $p\text{-value} = P(F > 2.23) = 0.1241$

Compare α and the p-value: $\alpha = 0.01$

$$p\text{-value} = 0.1241$$

$$\alpha < p\text{-value}$$

Make a decision: Since $\alpha < p\text{-value}$, you cannot reject H_0 .

Conclusion: There is not sufficient evidence to conclude that there is a difference among the mean grades for the sororities.

Exercise 12.4

Four sports teams took a random sample of players regarding their GPAs for the last year. The results are shown in Table 12.4.8

Table 12.4.8 GPAs for four sports teams

Basketball	Baseball	Hockey	Lacrosse
3.6	2.1	4.0	2.0
2.9	2.6	2.0	3.6
2.5	3.9	2.6	3.9
3.3	3.1	3.2	2.7
3.8	3.4	3.2	2.5

Use a significance level of 5%, and determine if there is a difference in GPA among the teams.

Example 12.5

A fourth grade class is studying the environment. One of the assignments is to grow bean plants in different soils. Tommy chose to grow his bean plants in soil found outside his classroom mixed with dryer lint. Tara chose to grow her bean plants in potting soil bought at the local nursery. Nick chose to grow his bean plants in soil from his mother's garden. No chemicals were used on the plants, only water. They were grown inside the classroom next to a large window. Each child grew five plants. At the end of the growing period, each plant was measured, producing the data (in inches) in Table 12.4.9.

Table 12.4.9

Tommy's plants	Tara's plants	Nick's plants
24	25	23
21	31	27
23	23	22
30	20	30
23	28	20

Does it appear that the three media in which the bean plants were grown produce the same mean height? Test at a 3% level of significance.

Answer

This time, we will perform the calculations that lead to the F' statistic. Notice that each group has the same number of plants, so we will use the formula $F' = \frac{n \cdot s_x^2}{s_{pooled}^2}$.

First, calculate the sample mean and sample variance of each group.

Table 12.4.10

	Tommy's plants	Tara's plants	Nick's plants
Sample mean	24.2	25.4	24.4
Sample variance	11.7	18.3	16.3

Next, calculate the variance of the three group means (Calculate the variance of 24.2, 25.4, and 24.4). **Variance of the group means = $0.413 = s_x^2$**

Then $MS_{between} = n s_x^2 = (5)(0.413)$ where $n = 5$ is the sample size (number of plants each child grew).

Calculate the mean of the three sample variances (Calculate the mean of 11.7, 18.3, and 16.3). **Mean of the sample variances = $15.433 = s^2_{pooled}$**

Then $MS_{within} = s^2_{pooled} = 15.433$.

The F statistic (or F ratio) is $F = \frac{MS_{between}}{MS_{within}} = \frac{n s_x^2}{s^2_{pooled}} = \frac{(5)(0.413)}{15.433} = 0.134$

The dfs for the numerator = the number of groups $- 1 = 3 - 1 = 2$.

The dfs for the denominator = the total number of samples $-$ the number of groups $= 15 - 3 = 12$

The distribution for the test is $F_{2,12}$ and the F statistic is $F = 0.134$

The p -value is $P(F > 0.134) = 0.8759$

Decision: Since $\alpha = 0.03$ and the p -value $= 0.8759$, then you cannot reject H_0 . (Why?)

Conclusion: With a 3% level of significance, from the sample data, the evidence is not sufficient to conclude that the mean heights of the bean plants are different.

Notation

The notation for the F distribution is $F \sim F_{df(num), df(denom)}$ where $df(num) = df_{between}$ and $df(denom) = df_{within}$. The mean for the F distribution is $\mu = \frac{df(num)}{df(denom) - 2}$

This page titled [12.4: The F Distribution and the F-Ratio](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.3: The F Distribution and the F-Ratio](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.5: Facts About the F Distribution

Here are some facts about the **F** distribution.

1. The curve is not symmetrical but skewed to the right.
2. There is a different curve for each set of degrees of freedom.
3. The F statistic is greater than or equal to zero.
4. As the degrees of freedom for the numerator and for the denominator get larger, the curve approximates the normal as can be seen in the two figures below. Figure (b) with more degrees of freedom is more closely approaching the normal distribution, but remember that the F cannot ever be less than zero so the distribution does not have a tail that goes to infinity on the left as the normal distribution does.
5. Other uses for the F distribution include comparing two variances and two-way Analysis of Variance. Two-Way Analysis is beyond the scope of this chapter.




Figure 12.5.7

This page titled [12.5: Facts About the F Distribution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.4: Facts About the F Distribution](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.6: Chapter Key Terms

Analysis of Variance

also referred to as ANOVA, is a method of testing whether or not the means of three or more populations are equal. The method is applicable if:

- The test statistic for analysis of variance is the F -ratio.

One-Way ANOVA

a method of testing whether or not the means of three or more populations are equal; the method is applicable if:

- The test statistic for analysis of variance is the F -ratio.

Variance

mean of the squared deviations from the mean; the square of the standard deviation. For a set of data, a deviation can be represented as $x - \bar{x}$ where x is a value of the data and \bar{x} is the sample mean. The sample variance is equal to the sum of the squares of the deviations divided by the difference of the sample size and one.

This page titled [12.6: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.7: Chapter Key Terms](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.7: Chapter Review

12.1 Test of Two Variances

The F test for the equality of two variances rests heavily on the assumption of normal distributions. The test is unreliable if this assumption is not met. If both distributions are normal, then the ratio of the two sample variances is distributed as an F statistic, with numerator and denominator degrees of freedom that are one less than the samples sizes of the corresponding two groups. A **test of two variances** hypothesis test determines if two variances are the same. The distribution for the hypothesis test is the F distribution with two different degrees of freedom.

Assumptions:

1. Analysis of variance extends the comparison of two groups to several, each a level of a categorical variable (factor). Samples from each group are independent, and must be randomly selected from normal populations with equal variances. We test the null hypothesis of equal means of the response in every group versus the alternative hypothesis of one or more group means being different from the others. A one-way ANOVA hypothesis test determines if several population means are equal. The distribution for the test is the F distribution with two different degrees of freedom.

Assumptions:

1. Analysis of variance compares the means of a response variable for several groups. ANOVA compares the variation within each group to the variation of the mean of each group. The ratio of these two is the F statistic from an F distribution with (number of groups – 1) as the numerator degrees of freedom and (number of observations – number of groups) as the denominator degrees of freedom. These statistics are summarized in the ANOVA table.

12.4 Facts About the F Distribution

When the data have unequal group sizes (unbalanced data), then techniques from Figure 12.7.3 need to be used for hand calculations. In the case of balanced data (the groups are the same size) however, simplified calculations based on group means and variances may be used. In practice, of course, software is usually employed in the analysis. As in any analysis, graphs of various sorts should be used in conjunction with numerical techniques. Always look at your data!

This page titled [12.7: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **12.10: Chapter Review** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.8: Chapter Formula Review

12.1 Test of Two Variances

$$H_0 : \frac{\sigma_1^2}{\sigma_2^2} = \delta_0$$

$$H_a : \frac{\sigma_1^2}{\sigma_2^2} \neq \delta_0$$

if $\delta_0 = 1$ then

$$H_0 : \sigma_1^2 = \sigma_2^2$$

$$H_a : \sigma_1^2 \neq \sigma_2^2$$

Test statistic is :

$$F_c = \frac{S_1^2}{S_2^2}$$

12.3 The F Distribution and the F-Ratio

$$SS_{\text{between}} = \sum \left[\frac{(s_j)^2}{n_j} \right] - \frac{(\sum s_j)^2}{n}$$

$$SS_{\text{total}} = \sum x^2 - \frac{(\sum x)^2}{n}$$

$$SS_{\text{within}} = SS_{\text{total}} - SS_{\text{between}}$$

$$df_{\text{between}} = df(\text{num}) = k - 1$$

$$df_{\text{within}} = df(\text{denom}) = n - k$$

$$MS_{\text{between}} = \frac{SS_{\text{between}}}{df_{\text{between}}}$$

$$MS_{\text{within}} = \frac{SS_{\text{within}}}{df_{\text{within}}}$$

$$F = \frac{MS_{\text{between}}}{MS_{\text{within}}}$$

- k = the number of groups
- n_j = the size of the j th group
- s_j = the sum of the values in the j th group
- n = the total number of all values (observations) combined
- x = one value (one observation) from the data
- s_x^2 = the variance of the sample means
- s_{pooled}^2 = the mean of the sample variances (pooled variance)

This page titled [12.8: Chapter Formula Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.5: Chapter Formula Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.9: Chapter Practice

12.1 Test of Two Variances

Use the following information to answer the next two exercises. There are two assumptions that must be true in order to perform an F test of two variances.

1.

Name one assumption that must be true.

2.

What is the other assumption that must be true?

Use the following information to answer the next five exercises. Two coworkers commute from the same building. They are interested in whether or not there is any variation in the time it takes them to drive to work. They each record their times for 20 commutes. The first worker's times have a variance of 12.1. The second worker's times have a variance of 16.9. The first worker thinks that he is more consistent with his commute times. Test the claim at the 10% level. Assume that commute times are normally distributed.

3.

State the null and alternative hypotheses.

4.

What is s_1 in this problem?

5.

What is s_2 in this problem?

6.

What is n ?

7.

What is the F statistic?

8.

What is the critical value?

9.

Is the claim accurate?

Use the following information to answer the next four exercises. Two students are interested in whether or not there is variation in their test scores for math class. There are 15 total math tests they have taken so far. The first student's grades have a standard deviation of 38.1. The second student's grades have a standard deviation of 22.5. The second student thinks his scores are more consistent.

10.

State the null and alternative hypotheses.

11.

What is the F Statistic?

12.

What is the critical value?

13.

At the 5% significance level, do we reject the null hypothesis?

Use the following information to answer the next three exercises. Two cyclists are comparing the variances of their overall paces going uphill. Each cyclist records his or her speeds going up 35 hills. The first cyclist has a variance of 23.8 and the second cyclist

has a variance of 32.1. The cyclists want to see if their variances are the same or different. Assume that commute times are normally distributed.

14.

State the null and alternative hypotheses.

15.

What is the F Statistic?

16.

At the 5% significance level, what can we say about the cyclists' variances?

12.2 One-Way ANOVA

Use the following information to answer the next five exercises. There are five basic assumptions that must be fulfilled in order to perform a one-way ANOVA test. What are they?

17.

Write one assumption.

18.

Write another assumption.

19.

Write a third assumption.

20.

Write a fourth assumption.

12.3 The F Distribution and the F-Ratio

Use the following information to answer the next eight exercises. Groups of men from three different areas of the country are to be tested for mean weight. The entries in Table 12.9.13 are the weights for the different groups.

Group 1	Group 2	Group 3
216	202	170
198	213	165
240	284	182
187	228	197
176	210	201

Table 12.13

21.

What is the Sum of Squares Factor?

22.

What is the Sum of Squares Error?

23.

What is the df for the numerator?

24.

What is the df for the denominator?

25.

What is the Mean Square Factor?

26.

What is the Mean Square Error?

27.

What is the F statistic?

Use the following information to answer the next eight exercises. Girls from four different soccer teams are to be tested for mean goals scored per game. The entries in Table 12.9.14 are the goals per game for the different teams.

Table 12.9.14

Team 1	Team 2	Team 3	Team 4
1	2	0	3
2	3	1	4
0	2	1	4
3	4	0	3
2	4	0	2

28.

What is $SS_{between}$?

29.

What is the df for the numerator?

30.

What is $MS_{between}$?

31.

What is SS_{within} ?

32.

What is the df for the denominator?

33.

What is MS_{within} ?

34.

What is the F statistic?

35.

Judging by the F statistic, do you think it is likely or unlikely that you will reject the null hypothesis?

12.4 Facts About the F Distribution

36.

An F statistic can have what values?

37.

What happens to the curves as the degrees of freedom for the numerator and the denominator get larger?

Use the following information to answer the next seven exercise. Four basketball teams took a random sample of players regarding how high each player can jump (in inches). The results are shown in Table 12.9.15

Table 12.9.15

Team 1	Team 2	Team 3	Team 4	Team 5
36	32	48	38	41
42	35	50	44	39

Team 1	Team 2	Team 3	Team 4	Team 5
51	38	39	46	40

38.

What is the $df(num)$?

39.

What is the $df(denom)$?

40.

What are the Sum of Squares and Mean Squares Factors?

41.

What are the Sum of Squares and Mean Squares Errors?

42.

What is the F statistic?

43.

What is the p -value?

44.

At the 5% significance level, is there a difference in the mean jump heights among the teams?

Use the following information to answer the next seven exercises. A video game developer is testing a new game on three different groups. Each group represents a different target market for the game. The developer collects scores from a random sample from each group. The results are shown in Table 12.9.16

Group A	Group B	Group C
101	151	101
108	149	109
98	160	198
107	112	186
111	126	160

Table 12.16

45.

What is the $df(num)$?

46.

What is the $df(denom)$?

47.

What are the $SS_{between}$ and $MS_{between}$?

48.

What are the SS_{within} and MS_{within} ?

49.

What is the F Statistic?

50.

What is the p -value?

51.

At the 10% significance level, are the scores among the different groups different?

Use the following information to answer the next three exercises. Suppose a group is interested in determining whether teenagers obtain their drivers licenses at approximately the same average age across the country. Suppose that the following data are randomly collected from five teenagers in each region of the country. The numbers represent the age at which teenagers obtained their drivers licenses.

Table 12.9.17

	Northeast	South	West	Central	East
	16.3	16.9	16.4	16.2	17.1
	16.1	16.5	16.5	16.6	17.2
	16.4	16.4	16.6	16.5	16.6
	16.5	16.2	16.1	16.4	16.8
$\bar{x} =$	_____	_____	_____	_____	_____
$s^2 =$	_____	_____	_____	_____	_____

Enter the data into your calculator or computer.

52.

p -value = _____

State the decisions and conclusions (in complete sentences) for the following preconceived levels of α .

53.

$\alpha = 0.05$

a. Decision: _____

b. Conclusion: _____

54.

$\alpha = 0.01$

a. Decision: _____

b. Conclusion: _____

This page titled [12.9: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **12.8: Chapter Practice** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.10: Chapter Homework

12.1 Test of Two Variances

55.

Three students, Linda, Tuan, and Javier, are given five laboratory rats each for a nutritional experiment. Each rat's weight is recorded in grams. Linda feeds her rats Formula A, Tuan feeds his rats Formula B, and Javier feeds his rats Formula C. At the end of a specified time period, each rat is weighed again and the net gain in grams is recorded.

Table 12.10.18

Linda's rats	Tuan's rats	Javier's rats
43.5	47.0	51.2
39.4	40.5	40.9
41.3	38.9	37.9
46.0	46.3	45.0
38.2	44.2	48.6

Determine whether or not the variance in weight gain is statistically the same among Javier's and Linda's rats. Test at a significance level of 10%.

56.

A grassroots group opposed to a proposed increase in the gas tax claimed that the increase would hurt working-class people the most, since they commute the farthest to work. Suppose that the group randomly surveyed 24 individuals and asked them their daily one-way commuting mileage. The results are as follows.

Table 12.10.19

Working-class	Professional (middle incomes)	Professional (wealthy)
17.8	16.5	8.5
26.7	17.4	6.3
49.4	22.0	4.6
9.4	7.4	12.6
65.4	9.4	11.0
47.1	2.1	28.6
19.5	6.4	15.4
51.2	13.9	9.3

Determine whether or not the variance in mileage driven is statistically the same among the working class and professional (middle income) groups. Use a 5% significance level.

Use the following information to answer the next two exercises. The following table lists the number of pages in four different types of magazines.

Table 12.10.20

Home decorating	News	Health	Computer
172	87	82	104
286	94	153	136
163	123	87	98
205	106	103	207

Home decorating	News	Health	Computer
197	101	96	146

57.

Which two magazine types do you think have the same variance in length?

58.

Which two magazine types do you think have different variances in length?

59.

Is the variance for the amount of money, in dollars, that shoppers spend on Saturdays at the mall the same as the variance for the amount of money that shoppers spend on Sundays at the mall? Suppose that the Table 12.10.21 shows the results of a study.

Saturday	Sunday	Saturday	Sunday
75	44	62	137
18	58	0	82
150	61	124	39
94	19	50	127
62	99	31	141
73	60	118	73
	89		

Table 12.21

60.

Are the variances for incomes on the East Coast and the West Coast the same? Suppose that Table 12.10.22 shows the results of a study. Income is shown in thousands of dollars. Assume that both distributions are normal. Use a level of significance of 0.05.

Table 12.10.22

East	West
38	71
47	126
30	42
82	51
75	44
52	90
115	88
67	

61.

Thirty men in college were taught a method of finger tapping. They were randomly assigned to three groups of ten, with each receiving one of three doses of caffeine: 0 mg, 100 mg, 200 mg. This is approximately the amount in no, one, or two cups of coffee. Two hours after ingesting the caffeine, the men had the rate of finger tapping per minute recorded. The experiment was double blind, so neither the recorders nor the students knew which group they were in. Does caffeine affect the rate of tapping, and if so how?

Here are the data:

0 mg	100 mg	200 mg	0 mg	100 mg	200 mg
------	--------	--------	------	--------	--------

0 mg	100 mg	200 mg	0 mg	100 mg	200 mg
242	248	246	245	246	248
244	245	250	248	247	252
247	248	248	248	250	250
242	247	246	244	246	248
246	243	245	242	244	250

Table 12.23

62.

King Manuel I, Komnenus ruled the Byzantine Empire from Constantinople (Istanbul) during the years 1145 to 1180 A.D. The empire was very powerful during his reign, but declined significantly afterwards. Coins minted during his era were found in Cyprus, an island in the eastern Mediterranean Sea. Nine coins were from his first coinage, seven from the second, four from the third, and seven from a fourth. These spanned most of his reign. We have data on the silver content of the coins:

Table 12.10.24

First coinage	Second coinage	Third coinage	Fourth coinage
5.9	6.9	4.9	5.3
6.8	9.0	5.5	5.6
6.4	6.6	4.6	5.5
7.0	8.1	4.5	5.1
6.6	9.3		6.2
7.7	9.2		5.8
7.2	8.6		5.8
6.9			
6.2			

Did the silver content of the coins change over the course of Manuel's reign?

Here are the means and variances of each coinage. The data are unbalanced.

	First	Second	Third	Fourth
Mean	6.7444	8.2429	4.875	5.6143
Variance	0.2953	1.2095	0.2025	0.1314

Table 12.25

63.

The American League and the National League of Major League Baseball are each divided into three divisions: East, Central, and West. Many years, fans talk about some divisions being stronger (having better teams) than other divisions. This may have consequences for the postseason. For instance, in 2012 Tampa Bay won 90 games and did not play in the postseason, while Detroit won only 88 and did play in the postseason. This may have been an oddity, but is there good evidence that in the 2012 season, the American League divisions were significantly different in overall records? Use the following data to test whether the mean number of wins per team in the three American League divisions were the same or not. Note that the data are not balanced, as two divisions had five teams, while one had only four.

Table 12.10.26

Division	Team	Wins
East	NY Yankees	95

Division	Team	Wins
East	Baltimore	93
East	Tampa Bay	90
East	Toronto	73
East	Boston	69

Table 12.10.27

Division	Team	Wins
Central	Detroit	88
Central	Chicago Sox	85
Central	Kansas City	72
Central	Cleveland	68
Central	Minnesota	66

Table 12.10.28

Division	Team	Wins
West	Oakland	94
West	Texas	93
West	LA Angels	89
West	Seattle	75

12.2 One-Way ANOVA

64.

Three different traffic routes are tested for mean driving time. The entries in the Table 12.10.29 are the driving times in minutes on the three different routes.

Table 12.10.29

Route 1	Route 2	Route 3
30	27	16
32	29	41
27	28	22
35	36	31

State $SS_{between}$, SS_{within} , and the F statistic.

65.

Suppose a group is interested in determining whether teenagers obtain their drivers licenses at approximately the same average age across the country. Suppose that the following data are randomly collected from five teenagers in each region of the country. The numbers represent the age at which teenagers obtained their drivers licenses.

Table 12.10.30

	Northeast	South	West	Central	East
	16.3	16.9	16.4	16.2	17.1
	16.1	16.5	16.5	16.6	17.2
	16.4	16.4	16.6	16.5	16.6
	16.5	16.2	16.1	16.4	16.8

	Northeast	South	West	Central	East
$\bar{x} =$	_____	_____	_____	_____	_____
$s^2 =$	_____	_____	_____	_____	_____

State the hypotheses.

H_0 : _____

H_a : _____

12.3 The F Distribution and the F-Ratio

Use the following information to answer the next three exercises. Suppose a group is interested in determining whether teenagers obtain their drivers licenses at approximately the same average age across the country. Suppose that the following data are randomly collected from five teenagers in each region of the country. The numbers represent the age at which teenagers obtained their drivers licenses.

Table 12.10.31

	Northeast	South	West	Central	East
	16.3	16.9	16.4	16.2	17.1
	16.1	16.5	16.5	16.6	17.2
	16.4	16.4	16.6	16.5	16.6
	16.5	16.2	16.1	16.4	16.8
$\bar{x} =$	_____	_____	_____	_____	_____
$s^2 =$	_____	_____	_____	_____	_____

$H_0 : \mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5$

H_a : At least any two of the group means $\mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5$ are not equal.

66.

degrees of freedom – numerator: $df(num) =$ _____

67.

degrees of freedom – denominator: $df(denom) =$ _____

68.

F statistic = _____

12.4 Facts About the F Distribution

69.

Three students, Linda, Tuan, and Javier, are given five laboratory rats each for a nutritional experiment. Each rat's weight is recorded in grams. Linda feeds her rats Formula A, Tuan feeds his rats Formula B, and Javier feeds his rats Formula C. At the end of a specified time period, each rat is weighed again, and the net gain in grams is recorded. Using a significance level of 10%, test the hypothesis that the three formulas produce the same mean weight gain.

Table 12.10.32 Weights of Student Lab Rats

Linda's rats	Tuan's rats	Javier's rats
43.5	47.0	51.2
39.4	40.5	40.9
41.3	38.9	37.9
46.0	46.3	45.0
38.2	44.2	48.6

70.

A grassroots group opposed to a proposed increase in the gas tax claimed that the increase would hurt working-class people the most, since they commute the farthest to work. Suppose that the group randomly surveyed 24 individuals and asked them their daily one-way commuting mileage. The results are in Table 12.10.33 Using a 5% significance level, test the hypothesis that the three mean commuting mileages are the same.

Working-class	Professional (middle incomes)	Professional (wealthy)
17.8	16.5	8.5
26.7	17.4	6.3
49.4	22.0	4.6
9.4	7.4	12.6
65.4	9.4	11.0
47.1	2.1	28.6
19.5	6.4	15.4
51.2	13.9	9.3

Table 12.33

Use the following information to answer the next two exercises. Table 12.10.34 lists the number of pages in four different types of magazines.

Table 12.10.34

Home decorating	News	Health	Computer
172	87	82	104
286	94	153	136
163	123	87	98
205	106	103	207
197	101	96	146

71.

Using a significance level of 5%, test the hypothesis that the four magazine types have the same mean length.

72.

Eliminate one magazine type that you now feel has a mean length different from the others. Redo the hypothesis test, testing that the remaining three means are statistically the same. Use a new solution sheet. Based on this test, are the mean lengths for the remaining three magazines statistically the same?

73.

A researcher wants to know if the mean times (in minutes) that people watch their favorite news station are the same. Suppose that Table 12.10.35 shows the results of a study.

Table 12.10.35

CNN	FOX	Local
45	15	72
12	43	37
18	68	56
38	50	60
23	31	51

CNN	FOX	Local
35	22	

Assume that all distributions are normal, the four population standard deviations are approximately the same, and the data were collected independently and randomly. Use a level of significance of 0.05.

74.

Are the means for the final exams the same for all statistics class delivery types? Table 12.10.36 shows the scores on final exams from several randomly selected classes that used the different delivery types.

Table 12.10.36

Online	Hybrid	Face-to-Face
72	83	80
84	73	78
77	84	84
80	81	81
81		86
		79
		82

Assume that all distributions are normal, the four population standard deviations are approximately the same, and the data were collected independently and randomly. Use a level of significance of 0.05.

75.

Are the mean number of times a month a person eats out the same for whites, blacks, Hispanics and Asians? Suppose that Table 12.10.38 shows the results of a study.

Table 12.10.38

Powder	Machine Made	Hard Packed
1,210	2,107	2,846
1,080	1,149	1,638
1,537	862	2,019
941	1,870	1,178
	1,528	2,233
	1,382	

Assume that all distributions are normal, the four population standard deviations are approximately the same, and the data were collected independently and randomly. Use a level of significance of 0.05.

77.

Sanjay made identical paper airplanes out of three different weights of paper, light, medium and heavy. He made four airplanes from each of the weights, and launched them himself across the room. Here are the distances (in meters) that his planes flew.

Paper type/Trial	Trial 1	Trial 2	Trial 3	Trial 4
Heavy	5.1 meters	3.1 meters	4.7 meters	5.3 meters
Medium	4 meters	3.5 meters	4.5 meters	6.1 meters
Light	3.1 meters	3.3 meters	2.1 meters	1.9 meters

Table 12.39




Figure 12.10.8

1. An experiment was conducted on the number of eggs (fecundity) laid by female fruit flies. There are three groups of flies. One group was bred to be resistant to DDT (the RS group). Another was bred to be especially susceptible to DDT (SS). Finally there was a control line of non-selected or typical fruitflies (NS). Here are the data:

Table 12.10.40

RS	SS	NS	RS	SS	NS
12.8	38.4	35.4	22.4	23.1	22.6
21.6	32.9	27.4	27.5	29.4	40.4
14.8	48.5	19.3	20.3	16	34.4
23.1	20.9	41.8	38.7	20.1	30.4
34.6	11.6	20.3	26.4	23.3	14.9
19.7	22.3	37.6	23.7	22.9	51.8
22.6	30.2	36.9	26.1	22.5	33.8
29.6	33.4	37.3	29.5	15.1	37.9
16.4	26.7	28.2	38.6	31	29.5
20.3	39	23.4	44.4	16.9	42.4
29.3	12.8	33.7	23.2	16.1	36.6
14.9	14.6	29.2	23.6	10.8	47.4
27.3	12.2	41.7			

Here is a chart of the three groups:




Figure 12.10.9

79.

The data shown is the recorded body temperatures of 130 subjects as estimated from available histograms.

Traditionally we are taught that the normal human body temperature is 98.6 F. This is not quite correct for everyone. Are the mean temperatures among the four groups different?

Calculate 95% confidence intervals for the mean body temperature in each group and comment about the confidence intervals.

Table 12.10.41

FL	FH	ML	MH	FL	FH	ML	MH
96.4	96.8	96.3	96.9	98.4	98.6	98.1	98.6
96.7	97.7	96.7	97	98.7	98.6	98.1	98.6
97.2	97.8	97.1	97.1	98.7	98.6	98.2	98.7
97.2	97.9	97.2	97.1	98.7	98.7	98.2	98.8
97.4	98	97.3	97.4	98.7	98.7	98.2	98.8
97.6	98	97.4	97.5	98.8	98.8	98.2	98.8
97.7	98	97.4	97.6	98.8	98.8	98.3	98.9
97.8	98	97.4	97.7	98.8	98.8	98.4	99
97.8	98.1	97.5	97.8	98.8	98.9	98.4	99
97.9	98.3	97.6	97.9	99.2	99	98.5	99
97.9	98.3	97.6	98	99.3	99	98.5	99.2
98	98.3	97.8	98		99.1	98.6	99.5
98.2	98.4	97.8	98		99.1	98.6	
98.2	98.4	97.8	98.3		99.2	98.7	
98.2	98.4	97.9	98.4		99.4	99.1	
98.2	98.4	98	98.4		99.9	99.3	
98.2	98.5	98	98.6		100	99.4	
98.2	98.6	98	98.6		100.8		

This page titled [12.10: Chapter Homework](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.6: Chapter Homework](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.11: Chapter Reference

12.1 Test of Two Variances

“MLB Vs. Division Standings – 2012.” Available online at http://espn.go.com/mlb/standings/_y...ion/order/true.

12.3 The F Distribution and the F-Ratio

Tomato Data, Marist College School of Science (unpublished student research)

12.4 Facts About the F Distribution

Data from a fourth grade classroom in 1994 in a private K – 12 school in San Jose, CA.

Hand, D.J., F. Daly, A.D. Lunn, K.J. McConway, and E. Ostrowski. *A Handbook of Small Datasets: Data for Fruitfly Fecundity*. London: Chapman & Hall, 1994.

Hand, D.J., F. Daly, A.D. Lunn, K.J. McConway, and E. Ostrowski. *A Handbook of Small Datasets*. London: Chapman & Hall, 1994, pg. 50.

Hand, D.J., F. Daly, A.D. Lunn, K.J. McConway, and E. Ostrowski. *A Handbook of Small Datasets*. London: Chapman & Hall, 1994, pg. 118.

“MLB Standings – 2012.” Available online at http://espn.go.com/mlb/standings/_year/2012.

Mackowiak, P. A., Wasserman, S. S., and Levine, M. M. (1992), "A Critical Appraisal of 98.6 Degrees F, the Upper Limit of the Normal Body Temperature, and Other Legacies of Carl Reinhold August Wunderlich," *Journal of the American Medical Association*, 268, 1578-1580.

This page titled [12.11: Chapter Reference](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.9: Chapter Reference](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

12.12: Chapter Solution (Practice + Homework)




Figure 12.12.10

While there are differences in spread, it is not unreasonable to use ANOVA techniques. Here is the completed ANOVA table:

Table 12.12.42

Figure 12.12.11

Table 12.12.43

Source of variation	Sum of squares (SS)	Degrees of freedom (df)	Mean square (MS)	F
Factor (Between)	37.748	$4 - 1 = 3$	12.5825	26.272
Error (Within)	11.015	$27 - 4 = 23$	0.4789	
Total	48.763	$27 - 1 = 26$		

$$P(F > 1.5521) = 0.2548$$

Since the p-value is so large, there is not good evidence against the null hypothesis of equal means. We cannot reject the null hypothesis. Thus, for 2012, there is not any have any good evidence of a significant difference in mean number of wins between the divisions of the American League.

64.

$$SS_{between} = 26$$

$$SS_{within} = 441$$

$$F = 0.2653$$

67.

$$df(\text{denom}) = 15$$

69.

1. 72.

1. 74.

1. 76.

1. 78.

The data appear normally distributed from the chart and of similar spread. There do not appear to be any serious outliers, so we may proceed with our ANOVA calculations, to see if we have good evidence of a difference between the three groups.

Define μ_1, μ_2, μ_3 , as the population mean number of eggs laid by the three groups of fruit flies.

$$F \text{ statistic} = 8.6657;$$

$$p\text{-value} = 0.0004$$




Figure 12.12.12

Decision: Since the p -value is less than the level of significance of 0.01, we reject the null hypothesis.

Conclusion: We have good evidence that the average number of eggs laid during the first 14 days of life for these three strains of fruitflies are different.

Interestingly, if you perform a two sample t -test to compare the RS and NS groups they are significantly different ($p = 0.0013$). Similarly, SS and NS are significantly different ($p = 0.0006$). However, the two selected groups, RS and SS are not significantly different ($p = 0.5176$). Thus we appear to have good evidence that selection either for resistance or for susceptibility involves a reduced rate of egg production (for these specific strains) as compared to flies that were not selected for resistance or susceptibility to DDT. Here, genetic selection has apparently involved a loss of fecundity.

This page titled [12.12: Chapter Solution \(Practice + Homework\)](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [12.11: Chapter Solution \(Practice + Homework\)](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

CHAPTER OVERVIEW

13: Linear Regression and Correlation

[13.1: Introduction to Linear Regression and Correlation](#)

[13.2: The Correlation Coefficient \$r\$](#)

[13.3: Testing the Significance of the Correlation Coefficient](#)

[13.4: Linear Equations](#)

[13.5: The Regression Equation](#)

[13.6: Interpretation of Regression Coefficients- Elasticity and Logarithmic Transformation](#)

[13.7: Predicting with a Regression Equation](#)

[13.8: How to Use Microsoft Excel® for Regression Analysis](#)

[13.9: Chapter Key Terms](#)

[13.10: Chapter Review](#)

[13.11: Chapter Practice](#)

[13.12: Chapter Solution](#)

This page titled [13: Linear Regression and Correlation](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

13.1: Introduction to Linear Regression and Correlation



Figure 13.1 Linear regression and correlation can help you determine if an auto mechanic's salary is related to his work experience. (credit: Joshua Rothhaas)

Professionals often want to know how two or more numeric variables are related. For example, is there a relationship between the grade on the second math exam a student takes and the grade on the final exam? If there is a relationship, what is the relationship and how strong is it?

In another example, your income may be determined by your education, your profession, your years of experience, and your ability, or your gender or color. The amount you pay a repair person for labor is often determined by an initial amount plus an hourly fee.

These examples may or may not be tied to a model, meaning that some theory suggested that a relationship exists. This link between a cause and an effect, often referred to as a model, is the foundation of the scientific method and is the core of how we determine what we believe about how the world works. Beginning with a theory and developing a model of the theoretical relationship should result in a prediction, what we have called a hypothesis earlier. Now the hypothesis concerns a full set of relationships. As an example, in Economics the model of consumer choice is based upon assumptions concerning human behavior: a desire to maximize something called utility, knowledge about the benefits of one product over another, likes and dislikes, referred to generally as preferences, and so on. These combined to give us the demand curve. From that we have the prediction that as prices rise the quantity demanded will fall. Economics has models concerning the relationship between what prices are charged for goods and the market structure in which the firm operates, monopoly verse competition, for example. Models for who would be most likely to be chosen for an on-the-job training position, the impacts of Federal Reserve policy changes and the growth of the economy and on and on.

Models are not unique to Economics, even within the social sciences. In political science, for example, there are models that predict behavior of bureaucrats to various changes in circumstances based upon assumptions of the goals of the bureaucrats. There are models of political behavior dealing with strategic decision making both for international relations and domestic politics.

The so-called hard sciences are, of course, the source of the scientific method as they tried through the centuries to explain the confusing world around us. Some early models today make us laugh; spontaneous generation of life for example. These early models are seen today as not much more than the foundational myths we developed to help us bring some sense of order to what seemed chaos.

The foundation of all model building is the perhaps the arrogant statement that we know what caused the result we see. This is embodied in the simple mathematical statement of the functional form that $y = f(x)$. The response, Y , is caused by the stimulus, X . Every model will eventually come to this final place and it will be here that the theory will live or die. Will the data support this hypothesis? If so then fine, we shall believe this version of the world until a better theory comes to replace it. This is the process by which we moved from flat earth to round earth, from earth-center solar system to sun-center solar system, and on and on.

The scientific method does not confirm a theory for all time: it does not prove "truth". All theories are subject to review and may be overturned. These are lessons we learned as we first developed the concept of the hypothesis test earlier in this book. Here, as we begin this section, these concepts deserve review because the tool we will develop here is the cornerstone of the scientific method and the stakes are higher. Full theories will rise or fall because of this statistical tool; regression and the more advanced versions call econometrics.

In this chapter we will begin with correlation, the investigation of relationships among variables that may or may not be founded on a cause and effect model. The variables simply move in the same, or opposite, direction. That is to say, they do not move randomly. Correlation provides a measure of the degree to which this is true. From there we develop a tool to measure cause and effect

relationships; regression analysis. We will be able to formulate models and tests to determine if they are statistically sound. If they are found to be so, then we can use them to make predictions: if as a matter of policy we changed the value of this variable what would happen to this other variable? If we imposed a gasoline tax of 50 cents per gallon how would that effect the carbon emissions, sales of Hummers/Hybrids, use of mass transit, etc.? The ability to provide answers to these types of questions is the value of regression as both a tool to help us understand our world and to make thoughtful policy decisions.

This page titled [13.1: Introduction to Linear Regression and Correlation](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [13.0: Introduction to Linear Regression and Correlation](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.2: The Correlation Coefficient r

As we begin this section we note that the type of data we will be working with has changed. Perhaps unnoticed, all the data we have been using is for a single variable. It may be from two samples, but it is still a univariate variable. The type of data described in the examples above and for any model of cause and effect is **bivariate** data — "bi" for two variables. In reality, statisticians use **multivariate** data, meaning many variables.

For our work we can classify data into three broad categories, time series data, cross-section data, and panel data. We met the first two very early on. Time series data measures a single unit of observation; say a person, or a company or a country, as time passes. What are measured will be at least two characteristics, say the person's income, the quantity of a particular good they buy and the price they paid. This would be three pieces of information in one time period, say 1985. If we followed that person across time we would have those same pieces of information for 1985, 1986, 1987, etc. This would constitute a times series data set. If we did this for 10 years we would have 30 pieces of information concerning this person's consumption habits of this good for the past decade and we would know their income and the price they paid.

A second type of data set is for cross-section data. Here the variation is not across time for a single unit of observation, but across units of observation during one point in time. For a particular period of time we would gather the price paid, amount purchased, and income of many individual people.

A third type of data set is panel data. Here a panel of units of observation is followed across time. If we take our example from above we might follow 500 people, the unit of observation, through time, ten years, and observe their income, price paid and quantity of the good purchased. If we had 500 people and data for ten years for price, income and quantity purchased we would have 15,000 pieces of information. These types of data sets are very expensive to construct and maintain. They do, however, provide a tremendous amount of information that can be used to answer very important questions. As an example, what is the effect on the labor force participation rate of women as their family of origin, mother and father, age? Or are there differential effects on health outcomes depending upon the age at which a person started smoking? Only panel data can give answers to these and related questions because we must follow multiple people across time. The work we do here however will not be fully appropriate for data sets such as these.

Beginning with a set of data with two independent variables we ask the question: are these related? One way to visually answer this question is to create a scatter plot of the data. We could not do that before when we were doing descriptive statistics because those data were univariate. Now we have bivariate data so we can plot in two dimensions. Three dimensions are possible on a flat piece of paper, but become very hard to fully conceptualize. Of course, more than three dimensions cannot be graphed although the relationships can be measured mathematically.

To provide mathematical precision to the measurement of what we see we use the correlation coefficient. The correlation tells us something about the co-movement of two variables, but **nothing** about why this movement occurred. Formally, correlation analysis assumes that both variables being analyzed are **independent** variables. This means that neither one causes the movement in the other. Further, it means that neither variable is dependent on the other, or for that matter, on any other variable. Even with these limitations, correlation analysis can yield some interesting results.

The correlation coefficient, ρ (pronounced rho), is the mathematical statistic for a population that provides us with a measurement of the strength of a linear relationship between the two variables. For a sample of data, the statistic, r , developed by Karl Pearson in the early 1900s, is an estimate of the population correlation and is defined mathematically as:

$$r = \frac{\frac{1}{n-1} \sum (X_{1i} - \bar{X}_1) (X_{2i} - \bar{X}_2)}{s_{x_1} s_{x_2}}$$

OR

$$r = \frac{\sum X_{1i} X_{2i} - n \bar{X}_1 \bar{X}_2}{\sqrt{(\sum X_{1i}^2 - n \bar{X}_1^2) (\sum X_{2i}^2 - n \bar{X}_2^2)}}$$

where s_{x_1} and s_{x_2} are the standard deviations of the two independent variables X_1 and X_2 , \bar{X}_1 and \bar{X}_2 are the sample means of the two variables, and X_{1i} and X_{2i} are the individual observations of X_1 and X_2 . The correlation coefficient r ranges in value from -1 to 1. The second equivalent formula is often used because it may be computationally easier. As scary as these formulas

look they are really just the ratio of the covariance between the two variables and the product of their two standard deviations. That is to say, it is a measure of relative variances.

In practice all correlation and regression analysis will be provided through computer software designed for these purposes. Anything more than perhaps one-half a dozen observations creates immense computational problems. It was because of this fact that correlation, and even more so, regression, were not widely used research tools until after the advent of “computing machines”. Now the computing power required to analyze data using regression packages is deemed almost trivial by comparison to just a decade ago.

To visualize any **linear** relationship that may exist review the plot of a scatter diagrams of the standardized data. Figure 13.2.2 presents several scatter diagrams and the calculated value of r . In panels (a) and (b) notice that the data generally trend together, (a) upward and (b) downward. Panel (a) is an example of a positive correlation and panel (b) is an example of a negative correlation, or relationship. The sign of the correlation coefficient tells us if the relationship is a positive or negative (inverse) one. If all the values of X_1 and X_2 are on a straight line the correlation coefficient will be either 1 or -1 depending on whether the line has a positive or negative slope and the closer to one or negative one the stronger the relationship between the two variables. **BUT ALWAYS REMEMBER THAT THE CORRELATION COEFFICIENT DOES NOT TELL US THE SLOPE.**




Figure 13.2.2

Remember, all the correlation coefficient tells us is whether or not the data are linearly related. In panel (d) the variables obviously have some type of very specific relationship to each other, but the correlation coefficient is zero, indicating no **linear** relationship exists.

If you suspect a linear relationship between X_1 and X_2 then r can measure how strong the linear relationship is.

What the VALUE of r tells us:

- What the SIGN of r tells us
 - "correlation does not imply causation."

This page titled 13.2: The Correlation Coefficient r is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by OpenStax via source content that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- 13.1: The Correlation Coefficient r by OpenStax is licensed CC BY 4.0. Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.3: Testing the Significance of the Correlation Coefficient

The correlation coefficient, r , tells us about the strength and direction of the linear relationship between X_1 and X_2 .

The sample data are used to compute r , the correlation coefficient for the sample. If we had data for the entire population, we could find the population correlation coefficient. But because we have only sample data, we cannot calculate the population correlation coefficient. The sample correlation coefficient, r , is our estimate of the unknown population correlation coefficient.

- The hypothesis test lets us decide whether the value of the population correlation coefficient ρ is "close to zero" or "significantly different from zero". We decide this based on the sample correlation coefficient r and the sample size n .

If the test concludes that the correlation coefficient is significantly different from zero, we say that the correlation coefficient is "significant."

- What the Hypotheses Mean in Words
 - **Drawing a Conclusion** There are two methods of making the decision concerning the hypothesis. The test statistic to test this hypothesis is:

$$t_c = \frac{r}{\sqrt{(1-r^2)/(n-2)}}$$
$$t_c = \frac{r\sqrt{n-2}}{\sqrt{1-r^2}}$$

Where the second formula is an equivalent form of the test statistic, n is the sample size and the degrees of freedom are $n - 2$. This is a t -statistic and operates in the same way as other t tests. Calculate the t -value and compare that with the critical value from the t -table at the appropriate degrees of freedom and the level of confidence you wish to maintain. If the calculated value is in the tail then cannot accept the null hypothesis that there is no linear relationship between these two independent random variables. If the calculated t -value is NOT in the tailed then cannot reject the null hypothesis that there is no linear relationship between the two variables.

A quick shorthand way to test correlations is the relationship between the sample size and the correlation. If:

$$|r| \geq \frac{2}{\sqrt{n}}$$

then this implies that the correlation between the two variables demonstrates that a linear relationship exists and is statistically significant at approximately the 0.05 level of significance. As the formula indicates, there is an inverse relationship between the sample size and the required correlation for significance of a linear relationship. With only 10 observations, the required correlation for significance is 0.6325, for 30 observations the required correlation for significance decreases to 0.3651 and at 100 observations the required level is only 0.2000.

Correlations may be helpful in visualizing the data, but are not appropriately used to "explain" a relationship between two variables. Perhaps no single statistic is more misused than the correlation coefficient. Citing correlations between health conditions and everything from place of residence to eye color have the effect of implying a cause and effect relationship. This simply cannot be accomplished with a correlation coefficient. The correlation coefficient is, of course, innocent of this misinterpretation. It is the duty of the analyst to use a statistic that is designed to test for cause and effect relationships and report only those results if they are intending to make such a claim. The problem is that passing this more rigorous test is difficult so lazy and/or unscrupulous "researchers" fall back on correlations when they cannot make their case legitimately.

This page titled [13.3: Testing the Significance of the Correlation Coefficient](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [13.2: Testing the Significance of the Correlation Coefficient](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.4: Linear Equations

Linear regression for two variables is based on a linear equation with one independent variable. The equation has the form:

$$y = a + bx$$

where a and b are constant numbers.

The variable x is **the independent variable**, and y is **the dependent variable**. Another way to think about this equation is a statement of cause and effect. The X variable is the cause and the Y variable is the hypothesized effect. Typically, you choose a value to substitute for the independent variable and then solve for the dependent variable.

Example 13.4.1

The following examples are linear equations.

$$y = 3 + 2x$$

$$y = -0.01 + 1.2x$$

The graph of a linear equation of the form $y = a + bx$ is a **straight line**. Any line that is not vertical can be described by this equation

Example 13.4.2

Graph the equation $y = -1 + 2x$.




Figure 13.4.3

Exercise 13.4.2

Is the following an example of a linear equation? Why or why not?




Figure 13.4.4

Example 13.4.3

Aaron's Word Processing Service (AWPS) does word processing. The rate for services is \$32 per hour plus a \$31.50 one-time charge. The total cost to a customer depends on the number of hours it takes to complete the job.

Find the equation that expresses the **total cost** in terms of the **number of hours** required to complete the job.

Answer

Solution 13.3

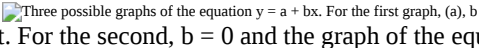
Let x = the number of hours it takes to get the job done.

Let y = the total cost to the customer.

The \$31.50 is a fixed cost. If it takes x hours to complete the job, then $(32)(x)$ is the cost of the word processing only. The total cost is: $y = 31.50 + 32x$

Slope and Y-Intercept of a Linear Equation

For the linear equation $y = a + bx$, b = slope and a = y -intercept. From algebra recall that the slope is a number that describes the steepness of a line, and the y -intercept is the y coordinate of the point $(0, a)$ where the line crosses the y -axis. From calculus the slope is the first derivative of the function. For a linear function the slope is $dy/dx = b$ where we can read the mathematical expression as "the change in y (dy) that results from a change in x (dx) = $b * dx$ ".

 Three possible graphs of the equation $y = a + bx$. For the first graph, (a), b

0 and so the line slopes upward to the right. For the second, $b = 0$ and the graph of the equation is a horizontal line. In the third graph, (c), $b < 0$ and the line slopes downward to the right." data-media-type="image/jpg" style="width: 856px; height: 203px; width="856px" height="203px" src="/@api/deki/files/8169/7096285f5e75a2c46961f54cfccefea4e79baef">

Figure 13.4.5 Three possible graphs of $y = a + bx$. (a) If $b > 0$, the line slopes upward to the right. (b) If $b = 0$, the line is horizontal. (c) If $b < 0$, the line slopes downward to the right.

Example 13.4.4

Svetlana tutors to make extra money for college. For each tutoring session, she charges a one-time fee of \$25 plus \$15 per hour of tutoring. A linear equation that expresses the total amount of money Svetlana earns for each session she tutors is $y = 25 + 15x$.

What are the independent and dependent variables? What is the y -intercept and what is the slope? Interpret them using complete sentences.

Answer

Solution 13.4

The independent variable (x) is the number of hours Svetlana tutors each session. The dependent variable (y) is the amount, in dollars, Svetlana earns for each session.

The y -intercept is 25 ($a = 25$). At the start of the tutoring session, Svetlana charges a one-time fee of \$25 (this is when $x = 0$). The slope is 15 ($b = 15$). For each session, Svetlana earns \$15 for each hour she tutors.

This page titled [13.4: Linear Equations](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [13.3: Linear Equations](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.5: The Regression Equation

Regression analysis is a statistical technique that can test the hypothesis that a variable is dependent upon one or more other variables. Further, regression analysis can provide an estimate of the magnitude of the impact of a change in one variable on another. This last feature, of course, is all important in predicting future values.

Regression analysis is based upon a functional relationship among variables and further, assumes that the relationship is linear. This linearity assumption is required because, for the most part, the theoretical statistical properties of non-linear estimation are not well worked out yet by the mathematicians and econometricians. This presents us with some difficulties in economic analysis because many of our theoretical models are nonlinear. The marginal cost curve, for example, is decidedly nonlinear as is the total cost function, if we are to believe in the effect of specialization of labor and the Law of Diminishing Marginal Product. There are techniques for overcoming some of these difficulties, exponential and logarithmic transformation of the data for example, but at the outset we must recognize that standard ordinary least squares (OLS) regression analysis will always use a linear function to estimate what might be a nonlinear relationship.

The general linear regression model can be stated by the equation:

$$y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \cdots + \beta_k X_{ki} + \varepsilon_i$$

where β_0 is the intercept, β_i 's are the slope between Y and the appropriate X_i , and ε (pronounced epsilon), is the error term that captures errors in measurement of Y and the effect on Y of any variables missing from the equation that would contribute to explaining variations in Y . This equation is the theoretical population equation and therefore uses Greek letters. The equation we will estimate will have the Roman equivalent symbols. This is parallel to how we kept track of the population parameters and sample parameters before. The symbol for the population mean was μ and for the sample mean \bar{X} and for the population standard deviation was σ and for the sample standard deviation was s . The equation that will be estimated with a sample of data for two independent variables will thus be:

$$y_i = b_0 + b_1 x_{1i} + b_2 x_{2i} + e_i$$

As with our earlier work with probability distributions, this model works only if certain assumptions hold. These are that the Y is normally distributed, the errors are also normally distributed with a mean of zero and a constant standard deviation, and that the error terms are independent of the size of X and independent of each other.

Assumptions of the Ordinary Least Squares Regression Model

Each of these assumptions needs a bit more explanation. If one of these assumptions fails to be true, then it will have an effect on the quality of the estimates. Some of the failures of these assumptions can be fixed while others result in estimates that quite simply provide no insight into the questions the model is trying to answer or worse, give biased estimates.

1. The independent variables, x_i , are all measured without error, and are fixed numbers that are independent of the error term. This assumption is saying in effect that Y is deterministic, the result of a fixed component “ X ” and a random error component “ ε .”
2. The error term is a random variable with a mean of zero and a constant variance. The meaning of this is that the variances of the independent variables are independent of the value of the variable. Consider the relationship between personal income and the quantity of a good purchased as an example of a case where the variance is dependent upon the value of the independent variable, income. It is plausible that as income increases the variation around the amount purchased will also increase simply because of the flexibility provided with higher levels of income. The assumption is for constant variance with respect to the magnitude of the independent variable called homoscedasticity. If the assumption fails, then it is called heteroscedasticity. Figure 13.6 shows the case of homoscedasticity where all three distributions have the same variance around the predicted value of Y regardless of the magnitude of X .
3. While the independent variables are all fixed values they are from a probability distribution that is normally distributed. This can be seen in Figure 13.6 by the shape of the distributions placed on the predicted line at the expected value of the relevant value of Y .
4. The independent variables are independent of Y , but are also assumed to be independent of the other X variables. The model is designed to estimate the effects of independent variables on some dependent variable in accordance with a proposed theory. The case where some or more of the independent variables are correlated is not unusual. There may be no cause and effect relationship among the independent variables, but nevertheless they move together. Take the case of a simple supply curve

where quantity supplied is theoretically related to the price of the product and the prices of inputs. There may be multiple inputs that may over time move together from general inflationary pressure. The input prices will therefore violate this assumption of regression analysis. This condition is called multicollinearity, which will be taken up in detail later.

- The error terms are uncorrelated with each other. This situation arises from an effect on one error term from another error term. While not exclusively a time series problem, it is here that we most often see this case. An X variable in time period one has an effect on the Y variable, but this effect then has an effect in the next time period. This effect gives rise to a relationship among the error terms. This case is called autocorrelation, "self-correlated." The error terms are now not independent of each other, but rather have their own effect on subsequent error terms.

Figure 13.6 does not show all the assumptions of the regression model, but it helps visualize these important ones.





Figure 13.6



$$y = \beta_0 + \beta_1 X + \varepsilon$$

Figure 13.7

This is the general form that is most often called the multiple regression model. So-called "simple" regression analysis has only one independent (right-hand) variable rather than many independent variables. Simple regression is just a special case of multiple regression. There is some value in beginning with simple regression: it is easy to graph in two dimensions, difficult to graph in three dimensions, and impossible to graph in more than three dimensions. Consequently, our graphs will be for the simple regression case. Figure 13.7 presents the regression problem in the form of a scatter plot graph of the data set where it is hypothesized that Y is dependent upon the single independent variable X .

A basic relationship from Macroeconomic Principles is the consumption function. This theoretical relationship states that as a person's income rises, their consumption rises, but by a smaller amount than the rise in income. If Y is consumption and X is income in the equation below Figure 13.7, the regression problem is, first, to establish that this relationship exists, and second, to determine the impact of a change in income on a person's consumption. The parameter β_1 was called the Marginal Propensity to Consume in Macroeconomics Principles.

Each "dot" in Figure 13.7 represents the consumption and income of different individuals at some point in time. This was called cross-section data earlier; observations on variables at one point in time across different people or other units of measurement. This analysis is often done with time series data, which would be the consumption and income of one individual or country at different

points in time. For macroeconomic problems it is common to use times series aggregated data for a whole country. For this particular theoretical concept these data are readily available in the annual report of the President's Council of Economic Advisors.

Figure 13.8. Regression analysis is sometimes called "least squares" analysis because the method of determining which line best "fits" the data is to minimize the sum of the squared residuals of a line put through the data.




Figure 13.8

$$\text{Population Equation: } C = \beta_0 + \beta_1 \text{Income} + \varepsilon$$

$$\text{Estimated Equation: } C = b_0 + b_1 \text{Income} + e$$

This figure shows the assumed relationship between consumption and income from macroeconomic theory. Here the data are plotted as a scatter plot and an estimated straight line has been drawn. From this graph we can see an error term, e_1 . Each data point also has an error term. Again, the error term is put into the equation to capture effects on consumption that are not caused by income changes. Such other effects might be a person's savings or wealth, or periods of unemployment. We will see how by minimizing the sum of these errors we can get an estimate for the slope and intercept of this line.

Consider the graph below. The notation has returned to that for the more general model rather than the specific case of the Macroeconomic consumption function in our example.




Figure 13.9

The \hat{y} is read "**y hat**" and is the **estimated value of y**. (In Figure 13.8 \hat{C} represents the estimated value of consumption because it is on the estimated line.) It is the value of y obtained using the regression line. \hat{y} is not generally equal to y from the data.

The term $y_0 - \hat{y}_0 = e_0$ is called the "**error**" or **residual**. It is not an error in the sense of a mistake. The error term was put into the estimating equation to capture missing variables and errors in measurement that may have occurred in the dependent variables. The **absolute value of a residual** measures the vertical distance between the actual value of y and the estimated value of y . In other words, it measures the vertical distance between the actual data point and the predicted point on the line as can be seen on the graph at point X_0 .

If the observed data point lies above the line, the residual is positive, and the line underestimates the actual data value for y .

If the observed data point lies below the line, the residual is negative, and the line overestimates that actual data value for y .

In the graph, $y_0 - \hat{y}_0 = e_0$ is the residual for the point shown. Here the point lies above the line and the residual is positive. For each data point the residuals, or errors, are calculated $y_i - \hat{y}_i = e_i$ for $i = 1, 2, 3, \dots, n$ where n is the sample size. Each $|e|$ is a vertical distance.

The sum of the errors squared is the term obviously called **Sum of Squared Errors (SSE)**.

Using calculus, you can determine the straight line that has the parameter values of b_0 and b_1 that minimizes the **SSE**. When you make the **SSE** a minimum, you have determined the points that are on the line of best fit. It turns out that the line of best fit has the equation:

$$\hat{y} = b_0 + b_1 x$$

where $b_0 = \bar{y} - b_1 \bar{x}$ and $b_1 = \frac{\Sigma(x-\bar{x})(y-\bar{y})}{\Sigma(x-\bar{x})^2} = \frac{\text{cov}(x,y)}{s_x^2}$

The sample means of the x values and the y values are \bar{x} and \bar{y} , respectively. The best fit line always passes through the point (\bar{y}, \bar{x}) called the points of means.

The slope b can also be written as:

$$b_1 = r_{y,x} \left(\frac{s_y}{s_x} \right)$$

where s_y = the standard deviation of the y values and s_x = the standard deviation of the x values and r is the correlation coefficient between x and y .

These equations are called the Normal Equations and come from another very important mathematical finding called the Gauss-Markov Theorem without which we could not do regression analysis. The Gauss-Markov Theorem tells us that the estimates we get from using the ordinary least squares (OLS) regression method will result in estimates that have some very important properties. In the Gauss-Markov Theorem it was proved that a least squares line is BLUE, which is, **B**est, **L**inear, **U**nbiased, **E**stimator. Best is the statistical property that an estimator is the one with the minimum variance. Linear refers to the property of the type of line being estimated. An unbiased estimator is one whose estimating function has an expected mean equal to the mean of the population. (You will remember that the expected value of $\mu_{\bar{x}}$ was equal to the population mean μ in accordance with the Central Limit Theorem. This is exactly the same concept here).

Both Gauss and Markov were giants in the field of mathematics, and Gauss in physics too, in the 18th century and early 19th century. They barely overlapped chronologically and never in geography, but Markov's work on this theorem was based extensively on the earlier work of Carl Gauss. The extensive applied value of this theorem had to wait until the middle of this last century.

Using the OLS method we can now find the **estimate of the error variance** which is the variance of the squared errors, e^2 . This is sometimes called the **standard error of the estimate**. (Grammatically this is probably best said as the estimate of the **error's**variance) The formula for the estimate of the error variance is:

$$s_e^2 = \frac{\Sigma(y_i - \hat{y}_i)^2}{n - k} = \frac{\Sigma e_i^2}{n - k}$$

where \hat{y} is the predicted value of y and y is the observed value, and thus the term $(y_i - \hat{y}_i)^2$ is the squared errors that are to be minimized to find the estimates of the regression line parameters. This is really just the variance of the error terms and follows our regular variance formula. One important note is that here we are dividing by $(n - k)$, which is the degrees of freedom. The degrees of freedom of a regression equation will be the number of observations, n , reduced by the number of estimated parameters, which includes the intercept as a parameter.

The variance of the errors is fundamental in testing hypotheses for a regression. It tells us just how "tight" the dispersion is about the line. As we will see shortly, the greater the dispersion about the line, meaning the larger the variance of the errors, the less probable that the hypothesized independent variable will be found to have a significant effect on the dependent variable. In short, the theory being tested will more likely fail if the variance of the error term is high. Upon reflection this should not be a surprise. As we tested hypotheses about a mean we observed that large variances reduced the calculated test statistic and thus it failed to

reach the tail of the distribution. In those cases, the null hypotheses could not be rejected. If we cannot reject the null hypothesis in a regression problem, we must conclude that the hypothesized independent variable has no effect on the dependent variable.

A way to visualize this concept is to draw two scatter plots of x and y data along a predetermined line. The first will have little variance of the errors, meaning that all the data points will move close to the line. Now do the same except the data points will have a large estimate of the error variance, meaning that the data points are scattered widely along the line. Clearly the confidence about a relationship between x and y is effected by this difference between the estimate of the error variance.

Testing the Parameters of the Line

The whole goal of the regression analysis was to test the hypothesis that the dependent variable, Y , was in fact dependent upon the values of the independent variables as asserted by some foundation theory, such as the consumption function example. Looking at the estimated equation under Figure 13.8, we see that this amounts to determining the values of b_0 and b_1 . Notice that again we are using the convention of Greek letters for the population parameters and Roman letters for their estimates.

The regression analysis output provided by the computer software will produce an estimate of b_0 and b_1 , and any other b 's for other independent variables that were included in the estimated equation. The issue is how good are these estimates? In order to test a hypothesis concerning any estimate, we have found that we need to know the underlying sampling distribution. It should come as no surprise at this stage in the course that the answer is going to be the normal distribution. This can be seen by remembering the assumption that the error term in the population, ϵ , is normally distributed. If the error term is normally distributed and the variance of the estimates of the equation parameters, b_0 and b_1 , are determined by the variance of the error term, it follows that the variances of the parameter estimates are also normally distributed. And indeed this is just the case.

We can see this by the creation of the test statistic for the test of hypothesis for the slope parameter, β_1 in our consumption function equation. To test whether or not Y does indeed depend upon X , or in our example, that consumption depends upon income, we need only test the hypothesis that β_1 equals zero. This hypothesis would be stated formally as:

$$H_0 : \beta_1 = 0$$

$$H_a : \beta_1 \neq 0$$

If we cannot reject the null hypothesis, we must conclude that our theory has no validity. If we cannot reject the null hypothesis that $\beta_1 = 0$ then b_1 , the coefficient of Income, is zero and zero times anything is zero. Therefore the effect of Income on Consumption is zero. There is no relationship as our theory had suggested.

Notice that we have set up the presumption, the null hypothesis, as "no relationship". This puts the burden of proof on the alternative hypothesis. In other words, if we are to validate our claim of finding a relationship, we must do so with a level of significance greater than 90, 95, or 99 percent. The status quo is ignorance, no relationship exists, and to be able to make the claim that we have actually added to our body of knowledge we must do so with significant probability of being correct. John Maynard Keynes got it right and thus was born Keynesian economics starting with this basic concept in 1936.

The test statistic for this test comes directly from our old friend the standardizing formula:

$$t_c = \frac{b_1 - \beta_1}{S_{b_1}}$$

where b_1 is the estimated value of the slope of the regression line, β_1 is the hypothesized value of beta, in this case zero, and S_{b_1} is the standard deviation of the estimate of b_1 . In this case we are asking how many standard deviations is the estimated slope away from the hypothesized slope. This is exactly the same question we asked before with respect to a hypothesis about a mean: how many standard deviations is the estimated mean, the sample mean, from the hypothesized mean?

The test statistic is written as a student's t distribution, but if the sample size is larger enough so that the degrees of freedom are greater than 30 we may again use the normal distribution. To see why we can use the student's t or normal distribution we have only to look at S_{b_1} , the formula for the standard deviation of the estimate of b_1 :

$$S_{b_1} = \frac{S_e^2}{\sqrt{(x_i - \bar{x})^2}}$$

or

$$S_{b_1} = \frac{S_e^2}{(n-1)S_x^2}$$

Where S_e is the estimate of the error variance and S_x^2 is the variance of x values of the coefficient of the independent variable being tested.

We see that S_e , the **estimate of the error variance**, is part of the computation. Because the estimate of the error variance is based on the assumption of normality of the error terms, we can conclude that the sampling distribution of the b 's, the coefficients of our hypothesized regression line, are also normally distributed.

One last note concerns the degrees of freedom of the test statistic, $\nu = n - k$. Previously we subtracted 1 from the sample size to determine the degrees of freedom in a student's t problem. Here we must subtract one degree of freedom for each parameter estimated in the equation. For the example of the consumption function we lose 2 degrees of freedom, one for b_0 , the intercept, and one for b_1 , the slope of the consumption function. The degrees of freedom would be $n - k - 1$, where k is the number of independent variables and the extra one is lost because of the intercept. If we were estimating an equation with three independent variables, we would lose 4 degrees of freedom: three for the independent variables, k , and one more for the intercept.

The decision rule for acceptance or rejection of the null hypothesis follows exactly the same form as in all our previous test of hypothesis. Namely, if the calculated value of t (or Z) falls into the tails of the distribution, where the tails are defined by α , the required significance level in the test, we cannot accept the null hypothesis. If on the other hand, the calculated value of the test statistic is within the critical region, we cannot reject the null hypothesis.

If we conclude that we cannot accept the null hypothesis, we are able to state with $(1 - \alpha)$ level of confidence that the slope of the line is given by b_1 . This is an extremely important conclusion. Regression analysis not only allows us to test if a cause and effect relationship exists, we can also determine the magnitude of that relationship, if one is found to exist. It is this feature of regression analysis that makes it so valuable. If models can be developed that have statistical validity, we are then able to simulate the effects of changes in variables that may be under our control with some degree of probability, of course. For example, if advertising is demonstrated to effect sales, we can determine the effects of changing the advertising budget and decide if the increased sales are worth the added expense.

Multicollinearity

Our discussion earlier indicated that like all statistical models, the OLS regression model has important assumptions attached. Each assumption, if violated, has an effect on the ability of the model to provide useful and meaningful estimates. The Gauss-Markov Theorem has assured us that the OLS estimates are unbiased and minimum variance, but this is true only under the assumptions of the model. Here we will look at the effects on OLS estimates if the independent variables are correlated. The other assumptions and the methods to mitigate the difficulties they pose if they are found to be violated are examined in Econometrics courses. We take up multicollinearity because it is so often prevalent in Economic models and it often leads to frustrating results.

The OLS model assumes that all the independent variables are independent of each other. This assumption is easy to test for a particular sample of data with simple correlation coefficients. Correlation, like much in statistics, is a matter of degree: a little is not good, and a lot is terrible.

The goal of the regression technique is to tease out the independent impacts of each of a set of independent variables on some hypothesized dependent variable. If two 2 independent variables are interrelated, that is, correlated, then we cannot isolate the effects on Y of one from the other. In an extreme case where x_1 is a linear combination of x_2 , correlation equal to one, both variables move in identical ways with Y . In this case it is impossible to determine the variable that is the true cause of the effect on Y . (If the two variables were actually perfectly correlated, then mathematically no regression results could actually be calculated.)

The normal equations for the coefficients show the effects of multicollinearity on the coefficients.

$$b_1 = \frac{s_y (r_{x_1y} - r_{x_1x_2} r_{x_2y})}{s_{x_1} (1 - r_{x_1x_2}^2)}$$

$$b_2 = \frac{s_y (r_{x_2y} - r_{x_1x_2} r_{x_1y})}{s_{x_2} (1 - r_{x_1x_2}^2)}$$

$$b_0 = \bar{y} - b_1 \bar{x}_1 - b_2 \bar{x}_2$$

The correlation between x_1 and x_2 , $r_{x_1x_2}^2$, appears in the denominator of both the estimating formula for b_1 and b_2 . If the assumption of independence holds, then this term is zero. This indicates that there is no effect of the correlation on the coefficient. On the other hand, as the correlation between the two independent variables increases the denominator decreases, and thus the estimate of the coefficient increases. The correlation has the same effect on both of the coefficients of these two variables. In essence, each variable is “taking” part of the effect on Y that should be attributed to the collinear variable. This results in biased estimates.

Multicollinearity has a further deleterious impact on the OLS estimates. The correlation between the two independent variables also shows up in the formulas for the estimate of the variance for the coefficients.

$$s_{b_1}^2 = \frac{s_e^2}{(n-1)s_{x_1}^2(1-r_{x_1x_2}^2)}$$

$$s_{b_2}^2 = \frac{s_e^2}{(n-1)s_{x_2}^2(1-r_{x_1x_2}^2)}$$

Here again we see the correlation between x_1 and x_2 in the denominator of the estimates of the variance for the coefficients for both variables. If the correlation is zero as assumed in the regression model, then the formula collapses to the familiar ratio of the variance of the errors to the variance of the relevant independent variable. If however the two independent variables are correlated, then the variance of the estimate of the coefficient increases. This results in a smaller t -value for the test of hypothesis of the coefficient. In short, multicollinearity results in failing to reject the null hypothesis that the X variable has no impact on Y when in fact X does have a statistically significant impact on Y . Said another way, the large standard errors of the estimated coefficient created by multicollinearity suggest statistical insignificance even when the hypothesized relationship is strong.

How Good is the Equation?

In the last section we concerned ourselves with testing the hypothesis that the dependent variable did indeed depend upon the hypothesized independent variable or variables. It may be that we find an independent variable that has some effect on the dependent variable, but it may not be the only one, and it may not even be the most important one. Remember that the error term was placed in the model to capture the effects of any missing independent variables. It follows that the error term may be used to give a measure of the “goodness of fit” of the equation taken as a whole in explaining the variation of the dependent variable, Y .

The **multiple correlation coefficient**, also called the **coefficient of multiple determination** or the **coefficient of determination**, is given by the formula:

$$R^2 = \frac{SSR}{SST}$$

where SSR is the regression sum of squares, the squared deviation of the predicted value of y from the mean value of y ($\hat{y} - \bar{y}$), and SST is the total sum of squares, the squared deviation of the observed value of y from the mean value of y ($y_i - \bar{y}$). Figure 13.10 shows how the total deviation of the dependent variable, y , is partitioned into these two pieces.




Figure 13.10

Figure 13.10 shows the estimated regression line and a single observation, x_1 . Regression analysis tries to explain the variation of the data about the mean value of the dependent variable, y . The question is, why do the observations of y vary from the average level of y ? The value of y at observation x_1 varies from the mean of y by the difference $(y_i - \bar{y})$. The sum of these differences squared is SST, the sum of squares total. The actual value of y at x_1 deviates from the estimated value, \hat{y} , by the difference between the estimated value and the actual value, $(y_i - \hat{y})$. We recall that this is the error term, e , and the sum of these errors is SSE, sum of squared errors. The deviation of the predicted value of y , \hat{y} , from the mean value of y is $(\hat{y} - \bar{y})$ and is the SSR, sum of squares regression. It is called “regression” because it is the deviation explained by the regression. (Sometimes the SSR is called SSM for sum of squares mean because it measures the deviation from the mean value of the dependent variable, y , as shown on the graph.).

Because the $SST = SSR + SSE$ we see that the multiple correlation coefficient is the percent of the variance, or deviation in y from its mean value, that is explained by the equation when taken as a whole. R^2 will vary between zero and 1, with zero indicating that none of the variation in y was explained by the equation and a value of 1 indicating that 100% of the variation in y was explained by the equation. For time series studies expect a high R^2 and for cross-section data expect low R^2 .

While a high R^2 is desirable, remember that it is the tests of the hypothesis concerning the existence of a relationship between a set of independent variables and a particular dependent variable that was the motivating factor in using the regression model. It is validating a cause and effect relationship developed by some theory that is the true reason that we chose the regression analysis. Increasing the number of independent variables will have the effect of increasing R^2 . To account for this effect the proper measure of the coefficient of determination is the \bar{R}^2 , adjusted for degrees of freedom, to keep down mindless addition of independent variables.

There is no statistical test for the R^2 and thus little can be said about the model using R^2 with our characteristic confidence level. Two models that have the same size of SSE, that is sum of squared errors, may have very different R^2 if the competing models have different SST, total sum of squared deviations. The goodness of fit of the two models is the same; they both have the same sum of squares unexplained, errors squared, but because of the larger total sum of squares on one of the models the R^2 differs. Again, the real value of regression as a tool is to examine hypotheses developed from a model that predicts certain relationships among the variables. These are tests of hypotheses on the coefficients of the model and not a game of maximizing R^2 .

Another way to test the general quality of the overall model is to test the coefficients as a group rather than independently. Because this is multiple regression (more than one X), we use the F-test to determine if our coefficients collectively affect Y. The hypothesis is:

$$H_o : \beta_1 = \beta_2 = \dots = \beta_i = 0$$

H_a : "at least one of the β_i is not equal to 0"

If the null hypothesis cannot be rejected, then we conclude that none of the independent variables contribute to explaining the variation in Y . Reviewing Figure 13.10 we see that SSR, the explained sum of squares, is a measure of just how much of the variation in Y is explained by all the variables in the model. SSE, the sum of the errors squared, measures just how much is unexplained. It follows that the ratio of these two can provide us with a statistical test of the model as a whole. Remembering that the F distribution is a ratio of Chi squared distributions and that variances are distributed according to Chi Squared, and the sum of squared errors and the sum of squares are both variances, we have the test statistic for this hypothesis as:

$$F_c = \frac{\left(\frac{SSR}{k} \right)}{\left(\frac{SSE}{n-k-1} \right)}$$

where n is the number of observations and k is the number of independent variables. It can be shown that this is equivalent to:

$$F_c = \frac{n-k-1}{k} \cdot \frac{R^2}{1-R^2}$$

Figure 13.10 where R^2 is the coefficient of determination which is also a measure of the “goodness” of the model.

As with all our tests of hypothesis, we reach a conclusion by comparing the calculated F statistic with the critical value given our desired level of confidence. If the calculated test statistic, an F statistic in this case, is in the tail of the distribution, then we cannot accept the null hypothesis. By not being able to accept the null hypotheses we conclude that this specification of this model has validity, because at least one of the estimated coefficients is significantly different from zero.

An alternative way to reach this conclusion is to use the p -value comparison rule. The p -value is the area in the tail, given the calculated F statistic. In essence, the computer is finding the F value in the table for us. The computer regression output for the calculated F statistic is typically found in the ANOVA table section labeled “significance F”. How to read the output of an Excel regression is presented below. This is the probability of NOT accepting a false null hypothesis. If this probability is less than our pre-determined alpha error, then the conclusion is that we cannot accept the null hypothesis.

Dummy Variables

Thus far the analysis of the OLS regression technique assumed that the independent variables in the models tested were continuous random variables. There are, however, no restrictions in the regression model against independent variables that are binary. This opens the regression model for testing hypotheses concerning categorical variables such as gender, race, region of the country, before a certain date, after a certain date and innumerable others. These categorical variables take on only two values, 1 and 0, success or failure, from the binomial probability distribution. The form of the equation becomes:

$$\hat{y} = b_0 + b_2x_2 + b_1x_1$$




Figure 13.11

where $x_2 = 0$. X_2 is the dummy variable and X_1 is some continuous random variable. The constant, b_0 , is the y -intercept, the value where the line crosses the y -axis. When the value of $X_2 = 0$, the estimated line crosses at b_0 . When the value of $X_2 = 1$ then the estimated line crosses at $b_0 + b_2$. In effect the dummy variable causes the estimated line to shift either up or down by the size of the effect of the characteristic captured by the dummy variable. Note that this is a simple parallel shift and does not affect the impact of the other independent variable; X_1 . This variable is a continuous random variable and predicts different values of y at different values of X_1 holding constant the condition of the dummy variable.

An example of the use of a dummy variable is the work estimating the impact of gender on salaries. There is a full body of literature on this topic and dummy variables are used extensively. For this example the salaries of elementary and secondary school teachers for a particular state is examined. Using a homogeneous job category, school teachers, and for a single state reduces many of the variations that naturally effect salaries such as differential physical risk, cost of living in a particular state, and other working conditions. The estimating equation in its simplest form specifies salary as a function of various teacher characteristic that economic theory would suggest could affect salary. These would include education level as a measure of potential productivity, age and/or experience to capture on-the-job training, again as a measure of productivity. Because the data are for school teachers employed in a public school districts rather than workers in a for-profit company, the school district’s average revenue per average daily student attendance is included as a measure of ability to pay. The results of the regression analysis using data on 24,916 school teachers are presented below.

Table 13.1 Earnings Estimate for Elementary and Secondary School Teachers

Variable	Regression Coefficients (b)	Standard Errors of the estimates for teacher's earnings function (sb)
Intercept	4269.9	
Gender (male = 1)	632.38	13.39
Total Years of Experience	52.32	1.10
Years of Experience in Current District	29.97	1.52

Variable	Regression Coefficients (b)	Standard Errors of the estimates for teacher's earnings function (sb)
Education	629.33	13.16
Total Revenue per ADA	90.24	3.76
\bar{R}^2	.725	
n	24,916	

The coefficients for all the independent variables are significantly different from zero as indicated by the standard errors. Dividing the standard errors of each coefficient results in a t-value greater than 1.96 which is the required level for 95% significance. The binary variable, our dummy variable of interest in this analysis, is gender where male is given a value of 1 and female given a value of 0. The coefficient is significantly different from zero with a dramatic t-statistic of 47 standard deviations. We thus cannot accept the null hypothesis that the coefficient is equal to zero. Therefore we conclude that there is a premium paid male teachers of \$632 after holding constant experience, education and the wealth of the school district in which the teacher is employed. It is important to note that these data are from some time ago and the \$632 represents a six percent salary premium at that time. A graph of this example of dummy variables is presented below.

TEACHER'S SALARY

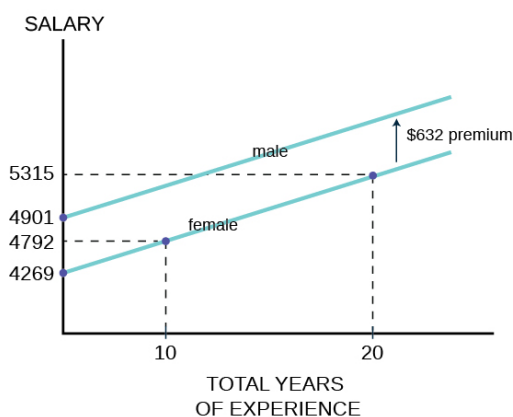



Figure 13.12

In two dimensions, salary is the dependent variable on the vertical axis and total years of experience was chosen for the continuous independent variable on horizontal axis. Any of the other independent variables could have been chosen to illustrate the effect of the dummy variable. The relationship between total years of experience has a slope of \$52.32 per year of experience and the estimated line has an intercept of \$4,269 if the gender variable is equal to zero, for female. If the gender variable is equal to 1, for male, the coefficient for the gender variable is added to the intercept and thus the relationship between total years of experience and salary is shifted upward parallel as indicated on the graph. Also marked on the graph are various points for reference. A female school teacher with 10 years of experience receives a salary of \$4,792 on the basis of her experience only, but this is still \$109 less than a male teacher with zero years of experience.

A more complex interaction between a dummy variable and the dependent variable can also be estimated. It may be that the dummy variable has more than a simple shift effect on the dependent variable, but also interacts with one or more of the other continuous independent variables. While not tested in the example above, it could be hypothesized that the impact of gender on salary was not a one-time shift, but impacted the value of additional years of experience on salary also. That is, female school teacher's salaries were discounted at the start, and further did not grow at the same rate from the effect of experience as for male school teachers. This would show up as a different slope for the relationship between total years of experience for males than for females. If this is so then females school teachers would not just start behind their male colleagues (as measured by the shift in the estimated regression line), but would fall further and further behind as time and experienced increased.

The graph below shows how this hypothesis can be tested with the use of dummy variables and an interaction variable.



$$\hat{y} = b_0 + b_2x_2 + b_1x_1 + b_3x_2x_1$$

Figure 13.13

The estimating equation shows how the slope of X_1 , the continuous random variable experience, contains two parts, b_1 and b_3 . This occurs because of the new variable $X_2 X_1$, called the interaction variable, was created to allow for an effect on the slope of X_1 from changes in X_2 , the binary dummy variable. Note that when the dummy variable, $X_2 = 0$ the interaction variable has a value of 0, but when $X_2 = 1$ the interaction variable has a value of X_1 . The coefficient b_3 is an estimate of the difference in the coefficient of X_1 when $X_2 = 1$ compared to when $X_2 = 0$. In the example of teacher's salaries, if there is a premium paid to male teachers that affects the rate of increase in salaries from experience, then the rate at which male teachers' salaries rises would be $b_1 + b_3$ and the rate at which female teachers' salaries rise would be simply b_1 . This hypothesis can be tested with the hypothesis:

$$H_0 : \beta_3 = 0 | \beta_1 = 0, \beta_2 = 0$$

$$H_a : \beta_3 \neq 0 | \beta_1 \neq 0, \beta_2 \neq 0$$

This is a t -test using the test statistic for the parameter β_3 . If we cannot accept the null hypothesis that $\beta_3 = 0$ we conclude there is a difference between the rate of increase for the group for whom the value of the binary variable is set to 1, males in this example. This estimating equation can be combined with our earlier one Figure 13.13 are drawn for this case with a shift in the earnings function and a difference in the slope of the function with respect to total years of experience.

Example 13.5

A random sample of 11 statistics students produced the following data, where x is the third exam score out of 80, and y is the final exam score out of 200. Can you predict the final exam score of a randomly selected student if you know the third exam score?

Table showing the scores on the final exam based on scores from the third exam.

Table 13.2

x (third exam score)	y (final exam score)
65	175
67	133
71	185
71	163
66	126
75	198
67	153
70	163
71	159

x (third exam score)	y (final exam score)
69	151
69	159




Figure 13.14 Scatter plot showing the scores on the final exam based on scores from the third exam.

This page titled [13.5: The Regression Equation](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **13.4: The Regression Equation** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.6: Interpretation of Regression Coefficients- Elasticity and Logarithmic Transformation

As we have seen, the coefficient of an equation estimated using OLS regression analysis provides an estimate of the slope of a straight line that is assumed to be the relationship between the dependent variable and at least one independent variable. From the calculus, the slope of the line is the first derivative and tells us the magnitude of the impact of a one unit change in the X variable upon the value of the Y variable measured in the units of the Y variable. As we saw in the case of dummy variables, this can show up as a parallel shift in the estimated line or even a change in the slope of the line through an interactive variable. Here we wish to explore the concept of elasticity and how we can use a regression analysis to estimate the various elasticities in which economists have an interest.

The concept of elasticity is borrowed from engineering and physics where it is used to measure a material's responsiveness to a force, typically a physical force such as a stretching/pulling force. It is from here that we get the term an "elastic" band. In economics, the force in question is some market force such as a change in price or income. Elasticity is measured as a percentage change/response in both engineering applications and in economics. The value of measuring in percentage terms is that the units of measurement do not play a role in the value of the measurement and thus allows direct comparison between elasticities. As an example, if the price of gasoline increased say 50 cents from an initial price of \$3.00 and generated a decline in monthly consumption for a consumer from 50 gallons to 48 gallons we calculate the elasticity to be 0.25. The price elasticity is the percentage change in quantity resulting from some percentage change in price. A 16 percent increase in price has generated only a 4 percent decrease in demand: 16% price change \rightarrow 4% quantity change or $.04/.16 = .25$. This is called an inelastic demand meaning a small response to the price change. This comes about because there are few if any real substitutes for gasoline; perhaps public transportation, a bicycle or walking. Technically, of course, the percentage change in demand from a price increase is a decline in demand thus price elasticity is a negative number. The common convention, however, is to talk about elasticity as the absolute value of the number. Some goods have many substitutes: pears for apples for plums, for grapes, etc. etc. The elasticity for such goods is larger than one and are called elastic in demand. Here a small percentage change in price will induce a large percentage change in quantity demanded. The consumer will easily shift the demand to the close substitute.

While this discussion has been about price changes, any of the independent variables in a demand equation will have an associated elasticity. Thus, there is an income elasticity that measures the sensitivity of demand to changes in income: not much for the demand for food, but very sensitive for yachts. If the demand equation contains a term for substitute goods, say candy bars in a demand equation for cookies, then the responsiveness of demand for cookies from changes in prices of candy bars can be measured. This is called the cross-price elasticity of demand and to an extent can be thought of as brand loyalty from a marketing view. How responsive is the demand for Coca-Cola to changes in the price of Pepsi?

Now imagine the demand for a product that is very expensive. Again, the measure of elasticity is in percentage terms thus the elasticity can be directly compared to that for gasoline: an elasticity of 0.25 for gasoline conveys the same information as an elasticity of 0.25 for \$25,000 car. Both goods are considered by the consumer to have few substitutes and thus have inelastic demand curves, elasticities less than one.

The mathematical formulae for various elasticities are:

$$\text{Price elasticity: } \eta_p = \frac{(\% \Delta Q)}{(\% \Delta P)}$$

Where η is the Greek small case letter eta used to designate elasticity. Δ is read as "change".

$$\text{Income elasticity: } \eta_Y = \frac{(\% \Delta Q)}{(\% \Delta Y)}$$

Where Y is used as the symbol for income.

$$\text{Cross-Price elasticity: } \eta_{p1} = \frac{(\% \Delta Q_1)}{(\% \Delta P_2)}$$

Where P_2 is the price of the substitute good.

Examining closer the price elasticity we can write the formula as:

$$\eta_p = \frac{(\% \Delta Q)}{(\% \Delta P)} = \frac{dQ}{dP} \left(\frac{P}{Q} \right) = b \left(\frac{P}{Q} \right)$$

Where b is the estimated coefficient for price in the OLS regression.

The first form of the equation demonstrates the principle that elasticities are measured in percentage terms. Of course, the ordinary least squares coefficients provide an estimate of the impact of a unit change in the independent variable, X , on the dependent variable measured in units of Y . These coefficients are not elasticities, however, and are shown in the second way of writing the formula for elasticity as $\left(\frac{dQ}{dP} \right)$, the derivative of the estimated demand function which is simply the slope of the regression line. Multiplying the slope times $\frac{P}{Q}$ provides an elasticity measured in percentage terms.

Along a straight-line demand curve the percentage change, thus elasticity, changes continuously as the scale changes, while the slope, the estimated regression coefficient, remains constant. Going back to the demand for gasoline. A change in price from \$3.00 to \$3.50 was a 16 percent increase in price. If the beginning price were \$5.00 then the same 50¢ increase would be only a 10 percent increase generating a different elasticity. Every straight-line demand curve has a range of elasticities starting at the top left, high prices, with large elasticity numbers, elastic demand, and decreasing as one goes down the demand curve, inelastic demand.

In order to provide a meaningful estimate of the elasticity of demand the convention is to estimate the elasticity at the point of means. Remember that all OLS regression lines will go through the point of means. At this point is the greatest weight of the data used to estimate the coefficient. The formula to estimate an elasticity when an OLS demand curve has been estimated becomes:

$$\eta_p = b \left(\frac{\bar{P}}{\bar{Q}} \right)$$

Where \bar{P} and \bar{Q} are the mean values of these data used to estimate b , the price coefficient.

The same method can be used to estimate the other elasticities for the demand function by using the appropriate mean values of the other variables; income and price of substitute goods for example.

Logarithmic Transformation of the Data

Ordinary least squares estimates typically assume that the population relationship among the variables is linear thus of the form presented in [The Regression Equation](#). In this form the interpretation of the coefficients is as discussed above; quite simply the coefficient provides an estimate of the impact of a one **unit** change in X on Y measured in **units** of Y . It does not matter just where along the line one wishes to make the measurement because it is a straight line with a constant slope thus constant estimated level of impact per unit change. It may be, however, that the analyst wishes to estimate not the simple unit measured impact on the Y variable, but the magnitude of the percentage impact on Y of a one unit change in the X variable. Such a case might be how a **unit change** in experience, say one year, effects not the absolute amount of a worker's wage, but the **percentage impact** on the worker's wage. Alternatively, it may be that the question asked is the unit measured impact on Y of a specific percentage increase in X . An example may be "by how many dollars will sales increase if the firm spends X percent more on advertising?" The third possibility is the case of elasticity discussed above. Here we are interested in the percentage impact on quantity demanded for a given percentage change in price, or income or perhaps the price of a substitute good. All three of these cases can be estimated by transforming the data to logarithms before running the regression. The resulting coefficients will then provide a percentage change measurement of the relevant variable.

To summarize, there are four cases:

1. Unit $\Delta X \rightarrow$ Unit ΔY (Standard OLS case)
2. Unit $\Delta X \rightarrow \% \Delta Y$
3. $\% \Delta X \rightarrow$ Unit ΔY
4. $\% \Delta X \rightarrow \% \Delta Y$ (elasticity case)

Case 1: The ordinary least squares case begins with the linear model developed above:

$$Y = a + bX$$

where the coefficient of the independent variable $b = \frac{dY}{dX}$ is the slope of a straight line and thus measures the impact of a unit change in X on Y measured in units of Y .

Case 2: The underlying estimated equation is:

$$\log(Y) = a + bX$$

The equation is estimated by converting the Y values to logarithms and using OLS techniques to estimate the coefficient of the X variable, b . This is called a semi-log estimation. Again, differentiating both sides of the equation allows us to develop the interpretation of the X coefficient b :

$$\begin{aligned} d(\log Y) &= b dX \\ \frac{dY}{Y} &= b dX \end{aligned}$$

Multiply by 100 to convert to percentages and rearranging terms gives:

$$100b = \frac{\% \Delta Y}{\text{Unit } \Delta X}$$

$100b$ is thus the percentage change in Y resulting from a unit change in X .

Case 3: In this case the question is “what is the unit change in Y resulting from a percentage change in X ?” What is the dollar loss in revenues of a five percent increase in price or what is the total dollar cost impact of a five percent increase in labor costs? The estimated equation for this case would be:

$$Y = a + B \log(X)$$

Here the calculus differential of the estimated equation is:

$$\begin{aligned} dY &= b d(\log X) \\ dY &= b \frac{dX}{X} \end{aligned}$$

Divide by 100 to get percentage and rearranging terms gives:

$$\frac{b}{100} = \frac{dY}{100 \frac{dX}{X}} = \frac{\text{Unit } \Delta Y}{\% \Delta X}$$

Therefore, $\frac{b}{100}$ is the increase in Y measured in units from a one percent increase in X .

Case 4: This is the elasticity case where both the dependent and independent variables are converted to logs before the OLS estimation. This is known as the log-log case or double log case, and provides us with direct estimates of the elasticities of the independent variables. The estimated equation is:

$$\log Y = a + b \log X$$

Differentiating we have:

$$\begin{aligned} d(\log Y) &= b d(\log X) \\ d(\log X) &= b \frac{1}{X} dX \end{aligned}$$

thus:

$$\frac{1}{Y} dY = b \frac{1}{X} dX \quad \text{OR} \quad \frac{dY}{Y} = b \frac{dX}{X} \quad \text{OR} \quad b = \frac{dY}{dX} \left(\frac{X}{Y} \right)$$

and $b = \frac{\% \Delta Y}{\% \Delta X}$ our definition of elasticity. We conclude that we can directly estimate the elasticity of a variable through double log transformation of the data. The estimated coefficient is the elasticity. It is common to use double log transformation of all variables in the estimation of demand functions to get estimates of all the various elasticities of the demand curve.

This page titled [13.6: Interpretation of Regression Coefficients- Elasticity and Logarithmic Transformation](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **13.5: Interpretation of Regression Coefficients- Elasticity and Logarithmic Transformation** by OpenStax is licensed CC BY 4.0.
Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.7: Predicting with a Regression Equation

One important value of an estimated regression equation is its ability to predict the effects on Y of a change in one or more values of the independent variables. The value of this is obvious. Careful policy cannot be made without estimates of the effects that may result. Indeed, it is the desire for particular results that drive the formation of most policy. Regression models can be, and have been, invaluable aids in forming such policies.

The Gauss-Markov theorem assures us that the point estimate of the impact on the dependent variable derived by putting in the equation the hypothetical values of the independent variables one wishes to simulate will result in an estimate of the dependent variable which is minimum variance and unbiased. That is to say that from this equation comes the best unbiased point estimate of y given the values of x .

$$\hat{y} = b_0 + b_1 X_{1i} + \cdots + b_k X_{ki}$$

Remember that point estimates do not carry a particular level of probability, or level of confidence, because points have no “width” above which there is an area to measure. This was why we developed confidence intervals for the mean and proportion earlier. The same concern arises here also. There are actually two different approaches to the issue of developing estimates of changes in the independent variable, or variables, on the dependent variable. The first approach wishes to measure the **expected mean** value of y from a specific change in the value of x : this specific value implies the expected value. Here the question is: what is the **mean** impact on y that would result from multiple hypothetical experiments on y at this specific value of x . Remember that there is a variance around the estimated parameter of x and thus each experiment will result in a bit of a different estimate of the predicted value of y .

The second approach to estimate the effect of a specific value of x on y treats the event as a single experiment: you choose x and multiply it times the coefficient and that provides a single estimate of y . Because this approach acts as if there were a single experiment the variance that exists in the parameter estimate is larger than the variance associated with the expected value approach.

The conclusion is that we have two different ways to predict the effect of values of the independent variable(s) on the dependent variable and thus we have two different intervals. Both are correct answers to the question being asked, but there are two different questions. To avoid confusion, the first case where we are asking for the **expected value** of the mean of the estimated y , is called a **confidence interval** as we have named this concept before. The second case, where we are asking for the estimate of the impact on the dependent variable y of a single experiment using a value of x , is called the **prediction interval**. The test statistics for these two interval measures within which the estimated value of y will fall are:

Confidence Interval for Expected Value of Mean Value of y for $x = x_p$

$$\hat{y} = \pm t_{\alpha/2} s_e \left(\sqrt{\frac{1}{n} + \frac{(x_p - \bar{x})^2}{s_x^2}} \right)$$

Prediction Interval for an Individual y for $x = x_p$

$$\hat{y} = \pm t_{\alpha/2} s_e \left(\sqrt{1 + \frac{1}{n} + \frac{(x_p - \bar{x})^2}{s_x^2}} \right)$$

Where s_e is the standard deviation of the error term and s_x is the standard deviation of the x variable.

The mathematical computations of these two test statistics are complex. Various computer regression software packages provide programs within the regression functions to Figure 13.7.15




Figure 13.15 Prediction and confidence intervals for regression equation; 95% confidence level.

Figure 13.7.15 shows visually the difference the standard deviation makes in the size of the estimated intervals. The confidence interval, measuring the expected value of the dependent variable, is smaller than the prediction interval for the same level of confidence. The expected value method assumes that the experiment is conducted multiple times rather than just once as in the other method. The logic here is similar, although not identical, to that discussed when developing the relationship between the sample size and the confidence interval using the Central Limit Theorem. There, as the number of experiments increased, the distribution narrowed and the confidence interval became tighter around the expected value of the mean.

It is also important to note that the intervals around a point estimate are highly dependent upon the range of data used to estimate the equation regardless of which approach is being used for prediction. Remember that all regression equations go through the point of means, that is, the mean value of y and the mean values of all independent variables in the equation. As the value of x chosen to estimate the associated value of y is further from the point of means the width of the estimated interval around Figure 13.7.16 shows this relationship.




Figure 13.16 Confidence interval for an individual value of x , X_p , at 95% level of confidence

Figure 13.7.16 demonstrates the concern for the quality of the estimated interval whether it is a prediction interval or a confidence interval. As the value chosen to predict y , X_p in the graph, is further from the central weight of the data, \bar{X} , we see the interval expand in width even while holding constant the level of confidence. This shows that the precision of any estimate will diminish as one tries to predict beyond the largest weight of the data and most certainly will degrade rapidly for predictions beyond the range of the data. Unfortunately, this is just where most predictions are desired. They can be made, but the width of the confidence interval may be so large as to render the prediction useless. Only actual calculation and the particular application can determine this, however.

Example 13.7.6

Recall the [third exam/final exam example](#).

We found the equation of the best-fit line for the final exam grade as a function of the grade on the third-exam. We can now use the least-squares regression line for prediction. Assume the coefficient for X was determined to be significantly different from zero.

Suppose you want to estimate, or predict, the mean final exam score of statistics students who received 73 on the third exam. The exam scores (**x-values**) range from 65 to 75. Since 73 is between the x-values 65 and 75, we feel comfortable to substitute $x = 73$ into the equation. Then:

$$\hat{y} = -173.51 + 4.83(73) = 179.08$$

We predict that statistics students who earn a grade of 73 on the third exam will earn a grade of 179.08 on the final exam, on average.

a. What would you predict the final exam score to be for a student who scored a 66 on the third exam?

Answer

Solution 13.6

a. 145.27

b. What would you predict the final exam score to be for a student who scored a 90 on the third exam?

Answer

Solution 13.6

b. The x values in the data are between 65 and 75. Ninety is outside of the domain of the observed x values in the data (independent variable), so you cannot reliably predict the final exam score for this student. (Even though it is possible to enter 90 into the equation for x and calculate a corresponding y value, the y value that you get will have a confidence interval that may not be meaningful.)

To understand really how unreliable the prediction can be outside of the observed x values observed in the data, make the substitution $x = 90$ into the equation.

$$\hat{y} = -173.51 + 4.83(90) = 261.19$$

The final-exam score is predicted to be 261.19. The largest the final-exam score can be is 200.

This page titled [13.7: Predicting with a Regression Equation](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [13.6: Predicting with a Regression Equation](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.8: How to Use Microsoft Excel® for Regression Analysis

This section of this chapter is here in recognition that what we are now asking requires much more than a quick calculation of a ratio or a square root. Indeed, the use of regression analysis was almost non-existent before the middle of the last century and did not really become a widely used tool until perhaps the late 1960's and early 1970's. Even then the computational ability of even the largest IBM machines is laughable by today's standards. In the early days programs were developed by the researchers and shared. There was no market for something called "software" and certainly nothing called "apps", an entrant into the market only a few years old.

With the advent of the personal computer and the explosion of a vital software market we have a number of regression and statistical analysis packages to choose from. Each has their merits. We have chosen Microsoft Excel because of the wide-spread availability both on college campuses and in the post-college market place. Stata is an alternative and has features that will be important for more advanced econometrics study if you choose to follow this path. Even more advanced packages exist, but typically require the analyst to do some significant amount of programing to conduct their analysis. The goal of this section is to demonstrate how to use Excel to run a regression and then to do so with an example of a simple version of a demand curve.

The first step to doing a regression using Excel is to load the program into your computer. If you have Excel you have the Analysis ToolPak although you may not have it activated. The program calls upon a significant amount of space so is not loaded automatically.

To activate the Analysis ToolPak follow these steps:

Click "File" > "Options" > "Add-ins" to bring up a menu of the add-in "ToolPaks". Select "Analysis ToolPak" and click "GO" next to "Manage: excel add-ins" near the bottom of the window. This will open a new window where you click "Analysis ToolPak" (make sure there is a green check mark in the box) and then click "OK". Now there should be an Analysis tab under the data menu. These steps are presented in the following screen shots.

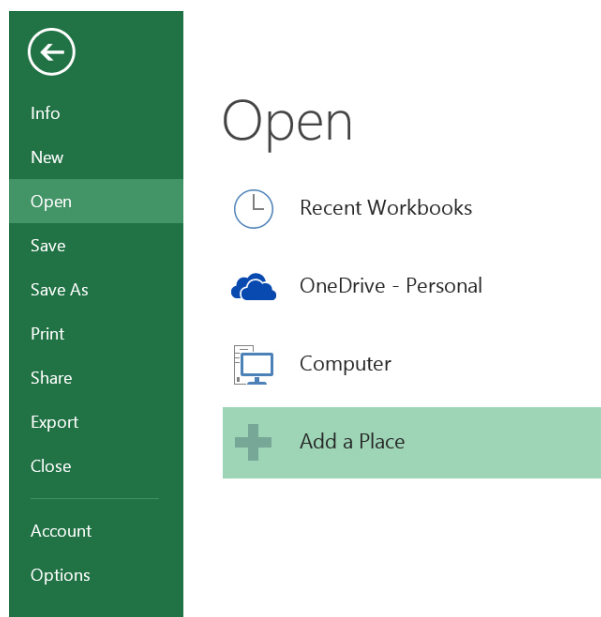


Figure 13.8.17

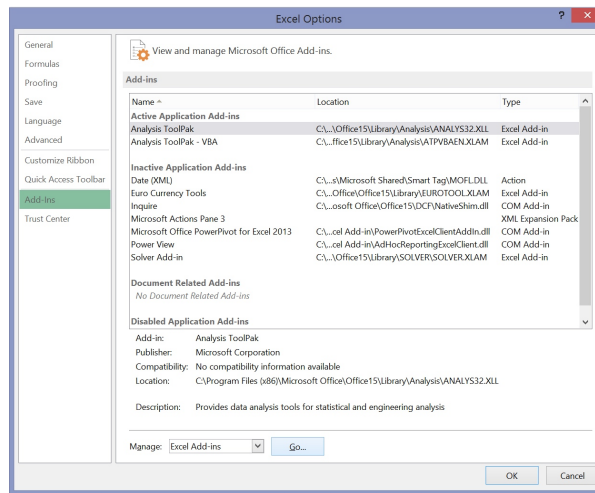


Figure 13.8.18

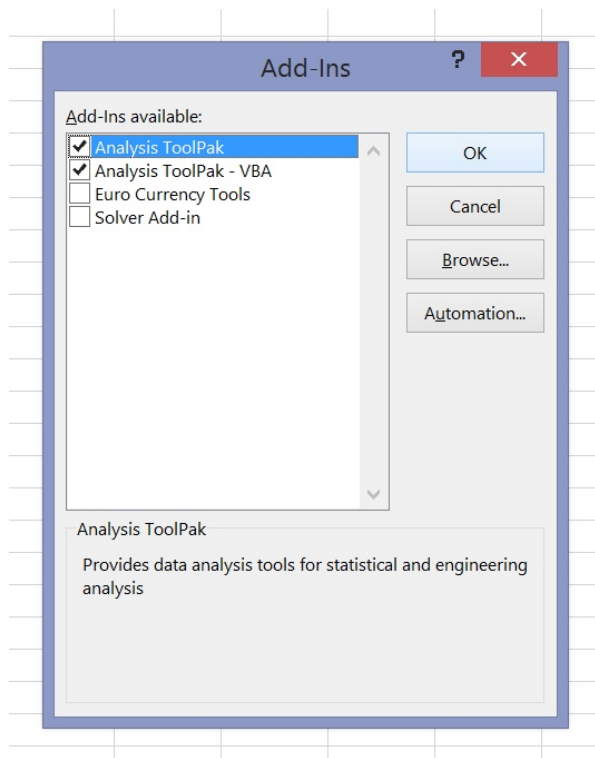


Figure 13.8.19




Figure 13.8.20

Click “Data” then “Data Analysis” and then click “Regression” and “OK”. Congratulations, you have made it to the regression window. The window asks for your inputs. Clicking the box next to the Y and X ranges will allow you to use the click and drag feature of Excel to select your input ranges. Excel has one odd quirk and that is the click and drop feature requires that the independent variables, the X variables, are all together, meaning that they form a single matrix. If your data are set up with the Y variable between two columns of X variables Excel will not allow you to use click and drag. As an example, say Column A and Column C are independent variables and Column B is the Y variable, the dependent variable. Excel will not allow you to click and drop the data ranges. The solution is to move the column with the Y variable to column A and then you can click and drag. The

same problem arises again if you want to run the regression with only some of the X variables. You will need to set up the matrix so all the X variables you wish to regress are in a tightly formed matrix. These steps are presented in the following scene shots.

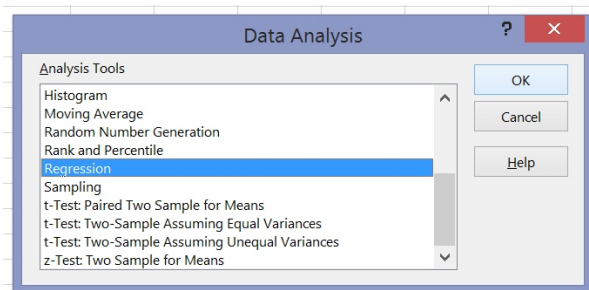


Figure 13.8.21

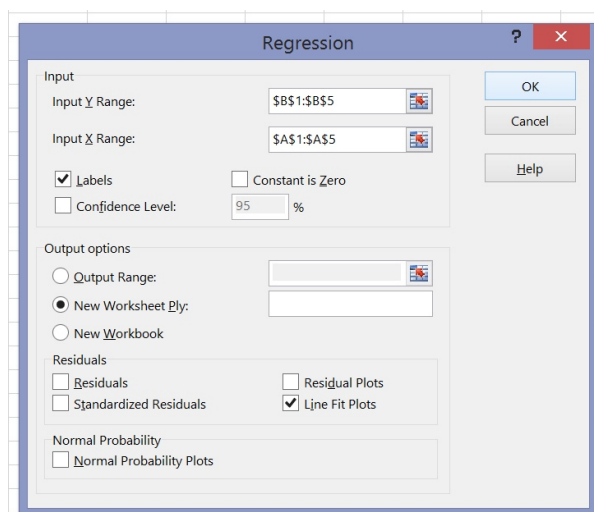


Figure 13.22

Once you have selected the data for your regression analysis and told Excel which one is the dependent variable (Y) and which ones are the independent variables (X 's), you have several choices as to the parameters and how the output will be displayed. Refer to screen shot Figure 13.8.22 under "Input" section. If you check the "labels" box the program will place the entry in the first column of each variable as its name in the output. You can enter an actual name, such as price or income in a demand analysis, in row one of the Excel spreadsheet for each variable and it will be displayed in the output.

The level of significance can also be set by the analyst. This will not change the calculated t statistic, called t stat, but will alter the p value for the calculated t statistic. It will also alter the boundaries of the confidence intervals for the coefficients. A 95 percent confidence interval is always presented, but with a change in this you will also get other levels of confidence for the intervals.

Excel also will allow you to suppress the intercept. This forces the regression program to minimize the residual sum of squares under the condition that the estimated line must go through the origin. This is done in cases where there is no meaning in the model at some value other than zero, zero for the start of the line. An example is an economic production function that is a relationship between the number of units of an input, say hours of labor, and output. There is no meaning of positive output with zero workers.

Once the data are entered and the choices are made click OK and the results will be sent to a separate new worksheet by default. The output from Excel is presented in a way typical of other regression package programs. The first block of information gives the overall statistics of the regression: Multiple R , R Squared, and the R squared adjusted for degrees of freedom, which is the one you want to report. You also get the Standard error (of the estimate) and the number of observations in the regression.

The second block of information is titled ANOVA which stands for Analysis of Variance. Our interest in this section is the column marked F . This is the calculated F statistics for the null hypothesis that all of the coefficients are equal to zero versus the alternative that at least one of the coefficients are not equal to zero. This hypothesis test was presented in 13.4 under "How Good is the Equation?" The next column gives the p value for this test under the title "Significance F". If the p value is less than say 0.05 (the

calculated F statistic is in the tail) we can say with 90 % confidence that we cannot accept the null hypotheses that all the coefficients are equal to zero. This is a good thing; it means that at least one of the coefficients is significantly different from zero thus do have an effect on the value of Y .

The last block of information contains the hypothesis tests for the individual coefficient. The estimated coefficients, the intercept and the slopes, are first listed and then each standard error (of the estimated coefficient) followed by the t stat (calculated student's t statistic for the null hypothesis that the coefficient is equal to zero). We compare the t stat and the critical value of the student's t , dependent on the degrees of freedom, and determine if we have enough evidence to reject the null that the variable has no effect on Y . Remember that we have set up the null hypothesis as the status quo and our claim that we know what caused the Y to change is in the alternative hypothesis. We want to reject the status quo and substitute our version of the world, the alternative hypothesis. The next column contains the p values for this hypothesis test followed by the estimated upper and lower bound of the confidence interval of the estimated slope parameter for various levels of confidence set by us at the beginning.

Estimating the Demand for Roses

Here is an example of using the Excel program to run a regression for a particular specific case: estimating the demand for roses. We are trying to estimate a demand curve, which from economic theory we expect certain variables affect how much of a good we buy. The relationship between the price of a good and the quantity demanded is the demand curve. Beyond that we have the demand function that includes other relevant variables: a person's income, the price of substitute goods, and perhaps other variables such as season of the year or the price of complimentary goods. Quantity demanded will be our Y variable, and Price of roses, Price of carnations and Income will be our independent variables, the X variables.

For all of these variables theory tells us the expected relationship. For the price of the good in question, roses, theory predicts an inverse relationship, the negatively sloped demand curve. Theory also predicts the relationship between the quantity demanded of one good, here roses, and the price of a substitute, carnations in this example. Theory predicts that this should be a positive or direct relationship; as the price of the substitute falls we substitute away from roses to the cheaper substitute, carnations. A reduction in the price of the substitute generates a reduction in demand for the good being analyzed, roses here. Reduction generates reduction is a positive relationship. For normal goods, theory also predicts a positive relationship; as our incomes rise we buy more of the good, roses. We expect these results because that is what is predicted by a hundred years of economic theory and research. Essentially we are testing these century-old hypotheses. The data gathered was determined by the model that is being tested. This should always be the case. One is not doing inferential statistics by throwing a mountain of data into a computer and asking the machine for a theory. Theory first, test follows.

These data here are national average prices and income is the nation's per capita personal income. Quantity demanded is total national annual sales of roses. These are annual time series data; we are tracking the rose market for the United States from 1984-2017, 33 observations.

Because of the quirky way Excel requires how the data are entered into the regression package it is best to have the independent variables, price of roses, price of carnations and income next to each other on the spreadsheet. Once your data are entered into the spreadsheet it is always good to look at the data. Examine the range, the means and the standard deviations. Use your understanding of descriptive statistics from the very first part of this course. In large data sets you will not be able to "scan" the data. The Analysis ToolPac makes it easy to get the range, mean, standard deviations and other parameters of the distributions. You can also quickly get the correlations among the variables. Examine for outliers. Review the history. Did something happen? Was here a labor strike, change in import fees, something that makes these observations unusual? Do not take the data without question. There may have been a typo somewhere, who knows without review.

Go to the regression window, enter the data and select 95% confidence level and click "OK". You can include the labels in the input range if you have put a title at the top of each column, but be sure to click the "labels" box on the main regression page if you do.

The regression output should show up automatically on a new worksheet.

SUMMARY OUTPUT						
Regression Statistics						
Multiple R	0.8560327					
R Square	0.732792					
Adjusted R Square	0.699391					
Standard Error	3629.3427					
Observations	33					
ANOVA						
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>	
Regression	3	577972629.2	2.89E+08	21.9392274	2.59893E-05	
Residual	29	210754050.4	13172128			
Total	32	788726679.5				
	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>
Intercept	183475.43	16791.81835	10.92648	7.89854E-09	147878.367	219072.5
Price of Roses	-1.7607	0.2982	-5.9043	5.20E-05	-2.4049	-1.1164
Price of Carnations	1.3397	0.5273	2.5407	0.0246	0.208	2.4789
Income (per capita)	3.0338	1.2308	2.464901	0.00886322	0.621432	5.4446

Figure 13.8.23

The first results presented is the R-Square, a measure of the strength of the correlation between Y and X_1 , X_2 , and X_3 taken as a group. Our R-square here of 0.699, adjusted for degrees of freedom, means that 70% of the variation in Y , demand for roses, can be explained by variations in X_1 , X_2 , and X_3 , Price of roses, Price of carnations and Income. There is no statistical test to determine the “significance” of an R^2 . Of course a higher R^2 is preferred, but it is really the significance of the coefficients that will determine the value of the theory being tested and which will become part of any policy discussion if they are demonstrated to be significantly different from zero.

Looking at the third panel of output we can write the equation as:

$$Y = b_0 + b_1 X_1 + b_2 X_2 + b_3 X_3 + e$$

where b_0 is the intercept, b_1 is the estimated coefficient on price of roses, and b_2 is the estimated coefficient on price of carnations, b_3 is the estimated effect of income and e is the error term. The equation is written in Roman letters indicating that these are the estimated values and not the population parameters, β 's.

Our estimated equation is:

$$\text{Quantity of roses sold} = 183,475 - 1.76 \text{ Price of roses} + 1.33 \text{ Price of carnations} + 3.03 \text{ Income}$$

We first observe that the signs of the coefficients are as expected from theory. The demand curve is downward sloping with the negative sign for the price of roses. Further the signs of both the price of carnations and income coefficients are positive as would be expected from economic theory.

Interpreting the coefficients can tell us the magnitude of the impact of a change in each variable on the demand for roses. It is the ability to do this which makes regression analysis such a valuable tool. The estimated coefficients tell us that an increase the price of roses by one dollar will lead to a 1.76 reduction in the number roses purchased. The price of carnations seems to play an important role in the demand for roses as we see that increasing the price of carnations by one dollar would increase the demand for roses by 1.33 units as consumers would substitute away from the now more expensive carnations. Similarly, increasing per capita income by one dollar will lead to a 3.03 unit increase in roses purchased.

These results are in line with the predictions of economics theory with respect to all three variables included in this estimate of the demand for roses. It is important to have a theory first that predicts the significance or at least the direction of the coefficients. Without a theory to test, this research tool is not much more helpful than the correlation coefficients we learned about earlier.

We cannot stop there, however. We need to first check whether our coefficients are statistically significant from zero. We set up a hypothesis of:

$$H_0 : \beta_1 = 0$$

$$H_a : \beta_1 \neq 0$$

for all three coefficients in the regression. Recall from earlier that we will not be able to definitively say that our estimated b_1 is the actual real population of β_1 , but rather only that with $(1 - \alpha)\%$ level of confidence that we cannot reject the null hypothesis that our estimated β_1 is significantly different from zero. The analyst is making a claim that the price of roses causes an impact on quantity demanded. Indeed, that each of the included variables has an impact on the quantity of roses demanded. The claim is therefore in the alternative hypotheses. It will take a very large probability, 0.95 in this case, to overthrow the null hypothesis, the

status quo, that $\beta = 0$. In all regression hypothesis tests the claim is in the alternative and the claim is that the theory has found a variable that has a significant impact on the Y variable.

The test statistic for this hypothesis follows the familiar standardizing formula which counts the number of standard deviations, t , that the estimated value of the parameter, b_1 , is away from the hypothesized value, β_0 , which is zero in this case:

$$t_c = \frac{b_1 - \beta_0}{S_{b_1}}$$

The computer calculates this test statistic and presents it as “t stat”. You can find this value to the right of the standard error of the coefficient estimate. The standard error of the coefficient for b_1 is S_{b_1} in the formula. To reach a conclusion we compare this test statistic with the critical value of the student’s t at degrees of freedom $n - 3 - 1 = 29$, and alpha = 0.025 (5% significance level for a two-tailed test). Our t stat for b_1 is approximately 5.90 which is greater than 1.96 (the critical value we looked up in the t -table), so we cannot accept our null hypotheses of no effect. We conclude that Price has a significant effect because the calculated t value is in the tail. We conduct the same test for b_2 and b_3 . For each variable, we find that we cannot accept the null hypothesis of no relationship because the calculated t -statistics are in the tail for each case, that is, greater than the critical value. All variables in this regression have been determined to have a significant effect on the demand for roses.

These tests tell us whether or not an individual coefficient is significantly different from zero, but does not address the overall quality of the model. We have seen that the R squared adjusted for degrees of freedom indicates this model with these three variables explains 70% of the variation in quantity of roses demanded. We can also conduct a second test of the model taken as a whole. This is the F test presented in section 13.4 of this chapter. Because this is a multiple regression (more than one X), we use the F -test to determine if our coefficients collectively affect Y . The hypothesis is:

$$H_0 : \beta_1 = \beta_2 = \dots = \beta_i = 0$$

$$H_a : \text{” at least one of the } \beta_i \text{ is not equal to 0 ”}$$

Under the ANOVA section of the output we find the calculated F statistic for this hypotheses. For this example the F statistic is 21.9. Again, comparing the calculated F statistic with the critical value given our desired level of significance and the degrees of freedom will allow us to reach a conclusion.

The best way to reach a conclusion for this statistical test is to use the p -value comparison rule. The p -value is the area in the tail, given the calculated F statistic. In essence the computer is finding the F value in the table for us and calculating the p -value. In the Summary Output under “significance F” is this probability. For this example, it is calculated to be 2.6×10^{-5} , or 2.6 then moving the decimal five places to the left. (.000026) This is an almost infinitesimal level of probability and is certainly less than our alpha level of .05 for a 5 percent level of significance.

By not being able to accept the null hypotheses we conclude that this specification of this model has validity because at least one of the estimated coefficients is significantly different from zero. Since F -calculated is greater than F -critical, we cannot accept H_0 , meaning that X_1 , X_2 and X_3 together has a significant effect on Y .

The development of computing machinery and the software useful for academic and business research has made it possible to answer questions that just a few years ago we could not even formulate. Data is available in electronic format and can be moved into place for analysis in ways and at speeds that were unimaginable a decade ago. The sheer magnitude of data sets that can today be used for research and analysis gives us a higher quality of results than in days past. Even with only an Excel spreadsheet we can conduct very high level research. This section gives you the tools to conduct some of this very interesting research with the only limit being your imagination.

This page titled [13.8: How to Use Microsoft Excel® for Regression Analysis](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **13.11: How to Use Microsoft Excel® for Regression Analysis** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.9: Chapter Key Terms

a is the symbol for the Y-Intercept

Sometimes written as b_0 , because when writing the theoretical linear model β_0 is used to represent a coefficient for a population.

b is the symbol for Slope

The word coefficient will be used regularly for the slope, because it is a number that will always be next to the letter “ x .” It will be written as b_1 when a sample is used, and β_1 will be used with a population or when writing the theoretical linear model.

Bivariate

two variables are present in the model where one is the “cause” or independent variable and the other is the “effect” of dependent variable.

Linear

a model that takes data and regresses it into a straight line equation.

Multivariate

a system or model where more than one independent variable is being used to predict an outcome. There can only ever be one dependent variable, but there is no limit to the number of independent variables.

R² – Coefficient of Determination

This is a number between 0 and 1 that represents the percentage variation of the dependent variable that can be explained by the variation in the independent variable. Sometimes calculated by the equation $R^2 = \frac{SSR}{SST}$ where SSR is the “Sum of Squares Regression” and SST is the “Sum of Squares Total.” The appropriate coefficient of determination to be reported should always be adjusted for degrees of freedom first.

Residual or “error”

the value calculated from subtracting $y_0 - \hat{y}_0 = e_0$. The absolute value of a residual measures the vertical distance between the actual value of y and the estimated value of y that appears on the best-fit line.

RR – Correlation Coefficient

A number between -1 and 1 that represents the strength and direction of the relationship between “ X ” and “ Y .” The value for “ r ” will equal 1 or -1 only if all the plotted points form a perfectly straight line.

Sum of Squared Errors (SSE)

the calculated value from adding up all the squared residual terms. The hope is that this value is very small when creating a model.

X – the independent variable

This will sometimes be referred to as the “predictor” variable, because these values were measured in order to determine what possible outcomes could be predicted.

Y – the dependent variable

Also, using the letter “ y ” represents actual values while \hat{y} represents predicted or estimated values. Predicted values will come from plugging in observed “ x ” values into a linear model.

This page titled [13.9: Chapter Key Terms](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

13.10: Chapter Review

13.3 Linear Equations

The most basic type of association is a linear association. This type of relationship can be defined algebraically by the equations used, numerically with actual or predicted data values, or graphically from a plotted curve. (Lines are classified as straight curves.) Algebraically, a linear equation typically takes the form $y = mx + b$, where m and b are constants, x is the independent variable, y is the dependent variable. In a statistical context, a linear equation is written in the form $y = a + bx$, where a and b are the constants. This form is used to help readers distinguish the statistical context from the algebraic context. In the equation $y = a + bx$, the constant b that multiplies the x variable (b is called a coefficient) is called the **slope**. The slope describes the rate of change between the independent and dependent variables; in other words, the slope describes the change that occurs in the dependent variable as the independent variable is changed. In the equation $y = a + bx$, the constant a is called the y-intercept.

The **slope of a line** is a value that describes the rate of change between the independent and dependent variables. The **slope** tells us how the dependent variable (y) changes for every one unit increase in the independent (x) variable, on average. The **y-intercept** is used to describe the dependent variable when the independent variable equals zero. Graphically, the slope is represented by three line types in elementary statistics.

13.4 The Regression Equation

It is hoped that this discussion of regression analysis has demonstrated the tremendous potential value it has as a tool for testing models and helping to better understand the world around us. The regression model has its limitations, especially the requirement that the underlying relationship be approximately linear. To the extent that the true relationship is nonlinear it may be approximated with a linear relationship or nonlinear forms of transformations that can be estimated with linear techniques. Double logarithmic transformation of the data will provide an easy way to test this particular shape of the relationship. A reasonably good quadratic form (the shape of the total cost curve from Microeconomics Principles) can be generated by the equation:

$$Y = a + b_1X + b_2X^2$$

where the values of X are simply squared and put into the equation as a separate variable.

There is much more in the way of econometric "tricks" that can bypass some of the more troublesome assumptions of the general regression model. This statistical technique is so valuable that further study would provide any student significant, statistically significant, dividends.

This page titled [13.10: Chapter Review](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- [13.9: Chapter Review](#) by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.11: Chapter Practice

13.1 The Correlation Coefficient r

1.

In order to have a correlation coefficient between traits A and B , it is necessary to have:

- one group of subjects, some of whom possess characteristics of trait A , the remainder possessing those of trait B
- measures of trait A on one group of subjects and of trait B on another group
- two groups of subjects, one which could be classified as A or not A , the other as B or not B
- two groups of subjects, one which could be classified as A or not A , the other as B or not B

2.

Define the Correlation Coefficient and give a unique example of its use.

3.

If the correlation between age of an auto and money spent for repairs is $+0.90$

- 81% of the variation in the money spent for repairs is explained by the age of the auto
- 81% of money spent for repairs is unexplained by the age of the auto
- 90% of the money spent for repairs is explained by the age of the auto
- none of the above

4.

Suppose that college grade-point average and verbal portion of an IQ test had a correlation of $.40$. What percentage of the variance do these two have in common?

- 20
- 16
- 40
- 80

5.

True or false? If false, explain why: The coefficient of determination can have values between -1 and $+1$.

6.

True or False: Whenever r is calculated on the basis of a sample, the value which we obtain for r is only an estimate of the true correlation coefficient which we would obtain if we calculated it for the entire population.

7.

Under a "scatter diagram" there is a notation that the coefficient of correlation is $.10$. What does this mean?

- plus and minus 10% from the means includes about 68% of the cases
- one-tenth of the variance of one variable is shared with the other variable
- one-tenth of one variable is caused by the other variable
- on a scale from -1 to $+1$, the degree of linear relationship between the two variables is $+0.10$

8.

The correlation coefficient for X and Y is known to be zero. We then can conclude that:

- X and Y have standard distributions
- the variances of X and Y are equal
- there exists no relationship between X and Y
- there exists no linear relationship between X and Y
- none of these

9.

What would you guess the value of the correlation coefficient to be for the pair of variables: "number of man-hours worked" and "number of units of work completed"?

- Approximately 0.9
- Approximately 0.4
- Approximately 0.0
- Approximately -0.4
- Approximately -0.9

10.

In a given group, the correlation between height measured in feet and weight measured in pounds is +.68. Which of the following would alter the value of r ?

- height is expressed centimeters.
- weight is expressed in Kilograms.
- both of the above will affect r .
- neither of the above changes will affect r .

13.2 Testing the Significance of the Correlation Coefficient

11.

Define a t Test of a Regression Coefficient, and give a unique example of its use.

12.

The correlation between scores on a neuroticism test and scores on an anxiety test is high and positive; therefore

- anxiety causes neuroticism
- those who score low on one test tend to score high on the other.
- those who score low on one test tend to score low on the other.
- no prediction from one test to the other can be meaningfully made.

13.3 Linear Equations

13.

True or False? If False, correct it: Suppose a 95% confidence interval for the slope β of the straight line regression of Y on X is given by $-3.5 < \beta < -0.5$. Then a two-sided test of the hypothesis $H_0 : \beta = -1$ would result in rejection of H_0 at the 1% level of significance.

14.

True or False: It is safer to interpret correlation coefficients as measures of association rather than causation because of the possibility of spurious correlation.

15.

We are interested in finding the linear relation between the number of widgets purchased at one time and the cost per widget. The following data has been obtained:

X : Number of widgets purchased – 1, 3, 6, 10, 15

Y : Cost per widget(in dollars) – 55, 52, 46, 32, 25

Suppose the regression line is $\hat{y} = -2.5x + 60$. We compute the average price per widget if 30 are purchased and observe which of the following?

- $\hat{y} = 15$ dollars ; obviously, we are mistaken; the prediction \hat{y} is actually +15 dollars.
- $\hat{y} = 15$ dollars , which seems reasonable judging by the data.
- $\hat{y} = -15$ dollars , which is obvious nonsense. The regression line must be incorrect.
- $\hat{y} = -15$ dollars , which is obvious nonsense. This reminds us that predicting Y outside the range of X values in our data is a very poor practice.

16.

Discuss briefly the distinction between correlation and causality.

17.

True or False: If r is close to + or -1, we shall say there is a strong correlation, with the tacit understanding that we are referring to a linear relationship and nothing else.

13.4 The Regression Equation

18.

Suppose that you have at your disposal the information below for each of 30 drivers. Propose a model (including a very brief indication of symbols used to represent independent variables) to explain how miles per gallon vary from driver to driver on the basis of the factors measured.

Information:

1. miles driven per day
2. weight of car
3. number of cylinders in car
4. average speed
5. miles per gallon
6. number of passengers

19.

Consider a sample least squares regression analysis between a dependent variable (Y) and an independent variable (X). A sample correlation coefficient of -1 (minus one) tells us that

- a. there is no relationship between Y and X in the sample
- b. there is no relationship between Y and X in the population
- c. there is a perfect negative relationship between Y and X in the population
- d. there is a perfect negative relationship between Y and X in the sample.

20.

In correlational analysis, when the points scatter widely about the regression line, this means that the correlation is

- a. negative.
- b. low.
- c. heterogeneous.
- d. between two measures that are unreliable.

13.5 Interpretation of Regression Coefficients: Elasticity and Logarithmic Transformation

21.

In a linear regression, why do we need to be concerned with the range of the independent (X) variable?

22.

Suppose one collected the following information where X is diameter of tree trunk and Y is tree height.

Table 13.11.3

X	Y
4	8
2	4
8	18
6	22
10	30
6	8

Regression equation: $\hat{y}_i = -3.6 + 3.1 \cdot X_i$

What is your estimate of the average height of all trees having a trunk diameter of 7 inches?

23.

The manufacturers of a chemical used in flea collars claim that under standard test conditions each additional unit of the chemical will bring about a reduction of 5 fleas (i.e. where X_j = amount of chemical and $Y_j = B_0 + B_1 \cdot X_j + E_j$, $H_0 : B_1 = -5$

Suppose that a test has been conducted and results from a computer include:

Intercept = 60

Slope = -4

Standard error of the regression coefficient = 1.0

Degrees of Freedom for Error = 2000

95% Confidence Interval for the slope -2.04, -5.96

Is this evidence consistent with the claim that the number of fleas is reduced at a rate of 5 fleas per unit chemical?

13.6 Predicting with a Regression Equation

24.

True or False? If False, correct it: Suppose you are performing a simple linear regression of Y on X and you test the hypothesis that the slope β is zero against a two-sided alternative. You have $n = 25$ observations and your computed test (t) statistic is 2.6. Then your P-value is given by $.01 < P < .02$, which gives borderline significance (i.e. you would reject H_0 at $\alpha = .02$ but fail to reject H_0 at $\alpha = .01$).

25.

An economist is interested in the possible influence of "Miracle Wheat" on the average yield of wheat in a district. To do so he fits a linear regression of average yield per year against year after introduction of "Miracle Wheat" for a ten year period.

The fitted trend line is

$$\hat{y}_j = 80 + 1.5 \cdot X_j$$

(Y_j : Average yield in j year after introduction)

(X_j : j year after introduction).

1. What is the estimated average yield for the fourth year after introduction?
2. Do you want to use this trend line to estimate yield for, say, 20 years after introduction? Why? What would your estimate be?

26.

An interpretation of $r = 0.5$ is that the following part of the Y -variation is associated with which variation in X :

- a. most
- b. half
- c. very little
- d. one quarter
- e. none of these

27.

Which of the following values of r indicates the most accurate prediction of one variable from another?

- a. $r = 1.18$
- b. $r = -.77$
- c. $r = .68$

13.7 How to Use Microsoft Excel® for Regression Analysis

28.

A computer program for multiple regression has been used to fit $\hat{y}_j = b_0 + b_1 \cdot X_{1j} + b_2 \cdot X_{2j} + b_3 \cdot X_{3j}$.

Part of the computer output includes:

Table 13.11.4

i	b_i	S_{b_i}
0	8	1.6
1	2.2	.24
2	-.72	.32
3	0.005	0.002

1. Calculation of confidence interval for b_2 consists of _____ \pm (a student's t value) (_____)
2. The confidence level for this interval is reflected in the value used for _____.
3. The degrees of freedom available for estimating the variance are directly concerned with the value used for _____

29.

An investigator has used a multiple regression program on 20 data points to obtain a regression equation with 3 variables. Part of the computer output is:

Table 13.11.5

Variable	Coefficient	Standard Error of bfb_i
1	0.45	0.21
2	0.80	0.10
3	3.10	0.86

1. 0.80 is an estimate of _____.
2. 0.10 is an estimate of _____.
3. Assuming the responses satisfy the normality assumption, we can be 95% confident that the value of β_2 is in the interval, _____ $\pm [t_{.025} \cdot \text{_____}]$, where $t_{.025}$ is the critical value of the student's t distribution with _____ degrees of freedom.

This page titled [13.11: Chapter Practice](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **13.8: Chapter Practice** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

13.12: Chapter Solution

1.

d

2.

A measure of the degree to which variation of one variable is related to variation in one or more other variables. The most commonly used correlation coefficient indicates the degree to which variation in one variable is described by a straight line relation with another variable.

Suppose that sample information is available on family income and Years of schooling of the head of the household. A correlation coefficient = 0 would indicate no linear association at all between these two variables. A correlation of 1 would indicate perfect linear association (where all variation in family income could be associated with schooling and vice versa).

3.

a. 81% of the variation in the money spent for repairs is explained by the age of the auto

4.

b. 16

5.

The coefficient of determination is r^2 with $0 \leq r^2 \leq 1$, since $-1 \leq r \leq 1$.

6.

True

7.

d. on a scale from -1 to +1, the degree of linear relationship between the two variables is +.10

8.

d. there exists no linear relationship between X and Y

9.

Approximately 0.9

10.

d. neither of the above changes will affect r .

11.

Definition: A t test is obtained by dividing a regression coefficient by its standard error and then comparing the result to critical values for Students' t with Error df . It provides a test of the claim that $\beta_i = 0$ when all other variables have been included in the relevant regression model.

Example: Suppose that 4 variables are suspected of influencing some response. Suppose that the results of fitting $Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \beta_3 X_{3i} + \beta_4 X_{4i} + e_i$ include:

Table 13.12.6

Variable	Regression coefficient	Standard error of regular coefficient
.5	1	-3
.4	2	+2
.02	3	+1
.6	4	-.5

t calculated for variables 1, 2, and 3 would be 5 or larger in absolute value while that for variable 4 would be less than 1. For most significance levels, the hypothesis $\beta_1 = 0$ would be rejected. But, notice that this is for the case when X_2 , X_3 , and X_4 have been

included in the regression. For most significance levels, the hypothesis $\beta_4 = 0$ would be continued (retained) for the case where X_1 , X_2 , and X_3 are in the regression. Often this pattern of results will result in computing another regression involving only X_1 , X_2 , X_3 , and examination of the t ratios produced for that case.

12.

c. those who score low on one test tend to score low on the other.

13.

False. Since $H_0 : \beta = -1$ would not be rejected at $\alpha = 0.05$, it would not be rejected at $\alpha = 0.01$.

14.

True

15.

d

16.

Some variables seem to be related, so that knowing one variable's status allows us to predict the status of the other. This relationship can be measured and is called correlation. However, a high correlation between two variables in no way proves that a cause-and-effect relation exists between them. It is entirely possible that a third factor causes both variables to vary together.

17.

True

18.

$$Y_j = b_0 + b_1 \cdot X_1 + b_2 \cdot X_2 + b_3 \cdot X_3 + b_4 \cdot X_4 + b_5 \cdot X_6 + e_j$$

19.

d. there is a perfect negative relationship between Y and X in the sample.

20.

b. low

21.

The precision of the estimate of the Y variable depends on the range of the independent (X) variable explored. If we explore a very small range of the X variable, we won't be able to make much use of the regression. Also, extrapolation is not recommended.

22.

$$\hat{y} = -3.6 + (3.1 \cdot 7) = 18.1$$

23.

Most simply, since -5 is included in the confidence interval for the slope, we can conclude that the evidence is consistent with the claim at the 95% confidence level.

$$\text{Using a t test: } H_0 : B_1 = -5 \quad H_A : B_1 \neq -5 \quad t_{\text{calculated}} = \frac{-5 - (-4)}{1} = -1 \quad t_{\text{critical}} = -1.96.$$

Since $t_{\text{calc}} < t_{\text{crit}}$ we retain the null hypothesis that $B_1 = -5$.

24.

True.

$$t_{(\text{critical}, df=23, \text{two-tailed}, \alpha=.02)} = \pm 2.5$$

$$t_{\text{critical}, df=23, \text{two-tailed}, \alpha=.01} = \pm 2.8$$

25.

$$1. 80 + 1.5 \cdot 4 = 86$$

2. No. Most business statisticians would not want to extrapolate that far. If someone did, the estimate would be 110, but some other factors probably come into play with 20 years.

26.

d. one quarter

27.

b. $r = -.77$

28.

1. $-.72, .32$
2. the t value
3. the t value

29.

1. The population value for β_2 , the change that occurs in Y with a unit change in X_2 , when the other variables are held constant.
2. The population value for the standard error of the distribution of estimates of β_2 .
3. $.8, .1, 16 = 20 - 4$.

This page titled [13.12: Chapter Solution](#) is shared under a [CC BY 4.0](#) license and was authored, remixed, and/or curated by [OpenStax](#) via [source content](#) that was edited to the style and standards of the LibreTexts platform; a detailed edit history is available upon request.

- **13.10: Chapter Solution** by [OpenStax](#) is licensed [CC BY 4.0](#). Original source: <https://openstax.org/details/books/introductory-business-statistics>.

Index

A

alternative hypothesis

- 9.2: Null and Alternative Hypotheses
- 9.4: Distribution Needed for Hypothesis Testing

average

- 1.2: Definitions of Statistics, Probability, and Key Terms

B

balanced design

- 12.4: The F Distribution and the F-Ratio

bar graph

- 1.3: Data, Sampling, and Variation in Data and Sampling
- 2.2: Display Data

Bernoulli trial

- 4.3: Binomial Distribution

binomial distribution

- 8.4: A Confidence Interval for A Population Proportion

- 9.5: Full Hypothesis Test Examples

binomial probability distribution

- 4.3: Binomial Distribution

bivariate

- 13.2: The Correlation Coefficient r

blinding

- 1.5: Experimental Design and Ethics

C

categorical variables

- 1.2: Definitions of Statistics, Probability, and Key Terms

central limit theorem

- 7.1: Introduction to the Central Limit Theorem

coefficient of determination

- 13.5: The Regression Equation

coefficient of multiple determination

- 13.5: The Regression Equation

Cohen's d

- 10.3: Cohen's Standards for Small, Medium, and Large Effect Sizes

complement

- 3.2: Probability Terminology

conditional probability

- 3.2: Probability Terminology

confidence interval

- 8.1: Introduction to Confidence Intervals
- 8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size
- 8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case
- 13.7: Predicting with a Regression Equation

confidence intervals

- 8.4: A Confidence Interval for A Population Proportion

- 9.1: Introduction to Hypothesis Testing

confidence level

- 8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size

contingency table

- 3.5: Contingency Tables and Probability Trees
- 11.5: Test of Independence

continuous

- 1.3: Data, Sampling, and Variation in Data and Sampling

control group

- 1.5: Experimental Design and Ethics

correlation coefficient

- 13.2: The Correlation Coefficient r

critical values

- 9.4: Distribution Needed for Hypothesis Testing

cumulative distribution function (CDF)

- 5.4: The Exponential Distribution

cumulative relative frequency

- 1.4: Levels of Measurement

D

data

- 1.2: Definitions of Statistics, Probability, and Key Terms

degrees of freedom

- 8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case

degrees of freedom (df)

- 10.2: Comparing Two Independent Population Means

dependent variable

- 1.5: Experimental Design and Ethics

descriptive statistics

- 1.2: Definitions of Statistics, Probability, and Key Terms

discrete

- 1.3: Data, Sampling, and Variation in Data and Sampling

E

Empirical Rule

- 6.2: The Standard Normal Distribution
- 8.1: Introduction to Confidence Intervals

equal standard deviations

- 12.3: One-Way ANOVA

equally likely

- 3.2: Probability Terminology

error bound mean

- 8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size

estimate of the error variance

- 13.5: The Regression Equation

event

- 3.2: Probability Terminology

expected mean

- 13.7: Predicting with a Regression Equation

expected value

- 13.7: Predicting with a Regression Equation

expected values

- 11.4: Goodness-of-Fit Test

experiment

- 3.2: Probability Terminology

experimental unit

- 1.5: Experimental Design and Ethics

explanatory variable

- 1.5: Experimental Design and Ethics

exponential distribution

- 5.4: The Exponential Distribution

F

f distribution

- 12.4: The F Distribution and the F-Ratio

f ratio

- 12.4: The F Distribution and the F-Ratio

fair

- 3.2: Probability Terminology

Finite Population Correction Factor

- 7.5: Finite Population Correction Factor

first moment

- 2.7: Skewness and the Mean, Median, and Mode

first quartile

- 2.3: Measures of the Location of the Data

frequency

- 1.4: Levels of Measurement
- 2.2: Display Data

G

geometric distribution

- 4.4: Geometric Distribution

H

histogram

- 2.2: Display Data

Hypergeometric Distribution

- 4.2: Hypergeometric Distribution

hypergeometric experiment

- 4.5: Poisson Distribution

hypotheses

- 9.2: Null and Alternative Hypotheses

hypothesis test

- 9.5: Full Hypothesis Test Examples

hypothesis testing

- 9.1: Introduction to Hypothesis Testing

I

independent

- 3.3: Independent and Mutually Exclusive Events
- 3.4: Two Basic Rules of Probability

independent groups

- 10.1: Introduction

independent variable

- 1.5: Experimental Design and Ethics

inferential statistics

- 1.2: Definitions of Statistics, Probability, and Key Terms

- 8.1: Introduction to Confidence Intervals

interquartile range

- 2.3: Measures of the Location of the Data

interval scale

- 1.4: Levels of Measurement

L

law of large numbers

- 3.2: Probability Terminology
- 7.3: Using the Central Limit Theorem

level of measurement

- 1.4: Levels of Measurement

Line Graph

- 2.2: Display Data

lurking variables

- 1.5: Experimental Design and Ethics

M

matched pairs

- 10.1: Introduction

mean

- 1.2: Definitions of Statistics, Probability, and Key Terms

- 2.4: Measures of the Center of the Data

mean square

12.4: The F Distribution and the F-Ratio

median

2.3: Measures of the Location of the Data
2.4: Measures of the Center of the Data

mode

2.4: Measures of the Center of the Data

multiple correlation coefficient

13.5: The Regression Equation

Multiplication Rule

3.4: Two Basic Rules of Probability

Multivariate

13.2: The Correlation Coefficient r

Mutually Exclusive

3.3: Independent and Mutually Exclusive Events
3.4: Two Basic Rules of Probability

N

Nominal Scale

1.4: Levels of Measurement

Normal Distribution

8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case
9.4: Distribution Needed for Hypothesis Testing

Null Hypothesis

9.2: Null and Alternative Hypotheses
9.4: Distribution Needed for Hypothesis Testing

Numerical Variables

1.2: Definitions of Statistics, Probability, and Key Terms

O

Observed Values

11.4: Goodness-of-Fit Test

Ordinal Scale

1.4: Levels of Measurement

Outcome

3.2: Probability Terminology

Outlier

2.2: Display Data
2.3: Measures of the Location of the Data

P

Paired Data Set

2.2: Display Data

Parameter

1.2: Definitions of Statistics, Probability, and Key Terms
8.1: Introduction to Confidence Intervals

Pareto Chart

1.3: Data, Sampling, and Variation in Data and Sampling

Pearson

1.2: Definitions of Statistics, Probability, and Key Terms

Percentage Impact

13.6: Interpretation of Regression Coefficients-Elasticity and Logarithmic Transformation

Percentiles

2.3: Measures of the Location of the Data

Pie Chart

1.3: Data, Sampling, and Variation in Data and Sampling

Placebo

1.5: Experimental Design and Ethics

Point Estimate

8.1: Introduction to Confidence Intervals

Poisson Probability Distribution

4.5: Poisson Distribution

Population

1.2: Definitions of Statistics, Probability, and Key Terms

1.3: Data, Sampling, and Variation in Data and Sampling

Population Variance

11.3: Test of a Single Variance

Power of the Test

9.3: Outcomes and the Type I and Type II Errors

Prediction Interval

13.7: Predicting with a Regression Equation

Preset or Preconceived b_{fa}

9.4: Distribution Needed for Hypothesis Testing

Probability

1.2: Definitions of Statistics, Probability, and Key Terms
3.2: Probability Terminology

Probability Density Function

4.1: Introduction to Discrete Random Variables
5.2: Properties of Continuous Probability Density Functions

Probability Distribution Function

4.1: Introduction to Discrete Random Variables

Proportion

1.2: Definitions of Statistics, Probability, and Key Terms

Q

Qualitative Data

1.3: Data, Sampling, and Variation in Data and Sampling

Quantitative Continuous Data

1.3: Data, Sampling, and Variation in Data and Sampling

Quantitative Data

1.3: Data, Sampling, and Variation in Data and Sampling

Quantitative Discrete Data

1.3: Data, Sampling, and Variation in Data and Sampling

Quartiles

2.3: Measures of the Location of the Data

R

Random Assignment

1.5: Experimental Design and Ethics

Random Variable

4.2: Hypergeometric Distribution
10.2: Comparing Two Independent Population Means
10.6: Two Population Means with Known Standard Deviations

Ratio Scale

1.4: Levels of Measurement

Regression Equation

13.5: The Regression Equation

Relative Frequency

1.4: Levels of Measurement
2.2: Display Data

Replacement

3.3: Independent and Mutually Exclusive Events

Representative Sample

1.2: Definitions of Statistics, Probability, and Key Terms

Response Variable

1.5: Experimental Design and Ethics

S

Sample

1.2: Definitions of Statistics, Probability, and Key Terms

Sample Space

3.2: Probability Terminology
3.4: Two Basic Rules of Probability
3.5: Contingency Tables and Probability Trees

Samples

1.3: Data, Sampling, and Variation in Data and Sampling

Sampling

1.2: Definitions of Statistics, Probability, and Key Terms

Second Moment

2.7: Skewness and the Mean, Median, and Mode

Significance Level

9.4: Distribution Needed for Hypothesis Testing

Skew

2.7: Skewness and the Mean, Median, and Mode

Standard Deviation

2.8: Measures of the Spread of the Data
8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case
9.4: Distribution Needed for Hypothesis Testing

Standard Error

10.2: Comparing Two Independent Population Means

Standard Error of the Estimate

13.5: The Regression Equation

Standard Normal Distribution

6.2: The Standard Normal Distribution

Standardizing Formula

6.3: Using the Normal Distribution

Statistic

1.2: Definitions of Statistics, Probability, and Key Terms

Statistics

1.2: Definitions of Statistics, Probability, and Key Terms

Stemplots

2.2: Display Data

Sum of Squared Errors (SSE)

13.5: The Regression Equation

Sum of Squares

12.4: The F Distribution and the F-Ratio

T

Test for Homogeneity

11.6: Test for Homogeneity

Test of a Single Variance

11.3: Test of a Single Variance

Test of Independence

11.5: Test of Independence

Test Statistic

9.4: Distribution Needed for Hypothesis Testing
10.6: Two Population Means with Known Standard Deviations

The Central Limit Theorem

7.2: The Central Limit Theorem for Sample Means

The Standard Deviation

10.6: Two Population Means with Known Standard Deviations

Third Quartile

2.3: Measures of the Location of the Data

Treatments

1.5: Experimental Design and Ethics

Tree Diagram

3.5: Contingency Tables and Probability Trees

type I

9.4: Distribution Needed for Hypothesis Testing

type I error

9.3: Outcomes and the Type I and Type II Errors

type II error

9.3: Outcomes and the Type I and Type II Errors

U

unfair

3.2: Probability Terminology

unit

13.6: Interpretation of Regression Coefficients-Elasticity and Logarithmic Transformation

unit change

13.6: Interpretation of Regression Coefficients-Elasticity and Logarithmic Transformation

units

13.6: Interpretation of Regression Coefficients-Elasticity and Logarithmic Transformation

V

variable

1.2: Definitions of Statistics, Probability, and Key Terms

variance

2.8: Measures of the Spread of the Data

variance between samples

12.4: The F Distribution and the F-Ratio

variance within samples

12.4: The F Distribution and the F-Ratio

variances

12.3: One-Way ANOVA

variation

1.3: Data, Sampling, and Variation in Data and Sampling

Venn diagram

3.6: Venn Diagrams

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: Introduction to Business Statistics - OER - Spring 2023

Webpages: 188

All licenses found:

- [CC BY 4.0](#): 94.7% (178 pages)
- [Undeclared](#): 5.3% (10 pages)

By Page

- [Introduction to Business Statistics - OER - Spring 2023 - Undeclared](#)
 - [Front Matter - CC BY 4.0](#)
 - [Licensing - Undeclared](#)
 - [TitlePage - Undeclared](#)
 - [InfoPage - Undeclared](#)
 - [Table of Contents - Undeclared](#)
 - [Licensing - Undeclared](#)
 - [1: Sampling and Data - CC BY 4.0](#)
 - [1.1: Introduction to Sampling and Data - CC BY 4.0](#)
 - [1.2: Definitions of Statistics, Probability, and Key Terms - CC BY 4.0](#)
 - [1.3: Data, Sampling, and Variation in Data and Sampling - CC BY 4.0](#)
 - [1.4: Levels of Measurement - CC BY 4.0](#)
 - [1.5: Experimental Design and Ethics - CC BY 4.0](#)
 - [1.6: Chapter Key Terms - CC BY 4.0](#)
 - [1.7: Chapter References - CC BY 4.0](#)
 - [1.H: Sampling and Data \(Homework\) - CC BY 4.0](#)
 - [1.R: Sampling and Data \(Review\) - CC BY 4.0](#)
 - [1.S: Sampling and Data \(Solutions\) - CC BY 4.0](#)
 - [2: Descriptive Statistics - CC BY 4.0](#)
 - [2.1: introduction to Descriptive Statistics - CC BY 4.0](#)
 - [2.2: Display Data - CC BY 4.0](#)
 - [2.3: Measures of the Location of the Data - CC BY 4.0](#)
 - [2.4: Measures of the Center of the Data - CC BY 4.0](#)
 - [2.5: Sigma Notation and Calculating the Arithmetic Mean - CC BY 4.0](#)
 - [2.6: Geometric Mean - CC BY 4.0](#)
 - [2.7: Skewness and the Mean, Median, and Mode - CC BY 4.0](#)
 - [2.8: Measures of the Spread of the Data - CC BY 4.0](#)
 - [2.9: Chapter Key Terms - CC BY 4.0](#)
 - [2.10: Chapter Formula Review - CC BY 4.0](#)
 - [2.11: Chapter Homework - CC BY 4.0](#)
 - [2.12: Chapter Practice - CC BY 4.0](#)
 - [2.13: Homework - CC BY 4.0](#)
 - [2.14: Chapter References - CC BY 4.0](#)
 - [2.15: Chapter Homework Solutions - CC BY 4.0](#)
 - [2.R: Descriptive Statistics \(Review\) - CC BY 4.0](#)
 - [3: Probability Topics - CC BY 4.0](#)
 - [3.1: Introduction to Probability - CC BY 4.0](#)
 - [3.2: Probability Terminology - CC BY 4.0](#)
 - [3.3: Independent and Mutually Exclusive Events - CC BY 4.0](#)
 - [3.4: Two Basic Rules of Probability - CC BY 4.0](#)
 - [3.5: Contingency Tables and Probability Trees - CC BY 4.0](#)
 - [3.6: Venn Diagrams - CC BY 4.0](#)
 - [3.7: Chapter Key Terms - CC BY 4.0](#)
 - [3.8: Chapter Review - CC BY 4.0](#)
 - [3.9: Chapter Formula Review - CC BY 4.0](#)
 - [3.10: Chapter Practice - CC BY 4.0](#)
 - [3.11: Chapter More Practice - CC BY 4.0](#)
 - [3.12: Chapter Homework - CC BY 4.0](#)
 - [3.13: Chapter Reference - CC BY 4.0](#)
 - [3.14: Chapter Solution \(Practice + Homework\) - CC BY 4.0](#)
 - [4: Discrete Random Variables - CC BY 4.0](#)
 - [4.1: Introduction to Discrete Random Variables - CC BY 4.0](#)
 - [4.2: Hypergeometric Distribution - CC BY 4.0](#)
 - [4.3: Binomial Distribution - CC BY 4.0](#)
 - [4.4: Geometric Distribution - CC BY 4.0](#)
 - [4.5: Poisson Distribution - CC BY 4.0](#)
 - [4.6: Chapter Key Items - CC BY 4.0](#)
 - [4.7: Chapter Review - CC BY 4.0](#)
 - [4.8: Chapter Formula Review - CC BY 4.0](#)
 - [4.9: Chapter Practice - CC BY 4.0](#)
 - [4.10: Chapter Homework - CC BY 4.0](#)
 - [4.11: Chapter References - CC BY 4.0](#)
 - [4.12: Chapter Solution \(Practice + Homework\) - CC BY 4.0](#)
 - [5: Continuous Random Variables - CC BY 4.0](#)
 - [5.1: Prelude to Continuous Random Variables - CC BY 4.0](#)

- 5.2: Properties of Continuous Probability Density Functions - *CC BY 4.0*
- 5.3: The Uniform Distribution - *CC BY 4.0*
- 5.4: The Exponential Distribution - *CC BY 4.0*
- 5.5: Chapter Key Terms - *CC BY 4.0*
- 5.6: Chapter Review - *CC BY 4.0*
- 5.7: Chapter Formula Review - *CC BY 4.0*
- 5.8: Chapter Practice - *CC BY 4.0*
- 5.9: Chapter Homework - *CC BY 4.0*
- 5.10: Chapter References - *CC BY 4.0*
- 5.11: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 6: The Normal Distribution - *CC BY 4.0*
 - 6.1: Introduction to Normal Distribution - *CC BY 4.0*
 - 6.2: The Standard Normal Distribution - *CC BY 4.0*
 - 6.3: Using the Normal Distribution - *CC BY 4.0*
 - 6.4: Estimating the Binomial with the Normal Distribution - *CC BY 4.0*
 - 6.5: Chapter Key Items - *CC BY 4.0*
 - 6.6: Chapter Review - *CC BY 4.0*
 - 6.7: Chapter Formula Review - *CC BY 4.0*
 - 6.8: Chapter Practice - *CC BY 4.0*
 - 6.9: Chapter Homework - *CC BY 4.0*
 - 6.10: Chapter References - *CC BY 4.0*
 - 6.11: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 7: The Central Limit Theorem - *CC BY 4.0*
 - 7.1: Introduction to the Central Limit Theorem - *CC BY 4.0*
 - 7.2: The Central Limit Theorem for Sample Means - *CC BY 4.0*
 - 7.3: Using the Central Limit Theorem - *CC BY 4.0*
 - 7.4: The Central Limit Theorem for Proportions - *CC BY 4.0*
 - 7.5: Finite Population Correction Factor - *CC BY 4.0*
 - 7.6: Chapter Key Terms - *CC BY 4.0*
 - 7.7: Chapter Review - *CC BY 4.0*
 - 7.8: Chapter Formula Review - *CC BY 4.0*
 - 7.9: Chapter Practice - *CC BY 4.0*
 - 7.10: Chapter Homework - *CC BY 4.0*
 - 7.11: Chapter References - *CC BY 4.0*
 - 7.12: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 8: Confidence Intervals - *CC BY 4.0*
 - 8.1: Introduction to Confidence Intervals - *CC BY 4.0*
 - 8.2: A Confidence Interval for a Population Standard Deviation, Known or Large Sample Size - *CC BY 4.0*
 - 8.3: A Confidence Interval for a Population Standard Deviation Unknown, Small Sample Case - *CC BY 4.0*
 - 8.4: A Confidence Interval for A Population Proportion - *CC BY 4.0*
- 8.5: Calculating the Sample Size n - Continuous and Binary Random Variables - *CC BY 4.0*
- 8.6: Chapter Key Terms - *CC BY 4.0*
- 8.7: Chapter Review - *CC BY 4.0*
- 8.8: Chapter Formula Review - *CC BY 4.0*
- 8.9: Chapter Practice - *CC BY 4.0*
- 8.10: Chapter Homework - *CC BY 4.0*
- 8.11: Chapter References - *CC BY 4.0*
- 8.12: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 9: Hypothesis Testing with One Sample - *CC BY 4.0*
 - 9.1: Introduction to Hypothesis Testing - *CC BY 4.0*
 - 9.2: Null and Alternative Hypotheses - *CC BY 4.0*
 - 9.3: Outcomes and the Type I and Type II Errors - *CC BY 4.0*
 - 9.4: Distribution Needed for Hypothesis Testing - *CC BY 4.0*
 - 9.5: Full Hypothesis Test Examples - *CC BY 4.0*
 - 9.6: Chapter Key Terms - *CC BY 4.0*
 - 9.7: Chapter Review - *CC BY 4.0*
 - 9.8: Chapter Formula Review - *CC BY 4.0*
 - 9.9: Chapter Practice - *CC BY 4.0*
 - 9.10: Chapter Homework - *CC BY 4.0*
 - 9.11: Chapter References - *CC BY 4.0*
 - 9.12: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 10: Hypothesis Testing with Two Samples - *CC BY 4.0*
 - 10.1: Introduction - *CC BY 4.0*
 - 10.2: Comparing Two Independent Population Means - *CC BY 4.0*
 - 10.3: Cohen's Standards for Small, Medium, and Large Effect Sizes - *CC BY 4.0*
 - 10.4: Test for Differences in Means- Assuming Equal Population Variances - *CC BY 4.0*
 - 10.5: Comparing Two Independent Population Proportions - *CC BY 4.0*
 - 10.6: Two Population Means with Known Standard Deviations - *CC BY 4.0*
 - 10.7: Matched or Paired Samples - *CC BY 4.0*
 - 10.8: Chapter Key Terms - *CC BY 4.0*
 - 10.9: Chapter Review - *CC BY 4.0*
 - 10.10: Chapter Formula Review - *CC BY 4.0*
 - 10.11: Chapter Practice - *CC BY 4.0*
 - 10.12: Homework - *CC BY 4.0*
 - 10.13: Chapter Homework - *CC BY 4.0*
 - 10.14: Chapter References - *CC BY 4.0*
 - 10.15: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 11: The Chi-Square Distribution - *CC BY 4.0*
 - 11.1: Prelude to the Chi-Square Distribution - *CC BY 4.0*

- 11.2: Facts About the Chi-Square Distribution - *CC BY 4.0*
- 11.3: Test of a Single Variance - *CC BY 4.0*
- 11.4: Goodness-of-Fit Test - *CC BY 4.0*
- 11.5: Test of Independence - *CC BY 4.0*
- 11.6: Test for Homogeneity - *CC BY 4.0*
- 11.7: Comparison of the Chi-Square Tests - *CC BY 4.0*
- 11.8: Chapter Key Terms - *CC BY 4.0*
- 11.9: Chapter Review - *CC BY 4.0*
- 11.10: Chapter Formula Review - *CC BY 4.0*
- 11.11: Chapter Practice - *CC BY 4.0*
- 11.12: Homework - *CC BY 4.0*
- 11.13: Chapter Homework - *CC BY 4.0*
- 11.14: Chapter References - *CC BY 4.0*
- 11.15: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 12: F Distribution and One-Way ANOVA - *CC BY 4.0*
 - 12.1: Introduction to F Distribution and One-Way ANOVA - *CC BY 4.0*
 - 12.2: Test of Two Variances - *CC BY 4.0*
 - 12.3: One-Way ANOVA - *CC BY 4.0*
 - 12.4: The F Distribution and the F-Ratio - *CC BY 4.0*
 - 12.5: Facts About the F Distribution - *CC BY 4.0*
 - 12.6: Chapter Key Terms - *CC BY 4.0*
 - 12.7: Chapter Review - *CC BY 4.0*
 - 12.8: Chapter Formula Review - *CC BY 4.0*
 - 12.9: Chapter Practice - *CC BY 4.0*
 - 12.10: Chapter Homework - *CC BY 4.0*
 - 12.11: Chapter Reference - *CC BY 4.0*
 - 12.12: Chapter Solution (Practice + Homework) - *CC BY 4.0*
- 13: Linear Regression and Correlation - *CC BY 4.0*
 - 13.1: Introduction to Linear Regression and Correlation - *CC BY 4.0*
 - 13.2: The Correlation Coefficient r - *CC BY 4.0*
 - 13.3: Testing the Significance of the Correlation Coefficient - *CC BY 4.0*
 - 13.4: Linear Equations - *CC BY 4.0*
 - 13.5: The Regression Equation - *CC BY 4.0*
 - 13.6: Interpretation of Regression Coefficients- Elasticity and Logarithmic Transformation - *CC BY 4.0*
 - 13.7: Predicting with a Regression Equation - *CC BY 4.0*
 - 13.8: How to Use Microsoft Excel® for Regression Analysis - *CC BY 4.0*
 - 13.9: Chapter Key Terms - *CC BY 4.0*
 - 13.10: Chapter Review - *CC BY 4.0*
 - 13.11: Chapter Practice - *CC BY 4.0*
 - 13.12: Chapter Solution - *CC BY 4.0*
- Back Matter - *Undeclared*
 - Index - *Undeclared*
 - Glossary - *Undeclared*
 - Detailed Licensing - *Undeclared*